

Internet Engineering Task Force (IETF)
Request for Comments: 7276
Category: Informational
ISSN: 2070-1721

T. Mizrahi
Marvell
N. Sprecher
Nokia Solutions and Networks
E. Bellagamba
Ericsson
Y. Weingarten
June 2014

An Overview of
Operations, Administration, and Maintenance (OAM) Tools

Abstract

Operations, Administration, and Maintenance (OAM) is a general term that refers to a toolset for fault detection and isolation, and for performance measurement. Over the years, various OAM tools have been defined for various layers in the protocol stack.

This document summarizes some of the OAM tools defined in the IETF in the context of IP unicast, MPLS, MPLS Transport Profile (MPLS-TP), pseudowires, and Transparent Interconnection of Lots of Links (TRILL). This document focuses on tools for detecting and isolating failures in networks and for performance monitoring. Control and management aspects of OAM are outside the scope of this document. Network repair functions such as Fast Reroute (FRR) and protection switching, which are often triggered by OAM protocols, are also out of the scope of this document.

The target audience of this document includes network equipment vendors, network operators, and standards development organizations. This document can be used as an index to some of the main OAM tools defined in the IETF. At the end of the document, a list of the OAM toolsets and a list of the OAM functions are presented as a summary.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see 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/rfc7276>.

Copyright Notice

Copyright (c) 2014 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. Background	5
1.2. Target Audience	6
1.3. OAM-Related Work in the IETF	6
1.4. Focusing on the Data Plane	7
2. Terminology	8
2.1. Abbreviations	8
2.2. Terminology Used in OAM Standards	10
2.2.1. General Terms	10
2.2.2. Operations, Administration, and Maintenance	10
2.2.3. Functions, Tools, and Protocols	11
2.2.4. Data Plane, Control Plane, and Management Plane	11
2.2.5. The Players	12
2.2.6. Proactive and On-Demand Activation	13
2.2.7. Connectivity Verification and Continuity Checks	14
2.2.8. Connection-Oriented vs. Connectionless Communication	15
2.2.9. Point-to-Point vs. Point-to-Multipoint Services	16
2.2.10. Failures	16
3. OAM Functions	17
4. OAM Tools in the IETF - A Detailed Description	18
4.1. IP Ping	18
4.2. IP Traceroute	19
4.3. Bidirectional Forwarding Detection (BFD)	20
4.3.1. Overview	20
4.3.2. Terminology	20
4.3.3. BFD Control	20
4.3.4. BFD Echo	21
4.4. MPLS OAM	21
4.4.1. LSP Ping	21
4.4.2. BFD for MPLS	22
4.4.3. OAM for Virtual Private Networks (VPNs) over MPLS ..	23
4.5. MPLS-TP OAM	23
4.5.1. Overview	23
4.5.2. Terminology	24
4.5.3. Generic Associated Channel	25
4.5.4. MPLS-TP OAM Toolset	25
4.5.4.1. Continuity Check and Connectivity Verification	26
4.5.4.2. Route Tracing	26
4.5.4.3. Lock Instruct	27
4.5.4.4. Lock Reporting	27
4.5.4.5. Alarm Reporting	27
4.5.4.6. Remote Defect Indication	27
4.5.4.7. Client Failure Indication	27

4.5.4.8. Performance Monitoring	28
4.5.4.8.1. Packet Loss Measurement (LM) ...	28
4.5.4.8.2. Packet Delay Measurement (DM) ..	28
4.6. Pseudowire OAM	29
4.6.1. Pseudowire OAM Using Virtual Circuit Connectivity Verification (VCCV)	29
4.6.2. Pseudowire OAM Using G-ACh	30
4.6.3. Attachment Circuit - Pseudowire Mapping	30
4.7. OWAMP and TWAMP	31
4.7.1. Overview	31
4.7.2. Control and Test Protocols	32
4.7.3. OWAMP	32
4.7.4. TWAMP	33
4.8. TRILL	33
5. Summary	34
5.1. Summary of OAM Tools	34
5.2. Summary of OAM Functions	37
5.3. Guidance to Network Equipment Vendors	38
6. Security Considerations	38
7. Acknowledgments	39
8. References	39
8.1. Normative References	39
8.2. Informative References	39
Appendix A. List of OAM Documents	46
A.1. List of IETF OAM Documents	46
A.2. List of Selected Non-IETF OAM Documents	50

1. Introduction

"OAM" is a general term that refers to a toolset for detecting, isolating, and reporting failures, and for monitoring network performance.

There are several different interpretations of the "OAM" acronym. This document refers to Operations, Administration, and Maintenance, as recommended in Section 3 of [OAM-Def].

This document summarizes some of the OAM tools defined in the IETF in the context of IP unicast, MPLS, MPLS Transport Profile (MPLS-TP), pseudowires, and TRILL.

This document focuses on tools for detecting and isolating failures and for performance monitoring. Hence, this document focuses on the tools used for monitoring and measuring the data plane; control and management aspects of OAM are outside the scope of this document. Network repair functions such as Fast Reroute (FRR) and protection switching, which are often triggered by OAM protocols, are also out of the scope of this document.

1.1. Background

OAM was originally used in traditional communication technologies such as E1 and T1, evolving into Plesiochronous Digital Hierarchy (PDH) and then later into Synchronous Optical Network / Synchronous Digital Hierarchy (SONET/SDH). ATM was probably the first technology to include inherent OAM support from day one, while in other technologies OAM was typically defined in an ad hoc manner after the technology was already defined and deployed. Packet-based networks were traditionally considered unreliable and best effort. As packet-based networks evolved, they have become the common transport for both data and telephony, replacing traditional transport protocols. Consequently, packet-based networks were expected to provide a similar "carrier grade" experience, and specifically to support more advanced OAM functions, beyond ICMP and router hellos, that were traditionally used for fault detection.

As typical networks have a multi-layer architecture, the set of OAM protocols similarly take a multi-layer structure; each layer has its own OAM protocols. Moreover, OAM can be used at different levels of hierarchy in the network to form a multi-layer OAM solution, as shown in the example in Figure 1.

Figure 1 illustrates a network in which IP traffic between two customer edges is transported over an MPLS provider network. MPLS OAM is used at the provider level for monitoring the connection between the two provider edges, while IP OAM is used at the customer level for monitoring the end-to-end connection between the two customer edges.

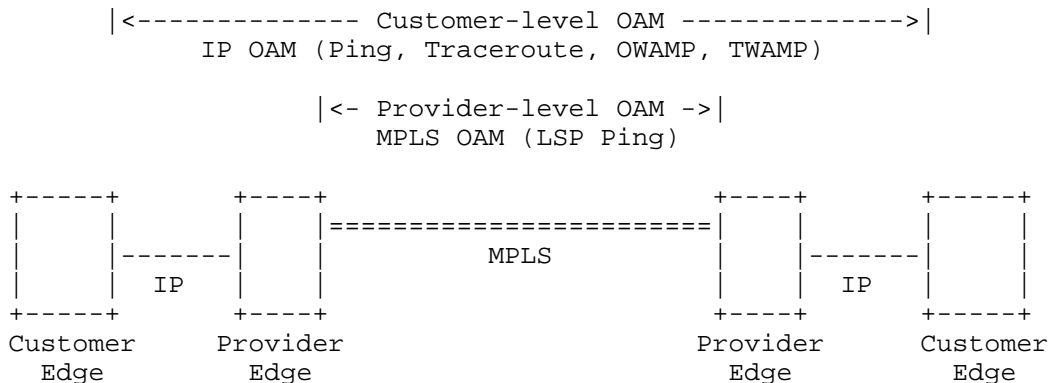


Figure 1: Example of Multi-layer OAM

1.2. Target Audience

The target audience of this document includes:

- o Standards development organizations - Both IETF working groups and non-IETF organizations can benefit from this document when designing new OAM protocols, or when looking to reuse existing OAM tools for new technologies.
- o Network equipment vendors and network operators can use this document as an index to some of the common IETF OAM tools.

It should be noted that some background in OAM is necessary in order to understand and benefit from this document. Specifically, the reader is assumed to be familiar with the term "OAM" [OAM-Def], the motivation for using OAM, and the distinction between OAM and network management [OAM-Mng].

1.3. OAM-Related Work in the IETF

This memo provides an overview of the different sets of OAM tools defined by the IETF. The set of OAM tools described in this memo are applicable to IP unicast, MPLS, pseudowires, MPLS Transport Profile (MPLS-TP), and TRILL. While OAM tools that are applicable to other technologies exist, they are beyond the scope of this memo.

This document focuses on IETF documents that have been published as RFCs, while other ongoing OAM-related work is outside the scope.

The IETF has defined OAM protocols and tools in several different contexts. We roughly categorize these efforts into a few sets of OAM-related RFCs, listed in Table 1. Each set defines a logically coupled set of RFCs, although the sets are in some cases intertwined by common tools and protocols.

The discussion in this document is ordered according to these sets (the acronyms and abbreviations are listed in Section 2.1).

Toolset	Transport Technology
IP Ping	IPv4/IPv6
IP Traceroute	IPv4/IPv6
BFD	generic
MPLS OAM	MPLS
MPLS-TP OAM	MPLS-TP
Pseudowire OAM	Pseudowires
OWAMP and TWAMP	IPv4/IPv6
TRILL OAM	TRILL

Table 1: OAM Toolset Packages in the IETF Documents

This document focuses on OAM tools that have been developed in the IETF. A short summary of some of the significant OAM standards that have been developed in other standard organizations is presented in Appendix A.2.

