Operations and Management Area Working Group
Internet Draft
Intended status: Informational
Expires: March 2013

T. Mizrahi
Marvell
N. Sprecher
Nokia Siemens Networks
E. Bellagamba
Ericsson
Y. Weingarten

September 12, 2012

An Overview of Operations, Administration, and Maintenance (OAM) Mechanisms
draft-ietf-opsawg-oam-overview-07.txt

Abstract

Operations, Administration, and Maintenance (OAM) is a general term that refers to a toolset that can be used for fault detection and isolation, and for performance measurement. OAM mechanisms have been defined for various layers in the protocol stack, and are used with a variety of protocols.

This document presents an overview of the OAM mechanisms that have been defined and are currently being defined by the IETF.

Table of Contents

1. Introduction

OAM is a general term that refers to a toolset that can be used for detecting, isolating and reporting connection failures or measurement of connection performance parameters. The term OAM has been used over the years in several different contexts, as discussed in [OAM-Def]. This term as been associated with the 3 logical abstraction layers: the forwarding plane, the control plane, and the management plane. In the context of this document OAM refers to Operations, Administration, and Maintenance. Hence, management aspects are outside the scope of this document.

1.1. Background

The communication of a network may be configured and maintained by use of various tools at different layers - these include use of a control plane or management plane to configure and maintain the connectivity of the network from the outside - looking in - and controlling the connections when the need arises. OAM, on the other hand, traditionally has been used to maintain the connectivity in-band with the actual data traffic, i.e. in the data plane.

While the OAM tools may, and quite often do, work in conjunction with a control-plane or management plane, they are usually defined to be independent of the control-plane. The OAM tools communicate with the management plane to raise alarms, and often the on-demand tools may be activated by the management, e.g. to locate and localize problems.

The considerations of the control-plane maintenance tools or the functionality of the management-plane are out of scope for this document, which will concentrate on presenting the data-plane tools that are used for OAM.

1.2. The OAM toolsets

This memo provides an overview of the different sets of OAM mechanisms defined by the IETF. It is intended for those with little or no familiarity with the described mechanisms. The set of OAM mechanisms described in this memo are applicable to IP unicast, MPLS, pseudowires, and MPLS for the transport environment (MPLS-TP). While OAM mechanisms 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 mechanisms in several different fronts:

- ICMP Ping:
  ICMP Echo request, also known as Ping, as defined in [ICMPv4], and [ICMPv6]. ICMP Ping is a very simple and basic mechanism in failure diagnosis. LSP Ping is to some extent based on ICMP Ping.

- IPPM:
  IP Performance Metrics (IPPM) is a working group in the IETF that defined common metrics for performance measurement, as well as a protocol for measuring delay and packet loss in IP networks.

- MPLS:
  MPLS LSP Ping, as defined in [MPLS-OAM], [MPLS-OAM-FW] and [LSP-Ping], is an OAM mechanism for point to point MPLS LSPs.

- MPLS-TP:
  The OAM requirements for MPLS Transport Profile (MPLS-TP) are defined in [MPLS-TP-OAM], and the toolset is described in [TP-OAM-FW].

- BFD:
  Bidirectional Forwarding Detection (BFD) is defined in [BFD] as a framework for a lightweight generic OAM mechanism. The intention is to define a base mechanism that can be used with various encapsulation types, network environments, and in various medium types.

This document summarizes the OAM mechanisms defined by the IETF. We first present a comparison of the terminology used in various OAM standards, and then summarize the OAM functions that each OAM standard provides.

1.3. IETF OAM Standards

Table 1 summarizes the IETF OAM standards discussed in this document.

The table includes a "Type" column, specifying the nature of each of the listed documents:

- Tool: documents that define an OAM tool or mechanism.
- Prof.: documents that define a profile or a variant for an OAM tool that is defined in other documents.
- Inf.: documents that define an infrastructure that is used by OAM tools.
- Misc.: other OAM related documents, e.g., OAM requirement and framework documents.

