Internet Engineering Task Force (IETF)
Request for Comments: 6669
Category: Informational
ISSN: 2070-1721

N. Sprecher
Nokia Siemens Networks
L. Fang
Cisco Systems

July 2012

An Overview of the Operations, Administration, and Maintenance (OAM) Toolset for MPLS-Based Transport Networks

Abstract

This document provides an overview of the Operations, Administration, and Maintenance (OAM) toolset for MPLS-based transport networks. The toolset consists of a comprehensive set of fault management and performance monitoring capabilities (operating in the data plane) that are appropriate for transport networks as required in RFC 5860 and support the network and services at different nested levels. This overview includes a brief recap of the MPLS Transport Profile (MPLS-TP) OAM requirements and functions and the generic mechanisms created in the MPLS data plane that allow the OAM packets to run in-band and share their fate with data packets. The protocol definitions for each of the MPLS-TP OAM tools are defined in separate documents (RFCs or Working Group documents), which are referenced by this document.

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/rfc6669.

Table of Contents

1. Introduction

1.1. Scope

The MPLS Transport Profile (MPLS-TP) architectural framework is defined in [RFC5921], and it describes a common set of protocol functions that supports the operational models and capabilities typical of such transport networks.

Operations, Administration, and Maintenance (OAM) plays a significant role in carrier networks. It provides methods for fault management and performance monitoring in both the transport and service layers, in order to improve their ability to support services with guaranteed and strict Service Level Agreements (SLAs) while reducing their operational costs.

[RFC5654], in general, and [RFC5860], in particular, define a set of requirements for the OAM functionality for MPLS-TP Label Switched Paths (LSPs), Pseudowires (PWs), and Sections.

The OAM solution, developed by the joint IETF and ITU-T MPLS-TP project, has three objectives:

- The OAM toolset should be developed based on existing MPLS architecture, technology, and toolsets.

- The OAM operational experience should be similar to that in other transport networks.

- The OAM toolset developed for MPLS-based transport networks needs to be fully interoperable with existing MPLS OAM tools as documented in Section 2.1.5. of [RFC5860].

The MPLS-TP OAM toolset is based on the following existing tools:

- LSP ping, as defined in [RFC4379].

- Bidirectional Forwarding Detection (BFD), as defined in [RFC5880] and refined in [RFC5884].

- ITU-T OAM for the Ethernet toolset, as defined in [Y.1731]. This has been used as functionality guidelines for the performance measurement tools that were not previously supported in MPLS.

Note that certain extensions and adjustments have been specified, relative to the existing MPLS tools, in order to conform to the transport environment and the requirements of MPLS-TP. However, compatibility with the existing MPLS tools has been maintained.

This document provides an overview of the MPLS-TP OAM toolset, which consists of tools for MPLS-TP fault management and performance monitoring. This overview includes a brief recap of MPLS-TP OAM requirements, their functions, and the generic mechanisms used to support the MPLS-TP OAM operation.

The protocol definitions for individual MPLS-TP OAM tools are specified in separate RFCs (or Working Group documents), which are referenced by this document.

In addition, this document includes a table that cross-references the solution documents of the OAM functionality supported. Finally, the document presents the applicability and utilization of each tool in the MPLS-TP OAM toolset.

1.2. Acronyms

This document uses the following acronyms:

ACH	Associated Channel Header
AIS	Alarm Indication Signal
BFD	Bidirectional Forwarding Detection
CC-CV	Continuity Check and Connectivity Verification
DM	Delay Measurement
FM	Fault Management
G-ACh	Generic Associated Channel
GAL	G-ACh Label
GMPLS	Generalized Multiprotocol Label Switching
IANA	Internet Assigned Numbers Authority
LDI	Link Down Indication
LKR	Lock Report
LM	Loss Measurement
LOC	Loss of Continuity
LSP	Label Switched Path
MEP	Maintenance Entity Group End Point
MEG	Maintenance Entity Group
MIP	Maintenance Entity Group Intermediate Point
MPLS	Multiprotocol Label Switching
MPLS-TP	Transport Profile for MPLS
OAM	Operations, Administration, and Maintenance
PM	Performance Monitoring
PW	Pseudowire
RDI	Remote Defect Indication
SLA	Service Level Agreement
TLV	Type, Length, Value
VCCV	Virtual Circuit Connectivity Verification

