

Internet Engineering Task Force (IETF)
Request for Comments: 5951
Category: Standards Track
ISSN: 2070-1721

K. Lam
Alcatel-Lucent
S. Mansfield
E. Gray
Ericsson
September 2010

Network Management Requirements for MPLS-based Transport Networks

Abstract

This document specifies the requirements for the management of equipment used in networks supporting an MPLS Transport Profile (MPLS-TP). The requirements are defined for specification of network management aspects of protocol mechanisms and procedures that constitute the building blocks out of which the MPLS Transport Profile is constructed. That is, these requirements indicate what management capabilities need to be available in MPLS for use in managing the MPLS-TP. This document is intended to identify essential network management capabilities, not to specify what functions any particular MPLS implementation supports.

Status of This Memo

This is an Internet Standards Track document.

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5951>.

Copyright Notice

Copyright (c) 2010 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. Terminology	5
2. Management Interface Requirements	7
3. Management Communication Channel (MCC) Requirements	7
4. Management Communication Network (MCN) Requirements	7
5. Fault Management Requirements	9
5.1. Supervision Function	9
5.2. Validation Function	10
5.3. Alarm Handling Function	11
5.3.1. Alarm Severity Assignment	11
5.3.2. Alarm Suppression	11
5.3.3. Alarm Reporting	11
5.3.4. Alarm Reporting Control	12
6. Configuration Management Requirements	12
6.1. System Configuration	12
6.2. Control Plane Configuration	13
6.3. Path Configuration	13
6.4. Protection Configuration	14
6.5. OAM Configuration	14
7. Performance Management Requirements	15
7.1. Path Characterization Performance Metrics	15
7.2. Performance Measurement Instrumentation	16
7.2.1. Measurement Frequency	16
7.2.2. Measurement Scope	17
8. Security Management Requirements	17
8.1. Management Communication Channel Security	17
8.2. Signaling Communication Channel Security	18
8.3. Distributed Denial of Service	18
9. Security Considerations	19
10. Acknowledgments	19
11. References	19
11.1. Normative References	19
11.2. Informative References	20
Appendix A. Communication Channel (CCh) Examples.....	22
Contributor's Address	24

1. Introduction

This document specifies the requirements for the management of equipment used in networks supporting an MPLS Transport Profile (MPLS-TP). The requirements are defined for specification of network management aspects of protocol mechanisms and procedures that constitute the building blocks out of which the MPLS Transport Profile is constructed. That is, these requirements indicate what management capabilities need to be available in MPLS for use in managing the MPLS-TP. This document is intended to identify essential network management capabilities, not to specify what functions any particular MPLS implementation supports.

This document also leverages management requirements specified in ITU-T G.7710/Y.1701 [1] and RFC 4377 [2], and attempts to comply with the guidelines defined in RFC 5706 [15].

ITU-T G.7710/Y.1701 defines generic management requirements for transport networks. RFC 4377 specifies the operations and management requirements, including operations-and-management-related network management requirements, for MPLS networks.

This document is a product of a joint ITU-T and IETF effort to include an MPLS Transport Profile (MPLS-TP) within the IETF MPLS and Pseudowire Emulation Edge-to-Edge (PWE3) architectures to support capabilities and functionality of a transport network as defined by the ITU-T.

The requirements in this document derive from two sources:

- 1) MPLS and PWE3 architectures as defined by the IETF, and
- 2) packet transport networks as defined by the ITU-T.

Requirements for management of equipment in MPLS-TP networks are defined herein. Related functions of MPLS and PWE3 are defined elsewhere (and are out of scope in this document).

This document expands on the requirements in ITU-T G.7710/Y.1701 [1] and RFC 4377 [2] to cover fault, configuration, performance, and security management for MPLS-TP networks, and the requirements for object and information models needed to manage MPLS-TP networks and network elements.

In writing this document, the authors assume the reader is familiar with RFCs 5921 [8] and 5950 [9].

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [5]. Although this document is not a protocol specification, the use of this language clarifies the instructions to protocol designers producing solutions that satisfy the requirements set out in this document.

Anomaly: The smallest discrepancy that can be observed between actual and desired characteristics of an item. The occurrence of a single anomaly does not constitute an interruption in ability to perform a required function. Anomalies are used as the input for the Performance Monitoring (PM) process and for detection of defects (from [21], Section 3.7).

Communication Channel (CCh): A logical channel between network elements (NEs) that can be used (for example) for management or control plane applications. The physical channel supporting the CCh is technology specific. See Appendix A.