	Title	Type	RFC
ICMPv4 Ping	Internet Control Message Protocol	Tool	RFC 792
ICMPv6 Ping	Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification	Tool	RFC 4443
Traceroute	A Primer On Internet and TCP/IP Tools and Utilities	Tool	RFC 2151
BFD	Bidirectional Forwarding Detection	Tool	RFC 5880
	Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)	Prof.	RFC 5881
	Generic Application of Bidirectional Forwarding Detection	Misc.	RFC 5882

	Bidirectional Forwarding Detection (BFD) for Multihop Paths	Prof.	RFC 5883
	Bidirectional Forwarding Detection for MPLS Label Switched Paths (LSPs)	Prof.	RFC 5884
	Bidirectional Forwarding Detection for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)	Prof.	RFC 5885
IETF MPLS OAM (LSP Ping)	Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks	Misc.	RFC 4377
	A Framework for Multi-Protocol Label Switching (MPLS) Operations and Management (OAM)	Misc.	RFC 4378
	Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures	Tool	RFC 4379
	Operations and Management (OAM) Requirements for Point-to-Multipoint MPLS Networks	Misc.	RFC 4687
	ICMP Extensions for Multiprotocol Label Switching	Tool	RFC 4950
MPLS-TP OAM	Requirements for OAM in MPLS-TP	Misc.	RFC 5860
	MPLS Generic Associated Channel	Inf.	RFC 5586
	MPLS-TP OAM Framework	Misc.	RFC 6371
	Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile	Tool	RFC 6428
	MPLS On-Demand Connectivity Verification and Route Tracing	Tool	RFC 6426

	MPLS Fault Management Operations, Administration, and Maintenance (OAM)	Tool	RFC 6427
	MPLS Transport Profile Lock Instruct and Loopback Functions	Tool	RFC 6435
	Packet Loss and Delay Measurement for MPLS Networks	Tool	RFC 6374
	A Packet Loss and Delay Measurement Profile for MPLS-Based Transport Networks	Prof.	RFC 6375
PW VCCV	Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires	Inf.	RFC 5085
IPPM	Framework for IP Performance Metrics	Misc.	RFC 2330
	IPPM Metrics for Measuring Connectivity	Misc.	RFC 2678
	A One-way Delay Metric for IPPM	Misc.	RFC 2679
	A One-way Packet Loss Metric for IPPM	Misc.	RFC 2680
	A Round-trip Delay Metric for IPPM	Misc.	RFC 2681
	A One-way Active Measurement Protocol (OWAMP)	Tool	RFC 4656
	A Two-Way Active Measurement Protocol (TWAMP)	Tool	RFC 5357

Table 1 Summary of IETF OAM Related Standards

1.4. Non-IETF OAM Standards

In addition to the OAM mechanisms defined by the IETF, the IEEE and ITU-T have also defined various OAM mechanisms that focus on Ethernet, and various other transport network environments. These various mechanisms, 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 mechanisms for MPLS LSPs, [ITU-T-Y.1711] and [LSP-Ping]. The following OAM standards by the IEEE and ITU-T are to some extent linked to IETF OAM mechanisms listed above and are mentioned here only as reference material:

- OAM mechanisms for Ethernet based networks have been defined by both the ITU-T in [ITU-T-Y.1731], and by the IEEE in [IEEE-802.1ag]. The IEEE 802.3 standard defines OAM for one-hop Ethernet links [IEEE-802.3ah].

- The ITU-T has defined OAM for MPLS LSPs in [ITU-T-Y.1711].

Table 2 summarizes the OAM standards mentioned in this document. This document focuses on IETF OAM standards, but these non-IETF standards are referenced where relevant.

	Title	Standard/Draft
ITU-T MPLS OAM	Operation & Maintenance mechanism for MPLS networks	[ITU-T-Y.1711]
	Assignment of the 'OAM Alert Label' for Multiprotocol Label Switching Architecture (MPLS) Operation and Maintenance (OAM) Functions 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 Y.1711.	RFC 3429
ITU-T Ethernet OAM	OAM Functions and Mechanisms for Ethernet-based Networks	[ITU-T-Y.1731]

IEEE CFM	Connectivity Fault Management	[IEEE-802.1ag]
IEEE 802.3 link level OAM	Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks	[IEEE-802.3ah]

Table 2 Non-IETF OAM Standards Mentioned in this Document

2. Basic Terminology

2.1. Abbreviations

ACH Associated Channel Header

AIS Alarm Indication Signal

BFD Bidirectional Forwarding Detection

CC Continuity Check

CCM Continuity Check Message

CV Connectivity Verification

DM Delay Measurement

FEC Forwarding Equivalence Class

GAL Generic Associated Label

ICMP Internet Control Message Protocol

L2TP Layer Two Tunneling Protocol

LCCE L2TP Control Connection Endpoint

LDP Label Distribution Protocol

LM Loss Measurement

LOC Loss Of Continuity

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

PE Provider Edge

PW Pseudowire

PWE3 Pseudowire Emulation Edge-to-Edge

RDI Remote Defect Indication

TTL Time To Live

VCCV Virtual Circuit Connectivity Verification

2.2. Terminology used in OAM Standards

2.2.1. General Terms

A wide variety of terms is used in various OAM standards. Each of the OAM standards listed in the reference section includes a section that defines terms relevant to that tool. A thesaurus of terminology for MPLS-TP terms is presented in [MPLS-TP-Term], and provides a good summary of some of the OAM related terminology.