1.4. Focusing on the Data Plane

OAM tools may, and quite often do, work in conjunction with a control plane and/or management plane. OAM provides instrumentation tools for measuring and monitoring the data plane. OAM tools often use control-plane functions, e.g., to initialize OAM sessions and to exchange various parameters. The OAM tools communicate with the management plane to raise alarms, and often OAM tools may be activated by the management plane (as well as by the control plane), e.g., to locate and localize problems.

The considerations of the control-plane maintenance tools and the functionality of the management plane are out of scope for this document, which concentrates on presenting the data-plane tools that are used for OAM. Network repair functions such as Fast Reroute (FRR) and protection switching, which are often triggered by OAM protocols, are also out of the scope of this document.

Since OAM protocols are used for monitoring the data plane, it is imperative for OAM tools to be capable of testing the actual data plane with as much accuracy as possible. Thus, it is important to enforce fate-sharing between OAM traffic that monitors the data plane and the data-plane traffic it monitors.

2. Terminology

2.1. Abbreviations

ACH	Associated Channel Header
AIS	Alarm Indication Signal
ATM	Asynchronous Transfer Mode
BFD	Bidirectional Forwarding Detection
CC	Continuity Check
CC-V	Continuity Check and Connectivity Verification
CV	Connectivity Verification
DM	Delay Measurement
ECMP	Equal-Cost Multipath
FEC	Forwarding Equivalence Class
FRR	Fast Reroute
G-ACh	Generic Associated Channel
GAL	Generic Associated Channel Label
ICMP	Internet Control Message Protocol
L2TP	Layer 2 Tunneling Protocol
L2VPN	Layer 2 Virtual Private Network
L3VPN	Layer 3 Virtual Private Network
LCCE	L2TP Control Connection Endpoint
LDP	Label Distribution Protocol

LER	Label Edge Router
LM	Loss Measurement
LSP	Label Switched Path
LSR	Label Switching Router
ME	Maintenance Entity
MEG	Maintenance Entity Group
MEP	MEG End Point
MIP	MEG Intermediate Point
MP	Maintenance Point
MPLS	Multiprotocol Label Switching
MPLS-TP	MPLS Transport Profile
MTU	Maximum Transmission Unit
OAM	Operations, Administration, and Maintenance
OWAMP	One-Way Active Measurement Protocol
PDH	Plesiochronous Digital Hierarchy
PE	Provider Edge
PSN	Public Switched Network
PW	Pseudowire
PWE3	Pseudowire Emulation Edge-to-Edge
RBridge	Routing Bridge
RDI	Remote Defect Indication
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
TRILL	Transparent Interconnection of Lots of Links

TTL	Time To Live
TWAMP	Two-Way Active Measurement Protocol
VCCV	Virtual Circuit Connectivity Verification
VPN	Virtual Private Network

2.2. Terminology Used in OAM Standards

2.2.1. General Terms

A wide variety of terms is used in various OAM standards. This section presents a comparison of the terms used in various OAM standards, without fully quoting the definition of each term.

An interesting overview of the term "OAM" and its derivatives is presented in [OAM-Def]. A thesaurus of terminology for MPLS-TP terms is presented in [TP-Term], which provides a good summary of some of the OAM-related terminology.

2.2.2. Operations, Administration, and Maintenance

The following definition of OAM is quoted from [OAM-Def]:

The components of the "OAM" acronym (and provisioning) are defined as follows:

- o Operations - Operation activities are undertaken to keep the network (and the services that the network provides) up and running. It includes monitoring the network and finding problems. Ideally these problems should be found before users are affected.
- o Administration - Administration activities involve keeping track of resources in the network and how they are used. It includes all the bookkeeping that is necessary to track networking resources and the network under control.
- o Maintenance - Maintenance activities are focused on facilitating repairs and upgrades -- for example, when equipment must be replaced, when a router needs a patch for an operating system image, or when a new switch is added to a network. Maintenance also involves corrective and preventive measures to make the managed network run more effectively, e.g., adjusting device configuration and parameters.

2.2.3. Functions, Tools, and Protocols

OAM Function

An OAM function is an instrumentation measurement type or diagnostic.

OAM functions are the atomic building blocks of OAM, where each function defines an OAM capability.

Typical examples of OAM functions are presented in Section 3.

OAM Protocol

An OAM protocol is a protocol used for implementing one or more OAM functions.

The OWAMP-Test [OWAMP] is an example of an OAM protocol.

OAM Tool

An OAM tool is a specific means of applying one or more OAM functions.

In some cases, an OAM protocol *is* an OAM tool, e.g., OWAMP-Test. In other cases, an OAM tool uses a set of protocols that are not strictly OAM related; for example, Traceroute (Section 4.2) can be implemented using UDP and ICMP messages, without using an OAM protocol per se.

2.2.4. Data Plane, Control Plane, and Management Plane

Data Plane

The data plane is the set of functions used to transfer data in the stratum or layer under consideration [ITU-Terms].

The data plane is also known as the forwarding plane or the user plane.

Control Plane

The control plane is the set of protocols and mechanisms that enable routers to efficiently learn how to forward packets towards their final destination (based on [Comp]).

Management Plane

The term "Management Plane", as described in [Mng], is used to describe the exchange of management messages through management protocols (often transported by IP and by IP transport protocols) between management applications and the managed entities such as network nodes.

Data Plane vs. Control Plane vs. Management Plane

The distinction between the planes is at times a bit vague. For example, the definition of "Control Plane" above may imply that OAM tools such as ping, BFD, and others are in fact in the control plane.

This document focuses on tools used for monitoring the data plane. While these tools could arguably be considered to be in the control plane, these tools monitor the data plane, and hence it is imperative to have fate-sharing between OAM traffic that monitors the data plane and the data-plane traffic it monitors.

Another potentially vague distinction is between the management plane and control plane. The management plane should be seen as separate from, but possibly overlapping with, the control plane (based on [Mng]).

2.2.5. The Players

An OAM tool is used between two (or more) peers. Various terms are used in IETF documents to refer to the players that take part in OAM. Table 2 summarizes the terms used in each of the toolsets discussed in this document.

Toolset	Terms
Ping / Traceroute ([ICMPv4], [ICMPv6], [TCPIP-Tools])	- Host - Node - Interface - Gateway
BFD [BFD]	- System
MPLS OAM [MPLS-OAM-FW]	- LSR
MPLS-TP OAM [TP-OAM-FW]	- End Point - MEP - Intermediate Point - MIP
Pseudowire OAM [VCCV]	- PE - LCCE
OWAMP and TWAMP ([OWAMP], [TWAMP])	- Host - End system
TRILL OAM [TRILL-OAM]	- RBridge

Table 2: Maintenance Point Terminology

2.2.6. Proactive and On-Demand Activation

The different OAM tools may be used in one of two basic types of activation:

Proactive

Proactive activation - indicates that the tool is activated on a continual basis, where messages are sent periodically, and errors are detected when a certain number of expected messages are not received.

On-demand

On-demand activation - indicates that the tool is activated "manually" to detect a specific anomaly.

2.2.7. Connectivity Verification and Continuity Checks

Two distinct classes of failure management functions are used in OAM protocols: Connectivity Verification and Continuity Checks. The distinction between these terms is defined in [MPLS-TP-OAM] and is used similarly in this document.

Continuity Check

Continuity Checks are used to verify that a destination is reachable, and are typically sent proactively, though they can be invoked on-demand as well.

Connectivity Verification

A Connectivity Verification function allows Alice to check whether she is connected to Bob or not. It is noted that while the CV function is performed in the data plane, the "expected path" is predetermined in either the control plane or the management plane. A Connectivity Verification (CV) protocol typically uses a CV message, followed by a CV reply that is sent back to the originator. A CV function can be applied proactively or on-demand.

Connectivity Verification tools often perform path verification as well, allowing Alice to verify that messages from Bob are received through the correct path, thereby verifying not only that the two MPs are connected, but also that they are connected through the expected path, allowing detection of unexpected topology changes.

Connectivity Verification functions can also be used for checking the MTU of the path between the two peers.

Connectivity Verification and Continuity Checks are considered complementary mechanisms and are often used in conjunction with each other.

2.2.8. Connection-Oriented vs. Connectionless Communication

Connection-Oriented

In connection-oriented technologies, an end-to-end connection is established (by a control protocol or provisioned by a management system) prior to the transmission of data.

Typically a connection identifier is used to identify the connection. In connection-oriented technologies, it is often the case (although not always) that all packets belonging to a specific connection use the same route through the network.

Connectionless

In connectionless technologies, data is typically sent between end points without prior arrangement. Packets are routed independently based on their destination address, and hence different packets may be routed in a different way across the network.

Discussion

The OAM tools described in this document include tools that support connection-oriented technologies, as well as tools for connectionless technologies.

In connection-oriented technologies, OAM is used to monitor a *specific* connection; OAM packets are forwarded through the same route as the data traffic and receive the same treatment. In connectionless technologies, OAM is used between a source and destination pair without defining a specific connection. Moreover, in some cases, the route of OAM packets may differ from the one of the data traffic. For example, the connectionless IP Ping (Section 4.1) tests the reachability from a source to a given destination, while the connection-oriented LSP Ping (Section 4.4.1) is used for monitoring a specific LSP (connection) and provides the capability to monitor all the available paths used by an LSP.