2. Basic OAM Infrastructure Functionality

[RFC5860] defines a set of requirements for OAM architecture and general principles of operations, which are evaluated below:

[RFC5860] requires that --

- OAM mechanisms in MPLS-TP are independent of the transmission media and the client service being emulated by the PW ([RFC5860], Section 2.1.2).

- 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 it must be possible to choose to make use of IP capabilities. On the other hand, in environments where IP functionality is not available, the OAM tools must still be able to operate independent of IP forwarding and routing ([RFC5860], Section 2.1.4). It is required to have OAM interoperability between distinct domains materializing the environments ([RFC5860], Section 2.1.5).

- All OAM protocols support identification information, at least in the form of IP addressing structure, and are extensible to support additional identification schemes ([RFC5860], Section 2.1.4).

- OAM packets and the user traffic are congruent (i.e., OAM packets are transmitted in-band) and there is a need to differentiate OAM packets from user-plane packets [RFC5860], Section 2.1.3. Inherent in this requirement is the principle that full operation of the MPLS-TP OAM must be possible independently of the control or management plane used to operate the network [RFC5860], Section 2.1.3.

- MPLS-TP OAM supports point-to-point bidirectional PWs, point-to-point co-routed bidirectional LSPs, and point-to-point bidirectional Sections ([RFC5860], Section 2.1.1). The applicability of particular MPLS-TP OAM functions to point-to-point associated bidirectional LSPs, point-to-point unidirectional LSPs, and point-to-multipoint LSPs, is described in [RFC5860], Section 2.2. In addition, MPLS-TP OAM supports these LSPs and PWs when they span either single or multiple domains ([RFC5860], Section 2.1.1).

- OAM packets may be directed to an intermediate point of an LSP/PW ([RFC5860], Sections 2.2.3, 2.2.4, and 2.2.5).

[RFC5860], Section 2.2 recommends that any protocol solution meeting one or more functional requirement(s) be the same for PWs, LSPs, and Sections.

The following document set addresses the basic requirements listed above:

- [RFC6371] describes the architectural framework for conformance to the basic requirements listed above. It also defines the basic relationships between the MPLS structures, e.g., LSP, PW, and the structures necessary for OAM functionality, i.e., the Maintenance Entity Group (MEG), its end points, and intermediate points.

- [RFC5586] specifies the use of the MPLS-TP in-band control channels. It generalizes the applicability of the PW ACH to MPLS LSPs and Sections by defining a Generic Associated Channel (G-ACh). The G-ACh allows control packets to be multiplexed transparently over LSPs and Sections similar to that of PW VCCV [RFC5085]. The Generic Association Label (GAL) is defined by assigning a reserved MPLS label value and is used to identify the OAM control packets. The value of the ACH Channel Type field indicates the specific protocol carried on the associated control channel. Each MPLS-TP OAM protocol has an IANA-assigned channel type allocated to it.

[RFC5085] defines an Associated Channel Header (ACH) that provides a PW associated control channel between a PW's end points, over which OAM and other control messages can be exchanged. [RFC5586] generalizes the PW Associated Channel Header (ACH) to provide common in-band control channels also at the LSP and MPLS-TP link levels. The G-ACh allows control packets to be multiplexed transparently over the same LSP or MPLS-TP link as in PW VCCV. Multiple control channels can exist between end points.

[RFC5085] also defines a label-based exception mechanism that helps a Label Switching Router (LSR) to identify the control packets and direct them to the appropriate entity for processing. The use of G-ACh and GAL provides the necessary mechanisms to allow OAM packets to run in-band and share their fate with data packets. It is expected that all of the OAM protocols will be used in conjunction with this Generic Associated Channel.