Data Communication Network (DCN): A network that supports Layer 1 (physical layer), Layer 2 (data-link layer), and Layer 3 (network layer) functionality for distributed management communications related to the management plane, for distributed signaling communications related to the control plane, and other operations communications (e.g., order-wire/voice communications, software downloads, etc.).

Defect: The density of anomalies has reached a level where the ability to perform a required function has been interrupted. Defects are used as input for performance monitoring, the control of consequent actions, and the determination of fault cause (from [21], Section 3.24).

Failure: The fault cause persisted long enough to consider the ability of an item to perform a required function to be terminated. The item may be considered as failed; a fault has now been detected (from [21], Section 3.25).

Fault: A fault is the inability of a function to perform a required action. This does not include an inability due to preventive maintenance, lack of external resources, or planned actions (from [21], Section 3.26).

Fault Cause: A single disturbance or fault may lead to the detection of multiple defects. A fault cause is the result of a correlation process that is intended to identify the defect that is representative of the disturbance or fault that is causing the problem (from [21], Section 3.27).

Fault Cause Indication (FCI): An indication of a fault cause.

Management Communication Channel (MCC): A CCh dedicated for management plane communications.

Management Communication Network (MCN): A DCN supporting management plane communication is referred to as a Management Communication Network (MCN).

MPLS-TP NE: A network element (NE) that supports the functions of MPLS necessary to participate in an MPLS-TP based transport service. See RFC 5645 [7] for further information on functionality required to support MPLS-TP.

MPLS-TP network: a network in which MPLS-TP NEs are deployed.

Operations, Administration and Maintenance (OAM), On-Demand and Proactive: One feature of OAM that is largely a management issue is control of OAM; on-demand and proactive are modes of OAM mechanism operation defined in (for example) Y.1731 ([22] - Sections 3.45 and 3.44, respectively) as:

- o On-demand OAM - OAM actions that are initiated via manual intervention for a limited time to carry out diagnostics. On-demand OAM can result in singular or periodic OAM actions during the diagnostic time interval.
- o Proactive OAM - OAM actions that are carried on continuously to permit timely reporting of fault and/or performance status.

(Note that it is possible for specific OAM mechanisms to only have a sensible use in either on-demand or proactive mode.)

Operations System (OS): A system that performs the functions that support processing of information related to operations, administration, maintenance, and provisioning (OAM&P) for the networks, including surveillance and testing functions to support customer access maintenance.

Signaling Communication Channel (SCC): A CCh dedicated for control plane communications. The SCC can be used for GMPLS/ASON signaling and/or other control plane messages (e.g., routing messages).

Signaling Communication Network (SCN): A DCN supporting control plane communication is referred to as a Signaling Communication Network (SCN).

2. Management Interface Requirements

This document does not specify a preferred management interface protocol to be used as the standard protocol for managing MPLS-TP networks. Managing an end-to-end connection across multiple operator domains where one domain is managed (for example) via NETCONF [16] or SNMP [17], and another domain via CORBA [18], is allowed.

- 1) For the management interface to the management system, an MPLS-TP NE MAY actively support more than one management protocol in any given deployment.

For example, an operator can use one protocol for configuration of an MPLS-TP NE and another for monitoring. The protocols to be supported are at the discretion of the operator.

3. Management Communication Channel (MCC) Requirements

- 1) Specifications SHOULD define support for management connectivity with remote MPLS-TP domains and NEs, as well as with termination points located in NEs under the control of a third party network operator. See ITU-T G.8601 [23] for example scenarios in multi-carrier, multi-transport technology environments.
- 2) For management purposes, every MPLS-TP NE MUST connect to an OS. The connection MAY be direct (e.g., via a software, hardware, or proprietary protocol connection) or indirect (via another MPLS-TP NE). In this document, any management connection that is not via another MPLS-TP NE is a direct management connection. When an MPLS-TP NE is connected indirectly to an OS, an MCC MUST be supported between that MPLS-TP NE and any MPLS-TP NE(s) used to provide the connection to an OS.

4. Management Communication Network (MCN) Requirements

Entities of the MPLS-TP management plane communicate via a DCN, or more specifically via the MCN. The MCN connects management systems with management systems, management systems with MPLS-TP NEs, and (in the indirect connectivity case discussed in section 3) MPLS-TP NEs with MPLS-TP NEs.

RFC 5586 [14] defines a Generic Associated Channel (G-ACh) to enable the realization of a communication channel (CCh) between adjacent MPLS-TP NEs for management and control. RFC 5718 [10] describes how the G-ACh can be used to provide infrastructure that forms part of the MCN and SCN. It also explains how MCN and SCN messages are encapsulated, carried on the G-ACh, and decapsulated for delivery to management or signaling/routing control plane components on a label switching router (LSR).