It should be noted that in some cases connectionless protocols are monitored by connection-oriented OAM protocols. For example, while IP is a connectionless protocol, it can be monitored by BFD (Section 4.3), which is connection oriented.

2.2.9. Point-to-Point vs. Point-to-Multipoint Services

Point-to-point (P2P)

A P2P service delivers data from a single source to a single destination.

Point-to-multipoint (P2MP)

A P2MP service delivers data from a single source to a one or more destinations (based on [Signal]).

An MP2MP service is a service that delivers data from more than one source to one or more receivers (based on [Signal]).

Note: the two definitions for P2MP and MP2MP are quoted from [Signal]. Although [Signal] describes a specific case of P2MP and MP2MP that is MPLS-specific, these two definitions also apply to non-MPLS cases.

Discussion

The OAM tools described in this document include tools for P2P services, as well as tools for P2MP services.

The distinction between P2P services and P2MP services affects the corresponding OAM tools. A P2P service is typically simpler to monitor, as it consists of a single pair of endpoints. P2MP and MP2MP services present several challenges. For example, in a P2MP service, the OAM mechanism not only verifies that each of the destinations is reachable from the source but also verifies that the P2MP distribution tree is intact and loop-free.

2.2.10. Failures

The terms "Failure", "Fault", and "Defect" are used interchangeably in the standards, referring to a malfunction that can be detected by a Connectivity Verification or a Continuity Check. In some standards, such as 802.1ag [IEEE802.1Q], there is no distinction between these terms, while in other standards each of these terms refers to a different type of malfunction.

The terminology used in IETF MPLS-TP OAM is based on the ITU-T terminology, which distinguishes between these three terms in [ITU-T-G.806] as follows:

Fault

The term "Fault" refers to an inability to perform a required action, e.g., an unsuccessful attempt to deliver a packet.

Defect

The term "Defect" refers to an interruption in the normal operation, such as a consecutive period of time where no packets are delivered successfully.

Failure

The term "Failure" refers to the termination of the required function. While a Defect typically refers to a limited period of time, a failure refers to a long period of time.

3. OAM Functions

This subsection provides a brief summary of the common OAM functions used in OAM-related standards. These functions are used as building blocks in the OAM standards described in this document.

- o Connectivity Verification (CV), Path Verification, and Continuity Check (CC):
As defined in Section 2.2.7.
- o Path Discovery / Fault Localization:
This function can be used to trace the route to a destination, i.e., to identify the nodes along the route to the destination. When more than one route is available to a specific destination, this function traces one of the available routes. When a failure occurs, this function attempts to detect the location of the failure.
Note that the term "route tracing" (or "Traceroute"), which is used in the context of IP and MPLS, is sometimes referred to as "path tracing" in the context of other protocols, such as TRILL.

- o Performance Monitoring:
Typically refers to:
 - * Loss Measurement (LM) - monitors the packet loss rate.
 - * Delay Measurement (DM) - monitors the delay and delay variation (jitter).

4. OAM Tools in the IETF - A Detailed Description

This section presents a detailed description of the sets of OAM-related tools in each of the toolsets in Table 1.

4.1. IP Ping

Ping is a common network diagnostic application for IP networks that use ICMP. According to [NetTerms], 'Ping' is an abbreviation for Packet internet groper, although the term has been so commonly used that it stands on its own. As defined in [NetTerms], it is a program used to test reachability of destinations by sending them an ICMP Echo request and waiting for a reply.

The ICMP Echo request/reply exchange in Ping is used as a Continuity Check function for the Internet Protocol. The originator transmits an ICMP Echo request packet, and the receiver replies with an Echo reply. ICMP Ping is defined in two variants: [ICMPv4] is used for IPv4, and [ICMPv6] is used for IPv6.

Ping can be invoked to either a unicast destination or a multicast destination. In the latter case, all members of the multicast group send an Echo reply back to the originator.

Ping implementations typically use ICMP messages. UDP Ping is a variant that uses UDP messages instead of ICMP Echo messages.

Ping is a single-ended Continuity Check, i.e., it allows the *initiator* of the Echo request to test the reachability. If it is desirable for both ends to test the reachability, both ends have to invoke Ping independently.

Note that since ICMP filtering is deployed in some routers and firewalls, the usefulness of Ping is sometimes limited in the wider Internet. This limitation is equally relevant to Traceroute.

4.2. IP Traceroute

Traceroute ([TCPIP-Tools], [NetTools]) is an application that allows users to discover a path between an IP source and an IP destination.

The most common way to implement Traceroute [TCPIP-Tools] is described as follows. Traceroute sends a sequence of UDP packets to UDP port 33434 at the destination. By default, Traceroute begins by sending three packets (the number of packets is configurable in most Traceroute implementations), each with an IP Time-To-Live (or Hop Limit in IPv6) value of one, to the destination. These packets expire as soon as they reach the first router in the path. Consequently, that router sends three ICMP Time Exceeded Messages back to the Traceroute application. Traceroute now sends another three UDP packets, each with the TTL value of 2. These messages cause the second router to return ICMP messages. This process continues, with ever-increasing values for the TTL field, until the packets actually reach the destination. Because no application listens to port 33434 at the destination, the destination returns ICMP Destination Unreachable Messages indicating an unreachable port. This event indicates to the Traceroute application that it is finished. The Traceroute program displays the round-trip delay associated with each of the attempts.

While Traceroute is a tool that finds *a* path from A to B, it should be noted that traffic from A to B is often forwarded through Equal-Cost Multipaths (ECMPs). Paris Traceroute [PARIS] is an extension to Traceroute that attempts to discover all the available paths from A to B by scanning different values of header fields (such as UDP ports) in the probe packets.

It is noted that Traceroute is an application, and not a protocol. As such, it has various different implementations. One of the most common ones uses UDP probe packets, as described above. Other implementations exist that use other types of probe messages, such as ICMP or TCP.

Note that IP routing may be asymmetric. While Traceroute discovers a path between a source and destination, it does not reveal the reverse path.

A few ICMP extensions ([ICMP-MP], [ICMP-Int]) have been defined in the context of Traceroute. These documents define several extensions, including extensions to the ICMP Destination Unreachable message, that can be used by Traceroute applications.

Traceroute allows path discovery to *unicast* destination addresses. A similar tool [mtrace] was defined for multicast destination addresses; it allows tracing the route that a multicast IP packet takes from a source to a particular receiver.

4.3. Bidirectional Forwarding Detection (BFD)

4.3.1. Overview

While multiple OAM tools have been defined for various protocols in the protocol stack, Bidirectional Forwarding Detection [BFD], defined by the IETF BFD working group, is a generic OAM tool that can be deployed over various encapsulating protocols, and in various medium types. The IETF has defined variants of the protocol for IP ([BFD-IP], [BFD-Multi]), for MPLS LSPs [BFD-LSP], and for pseudowires [BFD-VCCV]. The usage of BFD in MPLS-TP is defined in [TP-CC-CV].

BFD includes two main OAM functions, using two types of BFD packets: BFD Control packets and BFD Echo packets.

4.3.2. Terminology

BFD operates between *systems*. The BFD protocol is run between two or more systems after establishing a *session*.

4.3.3. BFD Control

BFD supports a bidirectional Continuity Check, using BFD Control packets that are exchanged within a BFD session. BFD sessions operate in one of two modes:

- o Asynchronous mode (i.e., proactive): in this mode, BFD Control packets are sent periodically. When the receiver detects that no BFD Control packets have been received during a predetermined period of time, a failure is reported.
- o Demand mode: in this mode, BFD Control packets are sent on demand. Upon need, a system initiates a series of BFD Control packets to check the continuity of the session. BFD Control packets are sent independently in each direction.

Each of the endpoints (referred to as systems) of the monitored path maintains its own session identification, called a Discriminator; both Discriminators are included in the BFD Control Packets that are exchanged between the endpoints. At the time of session establishment, the Discriminators are exchanged between the two endpoints. In addition, the transmission (and reception) rate is

negotiated between the two endpoints, based on information included in the control packets. These transmission rates may be renegotiated during the session.

During normal operation of the session, i.e., when no failures have been detected, the BFD session is in the Up state. If no BFD Control packets are received during a period of time called the Detection Time, the session is declared to be Down. The detection time is a function of the pre-configured or negotiated transmission rate and a parameter called Detect Mult. Detect Mult determines the number of missing BFD Control packets that cause the session to be declared as Down. This parameter is included in the BFD Control packet.

4.3.4. BFD Echo

A BFD Echo packet is sent to a peer system and is looped back to the originator. The echo function can be used proactively or on demand.

The BFD Echo function has been defined in BFD for IPv4 and IPv6 ([BFD-IP]), but it is not used in BFD for MPLS LSPs or PWS, or in BFD for MPLS-TP.

4.4. MPLS OAM

The IETF MPLS working group has defined OAM for MPLS LSPs. The requirements and framework of this effort are defined in [MPLS-OAM-FW] and [MPLS-OAM], respectively. The corresponding OAM tool defined, in this context, is LSP Ping [LSP-Ping]. OAM for P2MP services is defined in [MPLS-P2MP].

BFD for MPLS [BFD-LSP] is an alternative means for detecting data-plane failures, as described below.

4.4.1. LSP Ping

LSP Ping is modeled after the Ping/Traceroute paradigm, and thus it may be used in one of two modes:

- o "Ping" mode: In this mode, LSP Ping is used for end-to-end Connectivity Verification between two LERs.
- o "Traceroute" mode: This mode is used for hop-by-hop fault isolation.