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

formal definition of each term, refer to the references at the end of this document.

2.2.2. OAM Maintenance Entities

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 following related terms are also quoted from [TP-OAM-FW]:

o MEP: The two points that define a maintenance entity.

o MEG: The collection of one or more MEs that belongs to the same transport path and that are maintained and monitored as a group are known as a Maintenance Entity Group (MEG).

o MIP: In between MEPs, there are zero or more intermediate points, called Maintenance Entity Group Intermediate Points (MIPs).

A pair of MEPs engaged in an ME are connected by a communication link, which may be one of several types of connection, e.g. a single physical connection, a set of physical connections, or a virtual link such as an MPLS LSP.

The term Maintenance Entity (ME) is used in ITU-T Recommendations (e.g. [ITU-T-Y.1731]), as well as in the MPLS-TP terminology ([TP-OAM-FW]). Various terms are used to refer to an ME. For example, BFD does not explicitly use a term that is equivalent to ME, but rather uses the term "session", referring to the relationship between two nodes using a BFD protocol. The MPLS LSP Ping ([LSP-Ping]) terminology simply uses the term "LSP" in this context.

MPLS-TP has defined the terms ME and Maintenance Entity Group (MEG) in [TP-OAM-FW], similar to the terms defined by ITU-T. A MEG allows the monitoring of a compound set of MEs, for example when monitoring a p2mp MEG that is considered to be the set of MEs between the root and each individual destination MEP.

2.2.3. OAM Maintenance Points

A Maintenance Point (MP) is a functional entity that is defined at a node in the network, and either initiates or reacts to OAM messages. A Maintenance End Point (MEP) is one of the end points of an ME, and can initiate OAM messages and respond to them. A Maintenance Intermediate Point (MIP) is an intermediate point between two MEPs,

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 value of the TTL field in the OAM packet. The term Maintenance Point is a general term for MEPs and MIPs.

The 802.1ag defines a finer distinction between Up MPs and Down MPs. An MP is a bridge interface, that is monitored by an OAM protocol either in the direction facing the network, or in the direction facing the bridge. A Down MP is an MP that receives OAM packets from, and transmits them to the direction of the network. An Up MP 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 MP - either at the ingress, egress, or forwarding function of the node (Down / Up MPs). This placement is important for localization of a connection failure.

2.2.4. Proactive and On-demand activation

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

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

- On-demand activation - indicates that the tool is activated "manually" to detect a specific anomaly. In this activation a small number of OAM messages are sent by a MEP and the reply message is received.

2.2.5. 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 checks are used to verify the liveness of a connection or a path between two MPs, and are typically sent proactively, though they can be invoked on-demand as well.

A connectivity verification function allows an MP to check whether it is connected to a peer MP or not. This function also allows the MP to

verify that messages from the peer MP 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. This allows detection of unexpected topology changes. 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 and continuity checks are considered complementary mechanisms, and are often used in conjunction with each other.

2.2.6. Failures

The terms Failure, Fault, and Defect are used interchangeably in the standards, referring to a malfunction that can be detected by a connectivity or a continuity check. In some standards, such as [IEEE-802.1ag], 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 takes after the ITU-T, which distinguishes between these terms in [ITU-T-G.806]; The term Fault refers to an inability to perform a required action, e.g., an unsuccessful attempt to deliver a packet. The term Defect refers to an interruption in the normal operation, such as a consecutive period of time where no packets are delivered successfully. 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 Tools

3.1. ICMP Ping

ICMP provides a connectivity verification 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.

3.2. Traceroute

Traceroute ([TCPIP-Tools]) is an application that allows users to discover the path between an IP source and an IP destination. Traceroute sends a sequence of UDP packets to UDP port 33434 at the destination. By default, Traceroute begins by sending three packets, each with an IP Time-To-Live (TTL) value of one to the destination.