Section 7 of ITU-T G.7712/Y.1703 [6] describes the transport DCN architecture and requirements as follows:

- 1) The MPLS-TP MCN MUST support the requirements for:
 - a) CCh access functions specified in Section 7.1.1;
 - b) MPLS-TP SCC data-link layer termination functions specified in Section 7.1.2.3;
 - c) MPLS-TP MCC data-link layer termination functions specified in Section 7.1.2.4;
 - d) Network layer PDU into CCh data-link frame encapsulation functions specified in Section 7.1.3;
 - e) Network layer PDU forwarding (Section 7.1.6), interworking (Section 7.1.7), and encapsulation (Section 7.1.8) functions, as well as tunneling (Section 7.1.9) and routing (Section 7.1.10) functions.

As a practical matter, MCN connections will typically have addresses. See the section on Identifiers in RFC 5921 [8] for further information.

In order to have the MCN operate properly, a number of management functions for the MCN are needed, including:

- o Retrieval of DCN network parameters to ensure compatible functioning, e.g., packet size, timeouts, quality of service, window size, etc.;
- o Establishment of message routing between DCN nodes;
- o Management of DCN network addresses;
- o Retrieval of operational status of the DCN at a given node;

- o Capability to enable/disable access by an NE to the DCN. Note that this is to allow the isolation of a malfunctioning NE to keep it from impacting the rest of the network.

5. Fault Management Requirements

The Fault Management functions within an MPLS-TP NE enable the supervision, detection, validation, isolation, correction, and reporting of abnormal operation of the MPLS-TP network and its environment.

5.1. Supervision Function

The supervision function analyzes the actual occurrence of a disturbance or fault for the purpose of providing an appropriate indication of performance and/or detected fault condition to maintenance personnel and operations systems.

- 1) The MPLS-TP NE MUST support supervision of the OAM mechanisms that are deployed for supporting the OAM requirements defined in RFC 5860 [3].
- 2) The MPLS-TP NE MUST support the following data-plane forwarding path supervision functions:
 - a) Supervision of loop-checking functions used to detect loops in the data-plane forwarding path (which result in non-delivery of traffic, wasting of forwarding resources, and unintended self-replication of traffic);
 - b) Supervision of failure detection;
- 3) The MPLS-TP NE MUST support the capability to configure data-plane forwarding path related supervision mechanisms to perform on-demand or proactively.
- 4) The MPLS-TP NE MUST support supervision for software processing -- e.g., processing faults, storage capacity, version mismatch, corrupted data, and out of memory problems, etc.
- 5) The MPLS-TP NE MUST support hardware-related supervision for interchangeable and non-interchangeable unit, cable, and power problems.
- 6) The MPLS-TP NE SHOULD support environment-related supervision for temperature, humidity, etc.

5.2. Validation Function

Validation is the process of integrating Fault Cause indications into Failures. A Fault Cause Indication (FCI) indicates a limited interruption of the required transport function. A Fault Cause is not reported to maintenance personnel because it might exist only for a very short period of time. Note that some of these events are summed up in the Performance Monitoring process (see Section 7), and when this sum exceeds a configured value, a threshold crossing alert (report) can be generated.

When the Fault Cause lasts long enough, an inability to perform the required transport function arises. This failure condition is subject to reporting to maintenance personnel and/or an OS because corrective action might be required. Conversely, when the Fault Cause ceases after a certain time, clearing of the Failure condition is also subject to reporting.

- 1) The MPLS-TP NE MUST perform persistency checks on fault causes before it declares a fault cause a failure.
- 2) The MPLS-TP NE SHOULD provide a configuration capability for control parameters associated with performing the persistency checks described above.
- 3) An MPLS-TP NE MAY provide configuration parameters to control reporting and clearing of failure conditions.
- 4) A data-plane forwarding path failure MUST be declared if the fault cause persists continuously for a configurable time (Time-D). The failure MUST be cleared if the fault cause is absent continuously for a configurable time (Time-C).

Note: As an example, the default time values might be as follows:

Time-D = 2.5 +/- 0.5 seconds

Time-C = 10 +/- 0.5 seconds

These time values are as defined in G.7710 [1].

- 5) MIBs - or other object management semantics specifications - defined to enable configuration of these timers SHOULD explicitly provide default values and MAY provide guidelines on ranges and value determination methods for scenarios where the default value chosen might be inadequate. In addition, such specifications SHOULD define the level of granularity at which tables of these values are to be defined.