- [RFC6370] provides an IP-based identifier set for MPLS-TP that can be used to identify the transport entities in the network and referenced by the different OAM protocols.

Note: [MPLS-TP-ITU-Idents] augments that set of identifiers to include identifier information in a format used by the ITU-T. Other identifier sets may be defined as well.

3. MPLS-TP OAM Functions

The following sections discuss the OAM functions that are required in [RFC5860] and expanded upon in [RFC6371].

3.1. Continuity Check and Connectivity Verification

Continuity Check and Connectivity Verification (CC-CV) are OAM operations generally used in tandem and complement each other. These functions are generally run proactively, but may also be used on-demand for diagnoses of a specific condition. [RFC5860] states that the function should allow the MEPs to proactively monitor the liveliness and connectivity of a transport path (LSP, PW, or a Section) between them. In on-demand mode, this function should support monitoring between the MEPs and between a MEP and MIP. Note that as specified in [RFC6371], Sections 3.3 and 3.4, a MEP and a MIP can reside in an unspecified location within a node, or in a particular interface on a specific side of the forwarding engine.

[RFC6371] highlights the need for the CC-CV messages to include unique identification of the MEG that is being monitored and the MEP that originated the message. The function, both proactively and in on-demand mode, needs to be transmitted at regular transmission rates pre-configured by the operator.

3.1.1. Documents for CC-CV Tools

[RFC6428] defines BFD extensions to support proactive CC-CV applications.

[RFC6426] provides LSP ping extensions that are used to implement on-demand connectivity verification.

Both of these tools will be used within the basic functionality framework described in Section 2.

3.2. Remote Defect Indication

Remote Defect Indication (RDI) is used by a path end point to report that a defect is detected on a bidirectional connection to its peer end point. [RFC5860] points out that this function may be applied to a unidirectional LSP only if a return path exists. [RFC6371] points out that this function is associated with the proactive CC-CV function.

3.2.1. Documents for RDI

[RFC6428] provides an extension for BFD that includes the RDI indication in the BFD format and a specification of how this indication is to be used.

3.3. Route Tracing

[RFC5860] defines the need for functionality that would allow a path end point to identify the intermediate points (if any) and end point(s) along the path (LSP, PW, or a Section). 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.

3.3.1. Documents for Route Tracing

[RFC6426] specifies that the LSP ping enhancements for MPLS-TP on-demand connectivity verification include information on the use of LSP ping for route tracing of an MPLS-TP path.

3.4. Alarm Reporting

Alarm Reporting is a function used by an intermediate point of a path (LSP or PW) to report to the end points of the path that a fault exists on the path. [RFC6371] states that this may occur as a result of a defect condition discovered at a server layer. The intermediate point generates an Alarm Indication Signal (AIS) that continues until the fault is cleared. The consequent action of this function is detailed in [RFC6371].

3.4.1. Documents for Alarm Reporting

MPLS-TP defines a new protocol to address this functionality that is documented in [RFC6427]. This protocol uses all of the basic mechanisms detailed in Section 2.

3.5. Lock Instruct

The Lock Instruct function is an administrative control tool that allows a path end point to instruct its peer end point to lock the path (LSP, PW, or Section). The tool is necessary to support single-side provisioning for administrative locking, according to [RFC6371]. This function is used on-demand.

3.5.1. Documents for Lock Instruct

[RFC6435] describes the details of a new ACH-based protocol format for this functionality.

3.6. Lock Reporting

Lock Reporting, defined in [RFC5860], is similar to the Alarm Reporting function described above. It is used by an intermediate point to notify the end points of a transport path (LSP or PW) that an administrative lock condition exists for the transport path.

3.6.1. Documents for Lock Reporting

MPLS-TP defines a new protocol to address this functionality that is documented in [RFC6427]. This protocol uses all the basic mechanisms detailed in Section 2.