These packets expire as soon as they reach the first router in the path. That router responds by sending three ICMP Time Exceeded Messages 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.

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

3.3. Bidirectional Forwarding Detection (BFD)

3.3.1. Overview

While multiple OAM mechanisms 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 mechanism 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 PWE3 [BFD-VCCV]. The usage of BFD in MPLS-TP is defined in [MPLS-TP-CC-CV].

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

3.3.2. 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:

- Asynchronous mode (i.e. proactive): in this mode BFD control packets are sent periodically. When the receiver detects that no BFD control packet have been received during a predetermined period of time, a failure is detected.

- 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 verify the liveness of the session. BFD control packets are sent independently in each direction.

Each of the end-points of the monitored path maintains its own session identification, called a Discriminator, both of which are included in the BFD Control Packets that are exchanged between the end-points. At the time of session establishment, the Discriminators are exchanged between the two-end points. In addition, the transmission (and reception) rate is negotiated between the two end-points, 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. no failures are detected, the BFD session is in the Up state. If no BFD Control packets are received during a fixed period of time, called the Detection Time, the session is declared to be Down. The detection time is a function of the negotiated transmission time, 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.

3.3.3. 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 is not used in BFD for MPLS LSPs, PWs, or in BFD for MPLS-TP.

3.4. LSP Ping

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 mechanism defined, in this context, is LSP Ping [LSP-Ping].

LSP Ping is based on ICMP Ping and just like its predecessor 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 extends the basic ICMP Ping operation (of data-plane connectivity verification) with functionality to verify data-plane vs. control-plane consistency for a Forwarding Equivalence Class (FEC) and also Maximum Transmission Unit (MTU) problems. The traceroute functionality may be used to isolate and localize the 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.

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 easily extensible to include additional information needed to support new functionality, by use of Type-Length-Value (TLV) constructs. The usage of TLVs is typically not easy to perform in hardware, and is thus typically handled by the control plane.

LSP Ping supports both asynchronous, as well as, on-demand activation.

3.5. PWE3 Virtual Circuit Connectivity Verification (VCCV)

VCCV, as defined in [VCCV], provides a means for end-to-end fault detection and diagnostics tools to be extended for PWs (regardless of the underlying tunneling technology). The VCCV switching function provides a control channel associated with each PW (based on the PW Associated Channel Header (ACH) which is defined in [PW-ACH]), and allows transmitting the OAM packets in-band with PW data (using CC Type 1: In-band VCCV).

VCCV currently supports the following OAM mechanisms: ICMP Ping, LSP Ping, and BFD. ICMP and LSP Ping are IP encapsulated before being sent over the PW ACH. BFD for 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 pseudo wire (FEC), eliminating the need to exchange Discriminator values.

VCCV consists of two components: (1) signaled component to communicate VCCV capabilities as part of VC label, and (2) 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 negotiation 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.

3.6. IP Performance Metrics (IPPM)

3.6.1. 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].

Alternative protocols for performance measurement are defined, for example, in MPLS-TP OAM ([MPLS-LM-DM], [TP-LM-DM]), and in Ethernet OAM [ITU-T-Y.1731].

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] define a method and protocol for measuring delay and packet loss 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 end points.

TWAMP [TWAMP] is a similar protocol that enables measurement of two-way (round trip) characteristics. TWAMP does not require accurate time of day, and, furthermore, allows the use of a simple session reflector, making it an attractive alternative to OWAMP.

OWAMP and TWAMP use two separate protocols: a Control plane protocol, and a Test plane protocol.

3.6.2. 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 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 pseudo-random). Statistics recorded conform to the relevant IPPM RFCs.

OWAMP and TWAMP test traffic is designed with security in mind. 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.

3.6.3. 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 playing 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 modes. 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 and the Server acknowledges. 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 Fetch-Client sends a fetch request to the server, which responds with an acknowledgement and immediately thereafter the result data.

3.6.4. 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 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 insertion of timestamps.

3.7. MPLS-TP OAM

3.7.1. Overview

The MPLS working group is currently working on defining 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 mechanisms 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 as well as giving overviews of the functionality of the OAM toolset.

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

- MPLS-TP OAM must be able to support both an IP based and 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.

- 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 user-plane ones. 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.

3.7.2. 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 and VCCV mechanisms. However, to address the needs of LSPs as differentiated from PW, the following concepts were defined for [G-ACh]:

- An Associated Channel Header (ACH), that uses a format similar to the PW Control Word, is a 4-byte header that is prepended to OAM packets.