LSP Ping is based on the ICMP Ping operation (of data-plane Connectivity Verification) with additional functionality to verify data-plane vs. control-plane consistency for a Forwarding Equivalence Class (FEC) and also to identify Maximum Transmission Unit (MTU) problems.

The Traceroute functionality may be used to isolate and localize MPLS faults, using the Time-To-Live (TTL) indicator to incrementally identify the sub-path of the LSP that is successfully traversed before the faulty link or node.

The challenge in MPLS networks is that the traffic of a given LSP may be load-balanced across Equal-Cost Multipaths (ECMPs). LSP Ping monitors all the available paths of an LSP by monitoring its different FECs. Note that MPLS-TP does not use ECMP, and thus does not require OAM over multiple paths.

Another challenge is that an MPLS LSP does not necessarily have a return path; traffic that is sent back from the egress LSR to the ingress LSR is not necessarily sent over an MPLS LSP, but it can be sent through a different route, such as an IP route. Thus, responding to an LSP Ping message is not necessarily as trivial as in IP Ping, where the responder just swaps the source and destination IP addresses. Note that this challenge is not applicable to MPLS-TP, where a return path is always available.

It should be noted that LSP Ping supports unique identification of the LSP within an addressing domain. The identification is checked using the full FEC identification. LSP Ping is extensible to include additional information needed to support new functionality, by use of Type-Length-Value (TLV) constructs. The usage of TLVs is typically handled by the control plane, as it is not easy to implement in hardware.

LSP Ping supports both asynchronous and on-demand activation.

4.4.2. BFD for MPLS

BFD [BFD-LSP] can be used to detect MPLS LSP data-plane failures.

A BFD session is established for each MPLS LSP that is being monitored. BFD Control packets must be sent along the same path as the monitored LSP. If the LSP is associated with multiple FECs, a BFD session is established for each FEC.

While LSP Ping can be used for detecting MPLS data-plane failures and for verifying the MPLS LSP data plane against the control plane, BFD can only be used for the former. BFD can be used in conjunction with LSP Ping, as is the case in MPLS-TP (see Section 4.5.4).

4.4.3. OAM for Virtual Private Networks (VPNs) over MPLS

The IETF has defined two classes of VPNs: Layer 2 VPNs (L2VPNs) and Layer 3 VPNs (L3VPNs). [L2VPN-OAM] provides the requirements and framework for OAM in the context of L2VPNs, and specifically it also defines the OAM layering of L2VPNs over MPLS. [L3VPN-OAM] provides a framework for the operation and management of L3VPNs.

4.5. MPLS-TP OAM

4.5.1. Overview

The MPLS working group has defined the OAM toolset that fulfills the requirements for MPLS-TP OAM. The full set of requirements for MPLS-TP OAM are defined in [MPLS-TP-OAM] and include both general requirements for the behavior of the OAM tools and a set of operations that should be supported by the OAM toolset. The set of mechanisms required are further elaborated in [TP-OAM-FW], which describes the general architecture of the OAM system and also gives overviews of the functionality of the OAM toolset.

Some of the basic requirements for the OAM toolset for MPLS-TP are:

- o MPLS-TP OAM must be able to support both an IP-based environment and a non-IP-based environment. If the network is IP based, i.e., IP routing and forwarding are available, then the MPLS-TP OAM toolset should rely on the IP routing and forwarding capabilities. On the other hand, in environments where IP functionality is not available, the OAM tools must still be able to operate without dependence on IP forwarding and routing.
- o OAM packets and the user traffic are required to be congruent (i.e., OAM packets are transmitted in-band), and there is a need to differentiate OAM packets from ordinary user packets in the data plane. Inherent in this requirement is the principle that MPLS-TP OAM be independent of any existing control plane, although it should not preclude use of the control-plane functionality. OAM packets are identified by the Generic Associated Channel Label (GAL), which is a reserved MPLS label value (13).

4.5.2. Terminology

Maintenance Entity (ME)

The MPLS-TP OAM tools are designed to monitor and manage a Maintenance Entity (ME). An ME, as defined in [TP-OAM-FW], defines a relationship between two points of a transport path to which maintenance and monitoring operations apply.

The term "Maintenance Entity (ME)" is used in ITU-T Recommendations (e.g., [ITU-T-Y1731]), as well as in the MPLS-TP terminology ([TP-OAM-FW]).

Maintenance Entity Group (MEG)

The collection of one or more MEs that belong to the same transport path and that are maintained and monitored as a group are known as a Maintenance Entity Group (based on [TP-OAM-FW]).

Maintenance Point (MP)

A Maintenance Point (MP) is a functional entity that is defined at a node in the network and can initiate and/or react to OAM messages. This document focuses on the data-plane functionality of MPs, while MPs interact with the control plane and with the management plane as well.

The term "MP" is used in IEEE 802.lag and was similarly adopted in MPLS-TP ([TP-OAM-FW]).

MEG End Point (MEP)

A MEG End Point (MEP) is one of the endpoints of an ME, and can initiate OAM messages and respond to them (based on [TP-OAM-FW]).

MEG Intermediate Point (MIP)

In between MEPs, there are zero or more intermediate points, called MEG Intermediate Points (based on [TP-OAM-FW]).

A MEG Intermediate Point (MIP) is an intermediate point that does not generally initiate OAM frames (one exception to this is the use of AIS notifications) but is able to respond to OAM frames that are destined to it. A MIP in MPLS-TP identifies OAM packets destined to it by the expiration of the TTL field in the OAM packet. The term "Maintenance Point" is a general term for MEPs and MIPs.

Up and Down MEPs

IEEE 802.1ag [IEEE802.1Q] defines a distinction between Up MEPs and Down MEPs. A MEP monitors traffic in either the direction facing the network or the direction facing the bridge. A Down MEP is a MEP that receives OAM packets from and transmits them to the direction of the network. An Up MEP receives OAM packets from and transmits them to the direction of the bridging entity. MPLS-TP ([TP-OAM-FW]) uses a similar distinction on the placement of the MEP -- at either the ingress, egress, or forwarding function of the node (Down / Up MEPs). This placement is important for localization of a failure.

Note that the terms "Up MEP" and "Down MEP" are entirely unrelated to the conventional "Up"/"Down" terminology, where "Down" means faulty and "Up" means not faulty.

The distinction between Up and Down MEPs was defined in [TP-OAM-FW], but has not been used in other MPLS-TP RFCs, as of the writing of this document.

4.5.3. Generic Associated Channel

In order to address the requirement for in-band transmission of MPLS-TP OAM traffic, MPLS-TP uses a Generic Associated Channel (G-ACh), defined in [G-ACh] for LSP-based OAM traffic. This mechanism is based on the same concepts as the PWE3 ACH [PW-ACH] and VCCV [VCCV] mechanisms. However, to address the needs of LSPs as differentiated from PW, the following concepts were defined for [G-ACh]:

- o An Associated Channel Header (ACH), which uses a format similar to the PW Control Word [PW-ACH], is a 4-byte header that is prepended to OAM packets.
- o A Generic Associated Channel Label (GAL). The GAL is a reserved MPLS label value (13) that indicates that the packet is an ACH packet and the payload follows immediately after the label stack.

It should be noted that while the G-ACh was defined as part of the MPLS-TP definition effort, the G-ACh is a generic tool that can be used in MPLS in general, and not only in MPLS-TP.

4.5.4. MPLS-TP OAM Toolset

To address the functionality that is required of the OAM toolset, the MPLS WG conducted an analysis of the existing IETF and ITU-T OAM tools and their ability to fulfill the required functionality. The

conclusions of this analysis are documented in [OAM-Analys]. MPLS-TP uses a mixture of OAM tools that are based on previous standards and adapted to the requirements of [MPLS-TP-OAM]. Some of the main building blocks of this solution are based on:

- o Bidirectional Forwarding Detection ([BFD], [BFD-LSP]) for proactive Continuity Check and Connectivity Verification.
- o LSP Ping as defined in [LSP-Ping] for on-demand Connectivity Verification.
- o New protocol packets, using G-ACH, to address different functionality.
- o Performance measurement protocols.

The following subsections describe the OAM tools defined for MPLS-TP as described in [TP-OAM-FW].

4.5.4.1. Continuity Check and Connectivity Verification

Continuity Checks and Connectivity Verification are presented in Section 2.2.7 of this document. As presented there, these tools may be used either proactively or on demand. When using these tools proactively, they are generally used in tandem.

For MPLS-TP there are two distinct tools: the proactive tool is defined in [TP-CC-CV], while the on-demand tool is defined in [OnDemand-CV]. In on-demand mode, this function should support monitoring between the MEPs and, in addition, between a MEP and MIP. [TP-OAM-FW] highlights, when performing Connectivity Verification, the need for the CC-V messages to include unique identification of the MEG that is being monitored and the MEP that originated the message.

The proactive tool [TP-CC-CV] is based on extensions to BFD (see Section 4.3) with the additional limitation that the transmission and receiving rates are based on configuration by the operator. The on-demand tool [OnDemand-CV] is an adaptation of LSP Ping (see Section 4.4.1) for the required behavior of MPLS-TP.

4.5.4.2. Route Tracing

[MPLS-TP-OAM] defines that there is a need for functionality that would allow a path endpoint to identify the intermediate and endpoints of the path. This function would be used in on-demand mode. Normally, this path will be used for bidirectional PW, LSP,

and Sections; however, unidirectional paths may be supported only if a return path exists. The tool for this is based on the LSP Ping (see Section 4.4.1) functionality and is described in [OnDemand-CV].