- 6) Implementations MUST provide the ability to configure the preceding set of timers and SHOULD provide default values to enable rapid configuration. Suitable default values, timer ranges, and level of granularity are out of scope in this document and form part of the specification of fault management details. Timers SHOULD be configurable per NE for broad categories (for example, defects and/or fault causes), and MAY be configurable per-interface on an NE and/or per individual defect/fault cause.
- 7) The failure declaration and clearing MUST be time stamped. The time-stamp MUST indicate the time at which the fault cause is activated at the input of the fault cause persistency (i.e., defect-to-failure integration) function, and the time at which the fault cause is deactivated at the input of the fault cause persistency function.

5.3. Alarm Handling Function

5.3.1. Alarm Severity Assignment

Failures can be categorized to indicate the severity or urgency of the fault.

- 1) An MPLS-TP NE SHOULD support the ability to assign severity (e.g., Critical, Major, Minor, Warning) to alarm conditions via configuration.

See G.7710 [1], Section 7.2.2 for more detail on alarm severity assignment. For additional discussion of Alarm Severity management, see discussion of alarm severity in RFC 3877 [11].

5.3.2. Alarm Suppression

Alarms can be generated from many sources, including OAM, device status, etc.

- 1) An MPLS-TP NE MUST support suppression of alarms based on configuration.

5.3.3. Alarm Reporting

Alarm Reporting is concerned with the reporting of relevant events and conditions, which occur in the network (including the NE, incoming signal, and external environment).

Local reporting is concerned with automatic alarming by means of audible and visual indicators near the failed equipment.

- 1) An MPLS-TP NE MUST support local reporting of alarms.
- 2) The MPLS-TP NE MUST support reporting of alarms to an OS. These reports are either autonomous reports (notifications) or reports on request by maintenance personnel. The MPLS-TP NE SHOULD report local (environmental) alarms to a network management system.
- 3) An MPLS-TP NE supporting one or more other networking technologies (e.g., Ethernet, SDH/SONET, MPLS) over MPLS-TP MUST be capable of translating MPLS-TP defects into failure conditions that are meaningful to the client layer, as described in RFC 4377 [2], Section 4.7.

5.3.4. Alarm Reporting Control

Alarm Reporting Control (ARC) supports an automatic in-service provisioning capability. Alarm reporting can be turned off on a per-managed entity basis (e.g., LSP) to allow sufficient time for customer service testing and other maintenance activities in an "alarm free" state. Once a managed entity is ready, alarm reporting is automatically turned on.

- 1) An MPLS-TP NE SHOULD support the Alarm Reporting Control function for controlling the reporting of alarm conditions.

See G.7710 [1] (Section 7.1.3.2) and RFC 3878 [24] for more information about ARC.

6. Configuration Management Requirements

Configuration Management provides functions to identify, collect data from, provide data to, and control NEs. Specific configuration tasks requiring network management support include hardware and software configuration, configuration of NEs to support transport paths (including required working and protection paths), and configuration of required path integrity/connectivity and performance monitoring (i.e., OAM).

6.1. System Configuration

- 1) The MPLS-TP NE MUST support the configuration requirements specified in G.7710 [1], Section 8.1 for hardware.
- 2) The MPLS-TP NE MUST support the configuration requirements specified in G.7710 [1], Section 8.2 for software.

- 3) The MPLS-TP NE MUST support the configuration requirements specified in G.7710 [1], Section 8.13.2.1 for local real-time clock functions.
- 4) The MPLS-TP NE MUST support the configuration requirements specified in G.7710 [1], Section 8.13.2.2 for local real-time clock alignment with external time reference.
- 5) The MPLS-TP NE MUST support the configuration requirements specified in G.7710 [1], Section 8.13.2.3 for performance monitoring of the clock function.

6.2. Control Plane Configuration

- 1) If a control plane is supported in an implementation of MPLS-TP, the MPLS-TP NE MUST support the configuration of MPLS-TP control plane functions by the management plane. Further detailed requirements will be provided along with progress in defining the MPLS-TP control plane in appropriate specifications.

6.3. Path Configuration

- 1) In addition to the requirement to support static provisioning of transport paths (defined in RFC 5645 [7], Section 2.1 -- General Requirements, requirement 18), an MPLS-TP NE MUST support the configuration of required path performance characteristic thresholds (e.g., Loss Measurement <LM>, Delay Measurement <DM> thresholds) necessary to support performance monitoring of the MPLS-TP service(s).
- 2) In order to accomplish this, an MPLS-TP NE MUST support configuration of LSP information (such as an LSP identifier of some kind) and/or any other information needed to retrieve LSP status information, performance attributes, etc.
- 3) If a control plane is supported, and that control plane includes support for control-plane/management-plane hand-off for LSP setup/maintenance, the MPLS-TP NE MUST support management of the hand-off of Path control. For example, see RFCs 5943 [19] and 5852 [20].
- 4) Further detailed requirements SHALL be provided along with progress in defining the MPLS-TP control plane in appropriate specifications.