3.7. Diagnostic

[RFC5860] indicates a need to provide an OAM function that would enable conducting different diagnostic tests on a PW, LSP, or Section. [RFC6371] provides two types of specific tests to be used through this functionality:

- Throughput estimation - allowing the provider to verify the bandwidth/throughput of a transport path. This is an out-of-service tool that uses special packets of varying sizes to test the actual bandwidth and/or throughput of the path.

- Data-plane loopback - this out-of-service tool causes all traffic that reaches the target node, either a MEP or MIP, to be looped back to the originating MEP. For targeting MIPs, a co-routed bidirectional path is required.

3.7.1. Documents for Diagnostic Testing

[RFC6435] describes the details of a new ACH-based protocol format for the data-plane loopback functionality.

The tool for throughput estimation is under study.

3.8. Packet Loss Measurement

Packet Loss Measurement is required by [RFC5860] to provide a quantification of the packet loss ratio on a transport path. This is the ratio of the number of user packets lost to the total number of user packets during a defined time interval. To employ this

function, [RFC6371] defines that the two end points of the transport path should exchange counters of messages transmitted and received within a time period bounded by loss-measurement messages. The framework warns that there may be small errors in the computation, which result from various issues.

3.8.1. Documents for Packet Loss Measurement

[RFC6374] describes the protocol formats and procedures for using the tool and enabling efficient and accurate measurement of packet loss, delay, and throughput in MPLS networks. [RFC6375] describes a profile of the general MPLS loss, delay, and throughput measurement techniques that suffice to meet the specific requirements of MPLS-TP. Note that the tool logic is based on the behavior of the parallel function described in [Y.1731].

3.9. Packet Delay Measurement

Packet Delay Measurement is a function that is used to measure the one-way or two-way delay of packet transmission between a pair of the end points of a path (PW, LSP, or Section), as described in [RFC5860], 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.

- 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.

[RFC6371] describes how the tool could be used (both in proactive and on-demand modes) for either one-way or two-way measurement. However, it warns that the one-way delay option requires precise time synchronization between the end points.

3.9.1. Documents for Delay Measurement

[RFC6374] describes the protocol formats and procedures for using the tool and enabling efficient and accurate measurement of packet loss, delay, and throughput in MPLS networks. [RFC6375] describes a profile of the general MPLS loss, delay, and throughput measurement techniques that suffices to meet the specific requirements of MPLS-TP. Note that the tool logic is based on the behavior of the parallel function described in [Y.1731].

4. MPLS-TP OAM Documents Guide

The complete MPLS-TP OAM protocol suite is covered by a small set of existing IETF documents. This set of documents may be expanded in the future to cover additional OAM functionality. In order to allow the reader to understand this set of documents, a cross-reference of the existing documents (RFCs or Working Group documents) for the initial phase of the specification of MPLS-based transport networks is provided below.

[RFC5586] provides a specification of the basic structure of protocol messages for in-band data-plane OAM in an MPLS environment.

[RFC6370] provides definitions of different formats that may be used within OAM protocol messages to identify the network elements of an MPLS-based transport network.

The following table (Table 1) provides the summary of proactive MPLS-TP OAM Fault Management toolset functions, the associated tool/protocol, and the corresponding RFCs in which they are defined.

OAM Functions	OAM Tools/Protocols	RFCs
Continuity Check and Connectivity Verification	Bidirectional Forwarding Detection (BFD)	[RFC6428]
Remote Defect Indication (RDI)	Flag in Bidirectional Forwarding Detection (BFD) message	[RFC6428]
Alarm Indication Signal (AIS)	G-ACh-based AIS message	[RFC6427]
Link Down Indication (LDI)	Flag in AIS message	[RFC6427]
Lock Reporting (LKR)	G-ACh-based LKR message	[RFC6427]

Table 1. Proactive Fault Management OAM Toolset

The following table (Table 2) provides an overview of the on-demand MPLS-TP OAM Fault Management toolset functions, the associated tool/protocol, and the corresponding RFCs in which they are defined.