4.5.4.3. Lock Instruct

The Lock Instruct function [Lock-Loop] is used to notify a transport-path endpoint of an administrative need to disable the transport path. This functionality will generally be used in conjunction with some intrusive OAM function, e.g., performance measurement or diagnostic testing, to minimize the side-effect on user data traffic.

4.5.4.4. Lock Reporting

Lock Reporting is a function used by an endpoint of a path to report to its far-end endpoint that a lock condition has been affected on the path.

4.5.4.5. Alarm Reporting

Alarm reporting [TP-Fault] provides the means to suppress alarms following detection of defect conditions at the server sub-layer. Alarm reporting is used by an intermediate point of a path, that becomes aware of a fault on the path, to report to the endpoints of the path. [TP-OAM-FW] states that this may occur as a result of a defect condition discovered at a server sub-layer. This generates an Alarm Indication Signal (AIS) that continues until the fault is cleared. The consequent action of this function is detailed in [TP-OAM-FW].

4.5.4.6. Remote Defect Indication

Remote Defect Indication (RDI) is used proactively by a path endpoint to report to its peer endpoint that a defect is detected on a bidirectional connection between them. [MPLS-TP-OAM] points out that this function may be applied to a unidirectional LSP only if a return path exists. [TP-OAM-FW] points out that this function is associated with the proactive CC-V function.

4.5.4.7. Client Failure Indication

Client Failure Indication (CFI) is defined in [MPLS-TP-OAM] to allow the propagation information from one edge of the network to the other. The information concerns a defect to a client, in the case that the client does not support alarm notification.

4.5.4.8. Performance Monitoring

The definition of MPLS performance monitoring was motivated by the MPLS-TP requirements [MPLS-TP-OAM] but was defined generically for MPLS in [MPLS-LM-DM]. An additional document [TP-LM-DM] defines a performance monitoring profile for MPLS-TP.

4.5.4.8.1. Packet Loss Measurement (LM)

Packet Loss Measurement is a function used to verify the quality of the service. Packet loss, as defined in [IPPM-1LM] and [MPLS-TP-OAM], indicates the ratio of the number of user packets lost to the total number of user packets sent during a defined time interval.

There are two possible ways of determining this measurement:

- o Using OAM packets, it is possible to compute the statistics based on a series of OAM packets. This, however, has the disadvantage of being artificial and may not be representative since part of the packet loss may be dependent upon packet sizes and upon the implementation of the MEPs that take part in the protocol.
- o Delimiting messages can be sent at the start and end of a measurement period during which the source and sink of the path count the packets transmitted and received. After the end delimiter, the ratio would be calculated by the path OAM entity.

4.5.4.8.2. Packet Delay Measurement (DM)

Packet Delay Measurement is a function that is used to measure one-way or two-way delay of a packet transmission between a pair of the endpoints of a path (PW, LSP, or Section). Where:

- o One-way packet delay, as defined in [IPPM-1DM], is the time elapsed from the start of transmission of the first bit of the packet by a source node until the reception of the last bit of that packet by the destination node. Note that one-way delay measurement requires the clocks of the two endpoints to be synchronized.
- o Two-way packet delay, as defined in [IPPM-2DM], is the time elapsed from the start of transmission of the first bit of the packet by a source node until the reception of the last bit of the looped-back packet by the same source node, when the loopback is performed at the packet's destination node. Note that due to possible path asymmetry, the one-way packet delay from one endpoint to another is not necessarily equal to half of the

two-way packet delay. As opposed to one-way delay measurement, two-way delay measurement does not require the two endpoints to be synchronized.

For each of these two metrics, the DM function allows the MEP to measure the delay, as well as the delay variation. Delay measurement is performed by exchanging timestamped OAM packets between the participating MEPs.

4.6. Pseudowire OAM

4.6.1. Pseudowire OAM Using Virtual Circuit Connectivity Verification (VCCV)

VCCV, as defined in [VCCV], provides a means for end-to-end fault detection and diagnostic tools to be used for PWs (regardless of the underlying tunneling technology). The VCCV switching function provides a Control Channel associated with each PW. [VCCV] defines three Control Channel (CC) types, i.e., three possible methods for transmitting and identifying OAM messages:

- o Control Channel Type 1: In-band VCCV, as described in [VCCV], is also referred to as "PWE3 Control Word with 0001b as first nibble". It uses the PW Associated Channel Header [PW-ACH].
- o Control Channel Type 2: Out-of-band VCCV, as described in [VCCV], is also referred to as "MPLS Router Alert Label". In this case, the Control Channel is created by using the MPLS router alert label [MPLS-ENCAPS] immediately above the PW label.
- o Control Channel Type 3: TTL expiry VCCV, as described in [VCCV], is also referred to as "MPLS PW Label with TTL == 1", i.e., the Control Channel is identified when the value of the TTL field in the PW label is set to 1.

VCCV currently supports the following OAM tools: ICMP Ping, LSP Ping, and BFD. ICMP and LSP Ping are IP encapsulated before being sent over the PW ACH. BFD for VCCV [BFD-VCCV] supports two modes of encapsulation -- either IP/UDP encapsulated (with IP/UDP header) or PW-ACH encapsulated (with no IP/UDP header) -- and provides support to signal the AC status. The use of the VCCV Control Channel provides the context, based on the MPLS-PW label, required to bind and bootstrap the BFD session to a particular pseudowire (FEC), eliminating the need to exchange Discriminator values.

VCCV consists of two components: (1) the signaled component to communicate VCCV capabilities as part of the VC label, and (2) the switching component to cause the PW payload to be treated as a control packet.

VCCV is not directly dependent upon the presence of a control plane. The VCCV capability advertisement may be performed as part of the PW signaling when LDP is used. In case of manual configuration of the PW, it is the responsibility of the operator to set consistent options at both ends. The manual option was created specifically to handle MPLS-TP use cases where no control plane was a requirement. However, new use cases such as pure mobile backhaul find this functionality useful too.

The PWE3 working group has conducted an implementation survey of VCCV [VCCV-SURVEY] that analyzes which VCCV mechanisms are used in practice.

4.6.2. Pseudowire OAM Using G-ACh

As mentioned above, VCCV enables OAM for PWs by using a Control Channel for OAM packets. When PWs are used in MPLS-TP networks, rather than the Control Channels defined in VCCV, the G-ACh can be used as an alternative Control Channel. The usage of the G-ACh for PWs is defined in [PW-G-ACh].

4.6.3. Attachment Circuit - Pseudowire Mapping

The PWE3 working group has defined a mapping and notification of defect states between a pseudowire (PW) and the Attachment Circuits (ACs) of the end-to-end emulated service. This mapping is of key importance to the end-to-end functionality. Specifically, the mapping is provided by [PW-MAP], by [L2TP-EC] for L2TPv3 pseudowires, and by Section 5.3 of [ATM-L2] for ATM.

[L2VPN-OAM] provides the requirements and framework for OAM in the context of Layer 2 Virtual Private Networks (L2VPNs), and specifically it also defines the OAM layering of L2VPNs over pseudowires.

The mapping defined in [Eth-Int] allows an end-to-end emulated Ethernet service over pseudowires.

4.7. OWAMP and TWAMP

4.7.1. Overview

The IPPM working group in the IETF defines common criteria and metrics for measuring performance of IP traffic ([IPPM-FW]). Some of the key RFCs published by this working group have defined metrics for measuring connectivity [IPPM-Con], delay ([IPPM-1DM], [IPPM-2DM]), and packet loss [IPPM-1LM]. It should be noted that the work of the IETF in the context of performance metrics is not limited to IP networks; [PM-CONS] presents general guidelines for considering new performance metrics.

The IPPM working group has defined not only metrics for performance measurement but also protocols that define how the measurement is carried out. The One-Way Active Measurement Protocol [OWAMP] and the Two-Way Active Measurement Protocol [TWAMP] each define a method and protocol for measuring performance metrics in IP networks.

OWAMP [OWAMP] enables measurement of one-way characteristics of IP networks, such as one-way packet loss and one-way delay. For its proper operation, OWAMP requires accurate time-of-day setting at its endpoints.

TWAMP [TWAMP] is a similar protocol that enables measurement of both one-way and two-way (round-trip) characteristics.

OWAMP and TWAMP are each comprised of two separate protocols:

- o OWAMP-Control/TWAMP-Control: used to initiate, start, and stop test sessions and to fetch their results. Continuity Check and Connectivity Verification are tested and confirmed by establishing the OWAMP/TWAMP Control Protocol TCP connection.
- o OWAMP-Test/TWAMP-Test: used to exchange test packets between two measurement nodes. Enables the loss and delay measurement functions, as well as detection of other anomalies, such as packet duplication and packet reordering.

It should be noted that while [OWAMP] and [TWAMP] define tools for performance measurement, they do not define the accuracy of these tools. The accuracy depends on scale, implementation, and network configurations.

Alternative protocols for performance monitoring are defined, for example, in MPLS-TP OAM ([MPLS-LM-DM], [TP-LM-DM]) and in Ethernet OAM [ITU-T-Y1731].