- A Generic Associated 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.

3.7.3. 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 mechanisms and their ability to fulfill the required functionality. The conclusions of this analysis are documented in [OAM-Analysis]. The MPLS working group currently plans to use a mixture of OAM mechanisms that are based on various existing standards, and adapt them to the requirements of [MPLS-TP-OAM]. Some of the main building blocks of this solution are based on:

- Bidirectional Forwarding Detection ([BFD], [BFD-LSP]) for proactive continuity check and connectivity verification.

- LSP Ping as defined in [LSP-Ping] for on-demand connectivity verification.

- New protocol packets, using G-ACH, to address different functionality.

- Performance measurement protocols that are based on the functionality that is described in [ITU-T-Y.1731].

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

3.7.3.1. Continuity Check and Connectivity Verification

Continuity Check and Connectivity Verification are presented in Section 2.2.5 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 [MPLS-TP-CC-CV] while the on-demand tool is defined in [OnDemand-CV].Proactively [MPLS-TP-OAM] states that the function should allow the MEPs to monitor the liveness and connectivity of a transport path. 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 [MPLS-TP-CC-CV] is based on extensions to BFD (see Section 3.3) 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 3.4) for the required behavior of MPLS-TP.

3.7.3.2. Route Tracing

[MPLS-TP-OAM] defines that there is a need for functionality that would allow a path end-point to identify the intermediate and end-points 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 3.4) functionality and is described in [OnDemand-CV].

3.7.3.3. Lock Instruct

The Lock Instruct function is used to notify a transport path end-point 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, Diagnostic testing, to minimize the side-effect on user data traffic.

3.7.3.4. Lock Reporting

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

3.7.3.5. Alarm Reporting

Alarm Reporting is a function used by an intermediate point of a path, that becomes aware of a fault on the path, to report to the end-points 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].

3.7.3.6. Remote Defect Indication

Remote Defect Indication (RDI) is used proactively by a path end-point to report to its peer end-point 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 there a return path exists. [TP-OAM-FW] points out that this function is associated with the proactive CC-V function.

3.7.3.7. 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.

3.7.3.8. Packet Loss Measurement

Packet Loss Measurement is a function used to verify the quality of the service. This function indicates the ratio of packets that are not delivered out of all packets that are transmitted by the path source.

There are two possible ways of determining this measurement:

- 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.

- Sending delimiting messages for 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.

3.7.3.9. Packet Delay Measurement

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 end-points of a path (PW, LSP, or Section). Where:

- One-way packet delay 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.

o Two-way packet delay 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 loop-backed packet by the same source node, when the loopback is performed at the packet's destination node.

Similarly to the packet loss measurement this could be performed in either of the two ways outlined above.

3.8. Summary of OAM Functions

Table 3 summarizes the OAM functions that are supported in each of the standards that were analyzed in this section.

Standard	Continuity Check	Connectivity Verification	Path Discovery	Defect Indications	Performance Monitoring	Other Functions
ICMP Ping		Echo				
Traceroute			Traceroute			
BFD	BFD Control	BFD Echo				
LSP Ping		"Ping" mode	"Traceroute" mode			
IPPM					-Delay measurement -Packet loss measurement	
MPLS-TP OAM	CC	CV/pro-active or on-	Route Tracing	-Alarm Reporting -Client	-LM -DM	-Diagnostic Test

		demand		Failure		-Lock
				Indication		
				-Remote		
				Defect		
				Indication		

Table 3 Summary of OAM Functions

4. Security Considerations

This memo presents an overview of existing OAM mechanisms, and proposes no new OAM mechanisms. Therefore, this document introduces no security considerations. However, the OAM mechanism reviewed in this document can and do present security issues. The reader is encouraged to review the Security Considerations section of each document reference by this memo.

5. IANA Considerations

There are no new IANA considerations implied by this document.

6. Acknowledgments

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

7. References

7.1. Normative References

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

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

[MPLS-OAM-FW] Allan, D., Nadeau, T., "A Framework for Multi-Protocol Label Switching (MPLS) Operations and Management (OAM)", RFC 4378, February 2006.

[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.

[MPLS-TP-OAM] Vigoureux, M., Ward, D., Betts, M., "Requirements for OAM in MPLS Transport Networks", RFC 5860, May 2010.

[G-ACh] Bocci, M., Vigoureux, M., Bryant, S., "MPLS Generic Associated Channel", RFC 5586, June 2009.