OAM Functions	OAM Tools/Protocols	RFCs
Connectivity Verification	LSP Ping	[RFC6426]
Lock Instruct (LI)	(1) G-ACh-based Loopback, (2) Lock Instruct (LI)	[RFC6426]
Lock Report (LKR)	Flag in AIS message	[RFC6426]

Table 2. On Demand Fault Management OAM Toolset

The following table (Table 3) provides the Performance Monitoring Functions, the associated tool/protocol definitions, and the corresponding RFCs in which they are defined.

OAM Functions	OAM Tools/Protocols	RFCs
Packet Loss Measurement (LM)	G-ACh-based LM & DM query messages	[RFC6374] [RFC6375]
Packet Delay Measurement (DM)	G-ACh-based LM & DM query messages	[RFC6374] [RFC6375]
Throughput Measurement	derived from Loss Measurement	[RFC6374] [RFC6375]
Delay Variation Measurement	derived from Delay Measurement	[RFC6374] [RFC6375]

Table 3. Performance Monitoring OAM Toolset

5. OAM Toolset Applicability and Utilization

The following subsections present the MPLS-TP OAM toolset from the perspective of the specified protocols and identifies the required functionality that is supported by the particular protocol.

5.1. Connectivity Check and Connectivity Verification

Proactive Continuity Check and Connectivity Verification (CC-CV) functions are used to detect loss of continuity (LOC) and unintended connectivity between two MEPs. Loss of connectivity, mis-merging, mis-connectivity, or unexpected Maintenance Entity Group End Points (MEPs) can be detected using the CC-CV tools. See Sections 3.1, 3.2, 3.3 in this document for CC-CV protocol references.

The CC-CV tools are used to support MPLS-TP fault management, performance management, and protection switching. Proactive CC-CV control packets are sent by the source MEP to the sink MEP. The sink-MEP monitors the arrival of the CC-CV control packets and detects the defect. For bidirectional transport paths, the CC-CV protocol is usually transmitted simultaneously in both directions.

The transmission interval of the CC-CV control packets can be configured. For example:

- 3.3 ms is the default interval for protection switching.
- 100 ms is the default interval for performance monitoring.
- 1 s is the default interval for fault management.

5.2. Diagnostic Tests and Lock Instruct

[RFC6435] describes a protocol that provides a mechanism to Lock and Unlock traffic (e.g., data and control traffic or specific OAM traffic) at a specific LSR on the path of the MPLS-TP LSP to allow loopback of the traffic to the source.

These diagnostic functions apply to associated bidirectional MPLS-TP LSPs, including MPLS-TP LSPs, bidirectional RSVP-Traffic Engineering (RSVP-TE) tunnels (which is relevant for the MPLS-TP dynamic control-plane option with GMPLS), and single-segment and multi-segment Pseudowires. [RFC6435] provides the protocol definition for diagnostic tests functions.

[RFC6435] defines a mechanism where a lock instruction is sent by a management application to both ends of a point-to-point LSP, requesting them to take the LSP out-of-service. When an end point gets the management request, it locks the LSP and sends a Lock Instruct message to the other end of the LSP. The Lock Instruct message is carried in a Generic ACH message and is sent periodically. The time between successive messages is no longer than the value set in the Refresh Timer field of the Lock Instruct message. An LSP end point keeps the LSP locked while it is either receiving the periodic

Lock Instruct messages or has an in-force lock instruction from the management application.

Note that since the management application will receive a management plane response from both ends of the LSP confirming that the LSP has been locked, there is no requirement for the Lock Instruct message to have a response. Therefore, [RFC6435] does not define a Lock Instruct response message.

The loopback operations include:

- Lock: take an LSP out of service for maintenance.
- Unlock: Restore a previously locked LSP to service.
- Set_Full_Loopback and Set_OAM_Loopback.
- Unset_Full_Loopback and Set_OAM_Loopback.

Operators can use the loopback mode to test the connectivity or performance (loss, delay, delay variation, and throughput) of a given LSP up to a specific node on the path of the LSP.