4.7.2. Control and Test Protocols

OWAMP and TWAMP control protocols run over TCP, while the test protocols run over UDP. The purpose of the control protocols is to initiate, start, and stop test sessions, and for OWAMP to fetch results. The test protocols introduce test packets (which contain sequence numbers and timestamps) along the IP path under test according to a schedule, and they record statistics of packet arrival. Multiple sessions may be simultaneously defined, each with a session identifier, and defining the number of packets to be sent, the amount of padding to be added (and thus the packet size), the start time, and the send schedule (which can be either a constant time between test packets or exponentially distributed pseudorandomly). Statistics recorded conform to the relevant IPPM RFCs.

From a security perspective, OWAMP and TWAMP test packets are hard to detect because they are simply UDP streams between negotiated port numbers, with potentially nothing static in the packets. OWAMP and TWAMP also include optional authentication and encryption for both control and test packets.

4.7.3. OWAMP

OWAMP defines the following logical roles: Session-Sender, Session-Receiver, Server, Control-Client, and Fetch-Client. The Session-Sender originates test traffic that is received by the Session-Receiver. The Server configures and manages the session, as well as returning the results. The Control-Client initiates requests for test sessions, triggers their start, and may trigger their termination. The Fetch-Client requests the results of a completed session. Multiple roles may be combined in a single host -- for example, one host may play the roles of Control-Client, Fetch-Client, and Session-Sender, and a second may play the roles of Server and Session-Receiver.

In a typical OWAMP session, the Control-Client establishes a TCP connection to port 861 of the Server, which responds with a Server greeting message indicating supported security/integrity modes. The Control-Client responds with the chosen communications mode, and the Server accepts the mode. The Control-Client then requests and fully describes a test session to which the Server responds with its acceptance and supporting information. More than one test session may be requested with additional messages. The Control-Client then starts a test session; the Server acknowledges and then instructs the Session-Sender to start the test. The Session-Sender then sends test packets with pseudorandom padding to the Session-Receiver until the session is complete or until the Control-Client stops the session.

Once finished, the Session-Sender reports to the Server, which recovers data from the Session-Receiver. The Fetch-Client can then send a fetch request to the Server, which responds with an acknowledgement and, immediately thereafter, the result data.

4.7.4. TWAMP

TWAMP defines the following logical roles: Session-Sender, Session-Reflector, Server, and Control-Client. These are similar to the OWAMP roles, except that the Session-Reflector does not collect any packet information, and there is no need for a Fetch-Client.

In a typical TWAMP session, the Control-Client establishes a TCP connection to port 862 of the Server, and the mode is negotiated as in OWAMP. The Control-Client then requests sessions and starts them. The Session-Sender sends test packets with pseudorandom padding to the Session-Reflector, which returns them with timestamps inserted.

4.8. TRILL

The requirements of OAM in TRILL are defined in [TRILL-OAM]. The challenge in TRILL OAM, much like in MPLS networks, is that traffic between RBridges RB1 and RB2 may be forwarded through more than one path. Thus, an OAM protocol between RBridges RB1 and RB2 must be able to monitor all the available paths between the two RBridges.

During the writing of this document, the detailed definition of the TRILL OAM tools is still work in progress. This subsection presents the main requirements of TRILL OAM.

The main requirements defined in [TRILL-OAM] are:

- o Continuity Checking (CC) - the TRILL OAM protocol must support a function for CC between any two RBridges RB1 and RB2.
- o Connectivity Verification (CV) - connectivity between two RBridges RB1 and RB2 can be verified on a per-flow basis.
- o Path Tracing - allows an RBridge to trace all the available paths to a peer RBridge.
- o Performance monitoring - allows an RBridge to monitor the packet loss and packet delay to a peer RBridge.

5. Summary

This section summarizes the OAM tools and functions presented in this document. This summary is an index to some of the main OAM tools defined in the IETF. This compact index can be useful to all readers from network operators to standards development organizations. The summary includes a short subsection that presents some guidance to network equipment vendors.

5.1. Summary of OAM Tools

This subsection provides a short summary of each of the OAM toolsets described in this document.

A detailed list of the RFCs related to each toolset is given in Appendix A.1.

Toolset	Description	Transport Technology
IP Ping	Ping ([IntHost], [NetTerms]) is a simple application for testing reachability that uses ICMP Echo messages ([ICMPv4], [ICMPv6]).	IPv4/IPv6
IP Traceroute	Traceroute ([TCPIP-Tools], [NetTools]) is an application that allows users to trace the path between an IP source and an IP destination, i.e., to identify the nodes along the path. If more than one path exists between the source and destination, Traceroute traces *a* path. The most common implementation of Traceroute uses UDP probe messages, although there are other implementations that use different probes, such as ICMP or TCP. Paris Traceroute [PARIS] is an extension that attempts to discover all the available paths from A to B by scanning different values of header fields.	IPv4/IPv6
BFD	Bidirectional Forwarding Detection (BFD) is defined in [BFD] as a framework for a lightweight generic OAM tool. The intention is to define a base tool that can be used with various encapsulation types, network environments, and various medium types.	generic
MPLS OAM	MPLS LSP Ping, as defined in [MPLS-OAM], [MPLS-OAM-FW], and [LSP-Ping], is an OAM tool for point-to-point and point-to-multipoint MPLS LSPs. It includes two main functions: Ping and Traceroute. BFD [BFD-LSP] is an alternative means for detecting MPLS LSP data-plane failures.	MPLS

MPLS-TP OAM	MPLS-TP OAM is defined in a set of RFCs. The OAM requirements for MPLS Transport Profile (MPLS-TP) are defined in [MPLS-TP-OAM]. Each of the tools in the OAM toolset is defined in its own RFC, as specified in Appendix A.1.	MPLS-TP
Pseudowire OAM	The PWE3 OAM architecture defines Control Channels that support the use of existing IETF OAM tools to be used for a pseudowire (PW). The Control Channels that are defined in [VCCV] and [PW-G-ACh] may be used in conjunction with ICMP Ping, LSP Ping, and BFD to perform CC and CV functionality. In addition, the channels support use of any of the MPLS-TP-based OAM tools for completing their respective OAM functionality for a PW.	Pseudowire
OWAMP and TWAMP	The One-Way Active Measurement Protocol [OWAMP] and the Two-Way Active Measurement Protocol [TWAMP] are two protocols defined in the IP Performance Metrics (IPPM) working group in the IETF. These protocols allow various performance metrics to be measured, such as packet loss, delay, delay variation, duplication, and reordering.	IPv4/IPv6
TRILL OAM	The requirements of OAM in TRILL are defined in [TRILL-OAM]. These requirements include Continuity Checking, Connectivity Verification, path tracing, and performance monitoring. During the writing of this document, the detailed definition of the TRILL OAM tools is work in progress.	TRILL

Table 3: Summary of OAM-Related IETF Tools

5.2. Summary of OAM Functions

Table 4 summarizes the OAM functions that are supported in each of the toolsets that were analyzed in this section. The columns of this table are the typical OAM functions described in Section 1.3.

Toolset	Continuity Check	Connectivity Verification	Path Discovery	Perf. Monitoring	Other Functions
IP Ping	Echo				
IP Traceroute			Traceroute		
BFD	BFD Control/ Echo	BFD Control			RDI using BFD Control
MPLS OAM (LSP Ping)		"Ping" mode	"Trace-route" mode		
MPLS-TP OAM	CC	CV/proactive or on demand	Route Tracing	-LM -DM	-Diagnostic Test -Lock -Alarm Reporting -Client Failure Indication -RDI
Pseudowire OAM	BFD	-BFD -ICMP Ping -LSP Ping	LSP Ping		
OWAMP and TWAMP		- control protocol		-DM -LM	
TRILL OAM	CC	CV	Path tracing	-DM -LM	

Table 4: Summary of the OAM Functionality in IETF OAM Tools

5.3. Guidance to Network Equipment Vendors

As mentioned in Section 1.4, it is imperative for OAM tools to be capable of testing the actual data plane with as much accuracy as possible. While this guideline may appear obvious, it is worthwhile to emphasize the key importance of enforcing fate-sharing between OAM traffic that monitors the data plane and the data-plane traffic it monitors.

6. Security Considerations

OAM is tightly coupled with the stability of the network. A successful attack on an OAM protocol can create a false illusion of nonexistent failures or prevent the detection of actual ones. In both cases, the attack may result in denial of service.

Some of the OAM tools presented in this document include security mechanisms that provide integrity protection, thereby preventing attackers from forging or tampering with OAM packets. For example, [BFD] includes an optional authentication mechanism for BFD Control packets, using either SHA1, MD5, or a simple password. [OWAMP] and [TWAMP] have three modes of security: unauthenticated, authenticated, and encrypted. The authentication uses SHA1 as the HMAC algorithm, and the encrypted mode uses AES encryption.

Confidentiality is typically not considered a requirement for OAM protocols. However, the use of encryption (e.g., [OWAMP] and [TWAMP]) can make it difficult for attackers to identify OAM packets, thus making it more difficult to attack the OAM protocol.

OAM can also be used as a means for network reconnaissance; information about addresses, port numbers, and the network topology and performance can be gathered by either passively eavesdropping on OAM packets or actively sending OAM packets and gathering information from the respective responses. This information can then be used maliciously to attack the network. Note that some of this information, e.g., addresses and port numbers, can be gathered even when encryption is used ([OWAMP], [TWAMP]).

For further details about the security considerations of each OAM protocol, the reader is encouraged to review the Security Considerations section of each document referenced by this memo.