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

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

[ICMPv4] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

[ICMPv6] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.

[TCPIP-Tools] Kessler, G., Shepard, S., "A Primer On Internet and TCP/IP Tools and Utilities", RFC 2151, June 1997.

[IPPM-FW] Paxson, V., Almes, G., Mahdavi, J., and Mathis, M., "Framework for IP Performance Metrics", RFC 2330, May 1998.

[IPPM-Con] Mahdavi, J., Paxson, V., "IPPM Metrics for Measuring Connectivity", RFC 2678, September 1999.

[IPPM-1DM] Almes, G., Kalidindi, S., Zekauskas, M., "A One-way Delay Metric for IPPM", RFC 2679, September 1999.

[IPPM-1LM] Almes, G., Kalidindi, S., Zekauskas, M., "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999.

[IPPM-2DM] Almes, G., Kalidindi, S., Zekauskas, M., "A Round-trip Delay Metric for IPPM", RFC 2681, September 1999.

[OWAMP] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and Zekauskas, M., "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006.

[TWAMP] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and Babiarz, J., "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.

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

[BFD-IP] Katz, D., Ward, D., "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, June 2010.

[BFD-Gen] Katz, D., Ward, D., "Generic Application of Bidirectional Forwarding Detection (BFD)", RFC 5882, June 2010.

[BFD-Multi] Katz, D., Ward, D., "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, June 2010.

[BFD-LSP] Aggarwal, R., Kompella, K., Nadeau, T., and Swallow, G., "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010.

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

[TP-OAM-FW] Busi, I., Allan, D., "Operations, Administration and Maintenance Framework for MPLS-based Transport Networks ", RFC 6371, September 2011.

[MPLS-TP-CC-CV] Allan, D., Swallow, G., Drake, J., "Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile", RFC 6428, November 2011.

[OnDemand-CV] Gray, E., Bahadur, N., Boutros, S., Aggarwal, R. "MPLS On-Demand Connectivity Verification and Route Tracing", RFC 6426, November 2011.

[MPLS-LM-DM] Frost, D., Bryant, S., "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011.

[TP-LM-DM] Frost, D., Bryant, S., "A Packet Loss and Delay Measurement Profile for MPLS-Based Transport Networks", RFC 6375, September 2011.

[MPLS-TP-Fault] Swallow, G., Fulignoli, A., Vigoureux, M., Boutros, S., "MPLS Fault Management Operations, Administration, and Maintenance (OAM)", RFC 6427, November 2011.

[TP-Lock-Loop] Boutros, S., Sivabalan, S., Aggarwal, R., Vigoureux, M., Dai, X., "MPLS Transport Profile Lock Instruct and Loopback Functions", RFC 6435, November 2011.

7.2. Informative References

[OAM-Def] Andersson, L., Van Helvoort, H., Bonica, R., Romascanu, D., Mansfield, S., "Guidelines for the use of the OAM acronym in the IETF ", RFC 6291, June 2011.

[OAM-Analysis]Sprecher, N., Fang, L., "An Overview of the OAM Tool Set for MPLS based Transport Networks", RFC 6669, July 2012.

[MPLS-TP-Term]Van Helvoort, H., Andersson, L., Sprecher, N., "A Thesaurus for the Terminology used in Multiprotocol Label Switching Transport Profile (MPLS-TP) drafts/RFCs and ITU-T's Transport Network Recommendations", work-in-progress, draft-ietf-mpls-tp-rosetta-stone, January 2012.

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

[ITU-T-Y.1731]"OAM Functions and Mechanisms for Ethernet-based Networks", February 2008.

[ITU-T-Y.1711]"Operation & Maintenance mechanism for MPLS networks", February 2004.

[IEEE-802.3ah]"Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks", clause 57, September 2004.

[ITU-T-G.806] "Characteristics of transport equipment - Description methodology and generic functionality", January, 2009.

Authors' Addresses

Tal Mizrahi
Marvell
6 Hamada St.
Yokneam, 20692
Israel

Email: talmi@marvell.com

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

Email: nurit.sprecher@nsn.com

Elisa Bellagamba
Ericsson
6 Farogatan St.
Stockholm, 164 40
Sweden

Phone: +46 761440785
Email: elisa.bellagamba@ericsson.com

Yaacov Weingarten
34 Hagefen St.
Karnei Shomron, 4485500
Israel

Email: wyaacov@gmail.com