5.3. Lock Reporting

The Lock Report (LKR) function is used to communicate to the MEPS of the client (sub-)layer MEPs the administrative locking of a server (sub-)layer MEP, and consequential interruption of data traffic forwarding in the client layer. See Section 3.6 in this document for Lock Reporting protocol references.

When an operator is taking the LSP out of service for maintenance or another operational reason, using the LKR function can help to distinguish the condition as administrative locking from a defect condition.

The Lock Report function may also serve the purpose of alarm suppression in the MPLS-TP network above the level at which the Lock has occurred. The receipt of an LKR message may be treated as the equivalent of the loss of continuity at the client layer.

5.4. Alarm Reporting and Link Down Indication

Alarm Indication Signal (AIS) message is used to suppress alarms following detection of defect conditions at the server (sub-)layer. When the Link Down Indication (LDI) is set, the AIS message may be used to trigger recovery mechanisms.

When a server MEP detects the failure, it asserts LOC or signal fail, which sets the flag up to generate an OAM packet with the AIS message. The AIS message is forwarded to the downstream sink MEP in the client layer. This enables the client layer to suppress the generation of secondary alarms.

An LDI flag is defined in the AIS message. The LDI flag is set in the AIS message in response to detecting a fatal failure in the server layer. Receipt of an AIS message with this flag set may be interpreted by a MEP as an indication of signal fail at the client layer.

The protocols for AIS and LDI are defined in [RFC6427].

Fault OAM messages are generated by intermediate nodes where an LSP is switched and propagated to the end points (MEPs).

From a practical point of view, when both proactive Continuity Check functions and LDI are used, one may consider running the proactive Continuity Check functions at a slower rate (e.g., longer BFD hello intervals), and reply on LDI to trigger fast protection switch over upon failure detection in a given LSP.

5.5. Remote Defect Indication

The Remote Defect Indication (RDI) function enables an end point to report to its peer end point that a fault or defect condition is detected on the PW, LSP, or Section.

The RDI OAM function is supported by the use of BFD control packets [RFC6428]. RDI is only used for bidirectional connections and is associated with proactive CC-CV activation.

When an end point (MEP) detects a signal failure condition, it sets the flag up by setting the diagnostic field of the BFD control packet to a particular value to indicate the failure condition on the associated PW, LSP, or Section. Additionally, the BFD control packet is transmitted with the failure flag up to the other end point (its peer MEP).

The RDI function can be used to facilitate protection switching by synchronizing the two end points when unidirectional failure occurs and is detected by one end.

5.6. Packet Loss and Delay Measurement

The packet loss and delay measurement toolset enables operators to measure the quality of the packet transmission over a PW, LSP, or Section. Section 3.8 in this document defines the protocols for packet loss measurement, and Section 3.9 defines the protocols for packet delay measurement.

The loss and delay protocols have the following characteristics and capabilities:

- They support the measurement of packet loss, delay, and throughput over Label Switched Paths (LSPs), Pseudowires, and MPLS Sections.
- The same LM and DM protocols can be used for both continuous/proactive and selective/on-demand measurements.
- The LM and DM protocols use a simple query/response model for bidirectional measurement that allows a single node -- the querier -- to measure the loss or delay in both directions.
- The LM and DM protocols use query messages for unidirectional loss and delay measurement. The measurement can either be carried out at the downstream node(s), or at the querier if an out-of-band return path is available.
- The LM and DM protocols do not require that the transmit-and-receive interfaces be the same when performing bidirectional measurement.
- The LM supports test-message-based measurement (i.e., inferred mode) as well as measurement based on data-plane counters (i.e., direct mode).
- The LM protocol supports both 32-bit and 64-bit counters.
- The LM protocol supports measurement in terms of both packet counts and octet counts; although for simplicity, only packet counters are currently included in the MPLS-TP profile.
- The LM protocol can be used to measure channel throughput as well as packet loss.
- The DM protocol supports varying the measurement message size in order to measure delays associated with different packet sizes.
- The DM protocol uses IEEE 1588 timestamps [IEEE1588] by default but also supports other timestamp formats, such as NTP.