7. Acknowledgments

The authors gratefully acknowledge Sasha Vainshtein, Carlos Pignataro, David Harrington, Dan Romascanu, Ron Bonica, Benoit Claise, Stewart Bryant, Tom Nadeau, Elwyn Davies, Al Morton, Sam Aldrin, Thomas Narten, and other members of the OPSA WG for their helpful comments on the mailing list.

This document was originally prepared using 2-Word-v2.0.template.dot.

8. References

8.1. Normative References

[OAM-Def] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, June 2011.

8.2. Informative References

- [ATM-L2] Singh, S., Townsley, M., and C. Pignataro, "Asynchronous Transfer Mode (ATM) over Layer 2 Tunneling Protocol Version 3 (L2TPv3)", RFC 4454, May 2006.
- [BFD] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.
- [BFD-Gen] Katz, D. and D. Ward, "Generic Application of Bidirectional Forwarding Detection (BFD)", RFC 5882, June 2010.
- [BFD-IP] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, June 2010.
- [BFD-LSP] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010.
- [BFD-Multi] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, June 2010.

- [BFD-VCCV] Nadeau, T., Ed., and C. Pignataro, Ed., "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", RFC 5885, June 2010.
- [Comp] Bonaventure, O., "Computer Networking: Principles, Protocols and Practice", 2008.
- [Dup] Uijterwaal, H., "A One-Way Packet Duplication Metric", RFC 5560, May 2009.
- [Eth-Int] Mohan, D., Ed., Bitar, N., Ed., Sajassi, A., Ed., DeLord, S., Niger, P., and R. Qiu, "MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking", RFC 7023, October 2013.
- [G-ACh] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, June 2009.
- [ICMP-Ext] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "ICMP Extensions for Multiprotocol Label Switching", RFC 4950, August 2007.
- [ICMP-Int] Atlas, A., Ed., Bonica, R., Ed., Pignataro, C., Ed., Shen, N., and JR. Rivers, "Extending ICMP for Interface and Next-Hop Identification", RFC 5837, April 2010.
- [ICMP-MP] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "Extended ICMP to Support Multi-Part Messages", RFC 4884, April 2007.
- [ICMPv4] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.
- [ICMPv6] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.
- [IEEE802.1Q] IEEE, "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE 802.1Q, October 2012.

- [IEEE802.3ah] IEEE, "IEEE Standard for Information technology - Local and metropolitan area networks - Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications", IEEE 802.3ah, clause 57, December 2008.
- [IntHost] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.
- [IPPM-1DM] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", RFC 2679, September 1999.
- [IPPM-1LM] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999.
- [IPPM-2DM] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, September 1999.
- [IPPM-Con] Mahdavi, J. and V. Paxson, "IPPM Metrics for Measuring Connectivity", RFC 2678, September 1999.
- [IPPM-FW] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.
- [ITU-G8113.1] ITU-T, "Operations, Administration and Maintenance mechanism for MPLS-TP in Packet Transport Network (PTN)", ITU-T Recommendation G.8113.1/Y.1372.1, November 2012.
- [ITU-G8113.2] ITU-T, "Operations, administration and maintenance mechanisms for MPLS-TP networks using the tools defined for MPLS", ITU-T Recommendation G.8113.2/Y.1372.2, November 2012.
- [ITU-T-CT] Betts, M., "Allocation of a Generic Associated Channel Type for ITU-T MPLS Transport Profile Operation, Maintenance, and Administration (MPLS-TP OAM)", RFC 6671, November 2012.
- [ITU-T-G.806] ITU-T, "Characteristics of transport equipment - Description methodology and generic functionality", ITU-T Recommendation G.806, January 2009.
- [ITU-T-Y1711] ITU-T, "Operation & Maintenance mechanism for MPLS networks", ITU-T Recommendation Y.1711, February 2004.

- [ITU-T-Y1731] ITU-T, "OAM Functions and Mechanisms for Ethernet-based Networks", ITU-T Recommendation G.8013/Y.1731, July 2011.
- [ITU-Terms] ITU-R/ITU-T, "ITU-R/ITU-T Terms and Definitions", 2013, <<http://www.itu.int/pub/R-TER-DB>>.
- [L2TP-EC] McGill, N. and C. Pignataro, "Layer 2 Tunneling Protocol Version 3 (L2TPv3) Extended Circuit Status Values", RFC 5641, August 2009.
- [L2VPN-OAM] Sajassi, A., Ed., and D. Mohan, Ed., "Layer 2 Virtual Private Network (L2VPN) Operations, Administration, and Maintenance (OAM) Requirements and Framework", RFC 6136, March 2011.
- [L3VPN-OAM] El Mghazli, Y., Ed., Nadeau, T., Boucadair, M., Chan, K., and A. Gonguet, "Framework for Layer 3 Virtual Private Networks (L3VPN) Operations and Management", RFC 4176, October 2005.
- [Lock-Loop] Boutros, S., Ed., Sivabalan, S., Ed., Aggarwal, R., Ed., Vigoureux, M., Ed., and X. Dai, Ed., "MPLS Transport Profile Lock Instruct and Loopback Functions", RFC 6435, November 2011.
- [LSP-Ping] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.
- [Mng] Farrel, A., "Inclusion of Manageability Sections in Path Computation Element (PCE) Working Group Drafts", RFC 6123, February 2011.
- [MPLS-ENCAPS] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [MPLS-LM-DM] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011.
- [MPLS-OAM] 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.

- [MPLS-OAM-FW] Allan, D., Ed., and T. Nadeau, Ed., "A Framework for Multi-Protocol Label Switching (MPLS) Operations and Management (OAM)", RFC 4378, February 2006.
- [MPLS-P2MP] Yasukawa, S., Farrel, A., King, D., and T. Nadeau, "Operations and Management (OAM) Requirements for Point-to-Multipoint MPLS Networks", RFC 4687, September 2006.
- [MPLS-TP-OAM] Vigoureux, M., Ed., Ward, D., Ed., and M. Betts, Ed., "Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks", RFC 5860, May 2010.
- [mtrace] Fenner, W. and S. Casner, "A "traceroute" facility for IP Multicast", Work in Progress, July 2000.
- [NetTerms] Jacobsen, O. and D. Lynch, "A Glossary of Networking Terms", RFC 1208, March 1991.
- [NetTools] Enger, R. and J. Reynolds, "FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices", FYI 2, RFC 1470, June 1993.
- [OAM-Analys] Sprecher, N. and L. Fang, "An Overview of the Operations, Administration, and Maintenance (OAM) Toolset for MPLS-Based Transport Networks", RFC 6669, July 2012.
- [OAM-Label] Ohta, H., "Assignment of the 'OAM Alert Label' for Multiprotocol Label Switching Architecture (MPLS) Operation and Maintenance (OAM) Functions", RFC 3429, November 2002.
- [OAM-Mng] Ersue, M., Ed., and B. Claise, "An Overview of the IETF Network Management Standards", RFC 6632, June 2012.
- [OnDemand-CV] Gray, E., Bahadur, N., Boutros, S., and R. Aggarwal, "MPLS On-Demand Connectivity Verification and Route Tracing", RFC 6426, November 2011.
- [OWAMP] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006.

- [PARIS] Augustin, B., Friedman, T., and R. Teixeira, "Measuring Load-balanced Paths in the Internet", IMC '07 Proceedings of the 7th ACM SIGCOMM conference on Internet measurement, 2007.
- [PM-CONS] Clark, A. and B. Claise, "Guidelines for Considering New Performance Metric Development", BCP 170, RFC 6390, October 2011.
- [PW-ACH] 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.
- [PW-G-ACh] Li, H., Martini, L., He, J., and F. Huang, "Using the Generic Associated Channel Label for Pseudowire in the MPLS Transport Profile (MPLS-TP)", RFC 6423, November 2011.
- [PW-MAP] 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.
- [Reorder] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", RFC 4737, November 2006.
- [Signal] Yasukawa, S., Ed., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, April 2006.
- [TCPIP-Tools] Kessler, G. and S. Shepard, "A Primer On Internet and TCP/IP Tools and Utilities", FYI 30, RFC 2151, June 1997.
- [TP-CC-CV] Allan, D., Ed., Swallow Ed., G., and J. Drake Ed., "Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile", RFC 6428, November 2011.
- [TP-Fault] Swallow, G., Ed., Fulignoli, A., Ed., Vigoureux, M., Ed., Boutros, S., and D. Ward, "MPLS Fault Management Operations, Administration, and Maintenance (OAM)", RFC 6427, November 2011.
- [TP-LM-DM] Frost, D., Ed., and S. Bryant, Ed., "A Packet Loss and Delay Measurement Profile for MPLS-Based Transport Networks", RFC 6375, September 2011.

- [TP-OAM-FW] Busi, I., Ed., and D. Allan, Ed., "Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks", RFC 6371, September 2011.
- [TP-Term] van Helvoort, H., Ed., Andersson, L., Ed., and N. Sprecher, Ed., "A Thesaurus for the Interpretation of Terminology Used in MPLS Transport Profile (MPLS-TP) Internet-Drafts and RFCs in the Context of the ITU-T's Transport Network Recommendations", RFC 7087, December 2013.
- [TRILL-OAM] Senevirathne, T., Bond, D., Aldrin, S., Li, Y., and R. Watve, "Requirements for Operations, Administration, and Maintenance (OAM) in Transparent Interconnection of Lots of Links (TRILL)", RFC 6905, March 2013.
- [TWAMP] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarez, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.
- [VCCV] Nadeau, T., Ed., and C. Pignataro, Ed., "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, December 2007.
- [VCCV-SURVEY] Del Regno, N., Ed., and A. Malis, Ed., "The Pseudowire (PW) and Virtual Circuit Connectivity Verification (VCCV) Implementation Survey Results", RFC 7079, November 2013.

Appendix A. List of OAM Documents

A.1. List of IETF OAM Documents

Table 5 summarizes the OAM-related RFCs produced by the IETF.

It is important to note that the table lists various RFCs that are different by nature. For example, some of these documents define OAM tools or OAM protocols (or both), while others define protocols that are not strictly OAM related, but are used by OAM tools. The table also includes RFCs that define the requirements or the framework of OAM in a specific context (e.g., MPLS-TP).

The RFCs in the table are categorized in a few sets as defined in Section 1.3.

Toolset	Title	RFC
IP Ping	Requirements for Internet Hosts -- Communication Layers [IntHost]	RFC 1122
	A Glossary of Networking Terms [NetTerms]	RFC 1208
	Internet Control Message Protocol [ICMPv4]	RFC 792
	Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification [ICMPv6]	RFC 4443
IP Traceroute	A Primer On Internet and TCP/IP Tools and Utilities [TCPIP-Tools]	RFC 2151
	FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices [NetTools]	RFC 1470
	Internet Control Message Protocol [ICMPv4]	RFC 792
	Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification [ICMPv6]	RFC 4443

	Extended ICMP to Support Multi-Part Messages [ICMP-MP]	RFC 4884
	Extending ICMP for Interface and Next-Hop Identification [ICMP-Int]	RFC 5837
BFD	Bidirectional Forwarding Detection (BFD) [BFD]	RFC 5880
	Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop) [BFD-IP]	RFC 5881
	Generic Application of Bidirectional Forwarding Detection (BFD)[BFD-Gen]	RFC 5882
	Bidirectional Forwarding Detection (BFD) for Multihop Paths [BFD-Multi]	RFC 5883
	Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs) [BFD-LSP]	RFC 5884
	Bidirectional Forwarding Detection for the Pseudowire Virtual Circuit Connectivity Verification (VCCV) [BFD-VCCV]	RFC 5885
MPLS OAM	Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks [MPLS-OAM]	RFC 4377
	A Framework for Multi-Protocol Label Switching (MPLS) Operations and Management (OAM) [MPLS-OAM-FW]	RFC 4378
	Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures [LSP-Ping]	RFC 4379
	Operations and Management (OAM) Requirements for Point-to-Multipoint MPLS Networks [MPLS-P2MP]	RFC 4687
	ICMP Extensions for Multiprotocol Label Switching [ICMP-Ext]	RFC 4950

	Bidirectional Forwarding Detection for MPLS Label Switched Paths (LSPs) [BFD-LSP]	RFC 5884
MPLS-TP OAM	Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks [MPLS-TP-OAM]	RFC 5860
	MPLS Generic Associated Channel [G-ACh]	RFC 5586
	Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks [TP-OAM-FW]	RFC 6371
	Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile [TP-CC-CV]	RFC 6428
	MPLS On-Demand Connectivity Verification and Route Tracing [OnDemand-CV]	RFC 6426
	MPLS Fault Management Operations, Administration, and Maintenance (OAM) [TP-Fault]	RFC 6427
	MPLS Transport Profile Lock Instruct and Loopback Functions [Lock-Loop]	RFC 6435
	Packet Loss and Delay Measurement for MPLS Networks [MPLS-LM-DM]	RFC 6374
	A Packet Loss and Delay Measurement Profile for MPLS-Based Transport Networks [TP-LM-DM]	RFC 6375
Pseudowire OAM	Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires [VCCV]	RFC 5085

	Bidirectional Forwarding Detection for the Pseudowire Virtual Circuit Connectivity Verification (VCCV) [BFD-VCCV]	RFC 5885
	Using the Generic Associated Channel Label for Pseudowire in the MPLS Transport Profile (MPLS-TP) [PW-G-ACh]	RFC 6423
	Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping [PW-MAP]	RFC 6310
	MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking [Eth-Int]	RFC 7023
OWAMP and TWAMP	A One-way Active Measurement Protocol (OWAMP) [OWAMP]	RFC 4656
	A Two-Way Active Measurement Protocol (TWAMP) [TWAMP]	RFC 5357
	Framework for IP Performance Metrics [IPPM-FW]	RFC 2330
	IPPM Metrics for Measuring Connectivity [IPPM-Con]	RFC 2678
	A One-way Delay Metric for IPPM [IPPM-1DM]	RFC 2679
	A One-way Packet Loss Metric for IPPM [IPPM-1LM]	RFC 2680
	A Round-trip Delay Metric for IPPM [IPPM-2DM]	RFC 2681
	Packet Reordering Metrics [Reorder]	RFC 4737
	A One-Way Packet Duplication Metric [Dup]	RFC 5560

TRILL OAM	Requirements for Operations, Administration, and Maintenance (OAM) in Transparent Interconnection of Lots of Links (TRILL)	RFC 6905
-----------	--	----------

Table 5: Summary of IETF OAM-Related RFCs

A.2. List of Selected Non-IETF OAM Documents

In addition to the OAM tools defined by the IETF, the IEEE and ITU-T have also defined various OAM tools that focus on Ethernet and various other transport-network environments. These various tools, defined by the three standard organizations, are often tightly coupled and have had a mutual effect on each other. The ITU-T and IETF have both defined OAM tools for MPLS LSPs, [ITU-T-Y1711], and [LSP-Ping]. The following OAM standards by the IEEE and ITU-T are to some extent linked to the IETF OAM tools listed above and are mentioned here only as reference material.

- o OAM tools for Layer 2 have been defined by the ITU-T in [ITU-T-Y1731] and by the IEEE in 802.1ag [IEEE802.1Q]. The IEEE 802.3 standard defines OAM for one-hop Ethernet links [IEEE802.3ah].
- o The ITU-T has defined OAM for MPLS LSPs in [ITU-T-Y1711] and for MPLS-TP OAM in [ITU-G8113.1] and [ITU-G8113.2].

It should be noted that these non-IETF documents deal in many cases with OAM functions below the IP layer (Layer 2, Layer 2.5) and that in some cases operators use a multi-layered OAM approach, which is a function of the way their networks are designed.

Table 6 summarizes some of the main OAM standards published by non-IETF standard organizations. This document focuses on IETF OAM standards, but these non-IETF standards are referenced in this document where relevant.

	Title	Document
ITU-T MPLS OAM	Operation & Maintenance mechanism for MPLS networks [ITU-T-Y1711]	ITU-T Y.1711
	Assignment of the 'OAM Alert Label' for Multiprotocol Label Switching Architecture (MPLS) Operation and Maintenance (OAM) Functions [OAM-Label] Note: although this is an IETF document, it is listed as one of the non-IETF OAM standards, since it was defined as a complementary part of ITU-T Y.1711.	RFC 3429
ITU-T MPLS-TP OAM	Operations, administration and Maintenance mechanisms for MPLS-TP networks using the tools defined for MPLS [ITU-G8113.2] Note: this document describes the OAM toolset defined by the IETF for MPLS-TP, whereas ITU-T G.8113.1 describes the OAM toolset defined by the ITU-T.	ITU-T G.8113.2
	Operations, Administration and Maintenance mechanism for MPLS-TP in Packet Transport Network (PTN)	ITU-T G.8113.1

	<p>Allocation of a Generic Associated Channel Type for ITU-T MPLS Transport Profile Operation, Maintenance, and Administration (MPLS-TP OAM) [ITU-T-CT]</p> <p>Note: although this is an IETF document, it is listed as one of the non-IETF OAM standards, since it was defined as a complementary part of ITU-T G.8113.1.</p>	RFC 6671
ITU-T Ethernet OAM	OAM Functions and Mechanisms for Ethernet-based Networks [ITU-T-Y1731]	ITU-T Y.1731
IEEE CFM	<p>Connectivity Fault Management [IEEE802.1Q]</p> <p>Note: CFM was originally published as IEEE 802.1ag but is now incorporated in the 802.1Q standard.</p>	IEEE 802.1ag
IEEE DDCFM	<p>Management of Data Driven and Data Dependent Connectivity Faults [IEEE802.1Q]</p> <p>Note: DDCFm was originally published as IEEE 802.1Qaw but is now incorporated in the 802.1Q standard.</p>	IEEE 802.1ag
IEEE 802.3 link level OAM	<p>Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks [IEEE802.3ah]</p> <p>Note: link level OAM was originally defined in IEEE 802.3ah and is now incorporated in the 802.3 standard.</p>	IEEE 802.3ah

Table 6: Non-IETF OAM Standards Mentioned in This Document

Authors' Addresses

Tal Mizrahi
Marvell
6 Hamada St.
Yokneam 20692
Israel

EEmail: talmi@marvell.com

Nurit Sprecher
Nokia Solutions and Networks
3 Hanagar St. Neve Ne'eman B
Hod Hasharon 45241
Israel

EEmail: nurit.sprecher@nsn.com

Elisa Bellagamba
Ericsson
6 Farogatan St.
Stockholm 164 40
Sweden

Phone: +46 761440785

EEmail: elisa.bellagamba@ericsson.com

Yaacov Weingarten
34 Hagefen St.
Karnei Shomron 4485500
Israel

EEmail: wyaacov@gmail.com