6. Security Considerations

This document, as an overview of MPLS OAM tools, does not by itself raise any particular security considerations.

The general security considerations are provided in [RFC5920] and [MPLS-TP-SEC]. Security considerations for each function within the OAM toolset have been recorded in each document that specifies a particular functionality.

In general, OAM is always an area where the security risk is high. For example, confidential information may be intercepted by attackers to gain access to networks; therefore, authentication, authorization, and encryption must be enforced to prevent security breaches.

It is also important to strictly follow operational security procedures. For example, in the case of MPLS-TP static provisioning, the operator interacts directly with the Network Management System (NMS) and devices, and it is critical in order to prevent human errors and malicious attacks.

Since MPLS-TP OAM uses G-ACh, the security risks and mitigations described in [RFC5085] also apply here. In short, messages on the G-ACh could be intercepted, or false G-ACh packets could be inserted.

Additionally, DoS attacks can be mounted by flooding G-ACh messages to peer devices. To mitigate this type of attack, throttling mechanisms or rate limits can be used. For more details, please see [RFC5085].

7. Acknowledgements

The authors would like to thank the MPLS-TP experts from both the IETF and ITU-T for their helpful comments. In particular, we would like to thank Loa Andersson and the Area Directors for their suggestions and enhancements to the text.

Thanks to Tom Petch for useful comments and discussions.

Thanks to Rui Costa for his review and comments, which helped improve this document.

8. References

8.1. Normative References

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

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

[RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, June 2009.

[RFC5654] Niven-Jenkins, B., Ed., Brungard, D., Ed., Betts, M., Ed., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, September 2009.

[RFC5860] 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.

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

[RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010.

[RFC5921] Bocci, M., Ed., Bryant, S., Ed., Frost, D., Ed., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, July 2010.

[RFC6370] Bocci, M., Swallow, G., and E. Gray, "MPLS Transport Profile (MPLS-TP) Identifiers", RFC 6370, September 2011.

[RFC6371] Busi, I., Ed., and D. Allan, Ed., "Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks", RFC 6371, September 2011.

[RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011.

[RFC6375] Frost, D., Ed., and S. Bryant, Ed., "A Packet Loss and Delay Measurement Profile for MPLS-Based Transport Networks", RFC 6375, September 2011.

[RFC6426] Gray, E., Bahadur, N., Boutros, S., and R. Aggarwal, "MPLS On-Demand Connectivity Verification and Route Tracing", RFC 6426, November 2011.

[RFC6427] 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.

[RFC6428] 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.

[RFC6435] 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.

8.2. Informative References

[IEEE1588] IEEE, "1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", March 2008.

[MPLS-TP-ITU-Idents] Winter, R., van Helvoort, H., and M. Betts, "MPLS-TP Identifiers Following ITU-T Conventions", Work in Progress, March 2012.

[MPLS-TP-SEC] Fang, L., Niven-Jenkins, B., and S. Mansfield, "MPLS-TP Security Framework", Work in Progress, March 2012.

[RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

[Y.1731] International Telecommunications Union - Standardization, "OAM functions and mechanisms for Ethernet based networks", ITU Y.1731, May 2006.

Contributors

Elisa Bellagamba	Ericsson
Yaacov Weingarten	Nokia Siemens Networks
Dan Frost	Cisco
Nabil Bitar	Verizon
Raymond Zhang	Alcatel Lucent
Lei Wang	Telenor
Kam Lee Yap	XO Communications
John Drake	Juniper
Yaakov Stein	RAD
Anamaria Fulignoli	Ericsson
Italo Busi	Alcatel Lucent
Huub van Helvoort	Huawei
Thomas Nadeau	Computer Associate
Henry Yu	TW Telecom
Mach Chen	Huawei
Manuel Paul	Deutsche Telekom

Authors' Addresses

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

EMail: nurit.sprecher@nsn.com

Luyuan Fang
Cisco Systems
111 Wood Avenue South
Iselin, NJ 08830
USA

EMail: lufang@cisco.com