- 5) If MPLS-TP transport paths cannot be statically provisioned using MPLS LSP and pseudowire management tools (either already defined in standards or under development), further management specifications MUST be provided as needed.

6.4. Protection Configuration

- 1) The MPLS-TP NE MUST support configuration of required path protection information as follows:
 - o designate specifically identified LSPs as working or protecting LSPs;
 - o define associations of working and protecting paths;
 - o operate/release manual protection switching;
 - o operate/release force protection switching;
 - o operate/release protection lockout;
 - o set/retrieve Automatic Protection Switching (APS) parameters, including
 - o Wait to Restore time,
 - o Protection Switching threshold information.

6.5. OAM Configuration

- 1) The MPLS-TP NE MUST support configuration of the OAM entities and functions specified in RFC 5860 [3].
- 2) The MPLS-TP NE MUST support the capability to choose which OAM functions are enabled.
- 3) For enabled OAM functions, the MPLS-TP NE MUST support the ability to associate OAM functions with specific maintenance entities.
- 4) The MPLS-TP NE MUST support the capability to configure the OAM entities/functions as part of LSP setup and tear-down, including co-routed bidirectional point-to-point, associated bidirectional point-to-point, and uni-directional (both point-to-point and point-to-multipoint) connections.
- 5) The MPLS-TP NE MUST support the configuration of maintenance entity identifiers (e.g., MEP ID and MIP ID) for the purpose of LSP connectivity checking.

- 6) The MPLS-TP NE MUST support configuration of OAM parameters to meet their specific operational requirements, such as
 - a) one-time on-demand immediately or
 - b) one-time on-demand pre-scheduled or
 - c) on-demand periodically based on a specified schedule or
 - d) proactive on-going.
- 7) The MPLS-TP NE MUST support the enabling/disabling of the connectivity check processing. The connectivity check process of the MPLS-TP NE MUST support provisioning of the identifiers to be transmitted and the expected identifiers.

7. Performance Management Requirements

Performance Management provides functions for the purpose of maintenance, bring-into-service, quality of service, and statistics gathering.

This information could be used, for example, to compare behavior of the equipment, MPLS-TP NE, or network at different moments in time to evaluate changes in network performance.

ITU-T Recommendation G.7710 [1] provides transport performance monitoring requirements for packet-switched and circuit-switched transport networks with the objective of providing a coherent and consistent interpretation of the network behavior in a multi-technology environment. The performance management requirements specified in this document are driven by such an objective.

7.1. Path Characterization Performance Metrics

- 1) It MUST be possible to determine when an MPLS-TP-based transport service is available and when it is unavailable.

From a performance perspective, a service is unavailable if there is an indication that performance has degraded to the extent that a configurable performance threshold has been crossed and the degradation persists long enough (i.e., the indication persists for some amount of time, which is either configurable or well-known) to be certain it is not a measurement anomaly.

Methods, mechanisms, and algorithms for exactly how unavailability is to be determined -- based on collection of raw performance data -- are out of scope for this document.

- 2) The MPLS-TP NE MUST support collection and reporting of raw performance data that MAY be used in determining the unavailability of a transport service.
- 3) MPLS-TP MUST support the determination of the unavailability of the transport service. The result of this determination MUST be available via the MPLS-TP NE (at service termination points), and determination of unavailability MAY be supported by the MPLS-TP NE directly. To support this requirement, the MPLS-TP NE management information model MUST include objects corresponding to the availability-state of services.

Transport network unavailability is based on Severely Errored Seconds (SES) and Unavailable Seconds (UAS). The ITU-T is establishing definitions of unavailability that are generically applicable to packet transport technologies, including MPLS-TP, based on SES and UAS. Note that SES and UAS are already defined for Ethernet transport networks in ITU-T Recommendation Y.1563 [25].

- 4) The MPLS-TP NE MUST support collection of loss measurement (LM) statistics.
- 5) The MPLS-TP NE MUST support collection of delay measurement (DM) statistics.
- 6) The MPLS-TP NE MUST support reporting of performance degradation via fault management for corrective actions.

"Reporting" in this context could mean:

- o reporting to an autonomous protection component to trigger protection switching,
 - o reporting via a craft interface to allow replacement of a faulty component (or similar manual intervention),
 - o etc.
- 7) The MPLS-TP NE MUST support reporting of performance statistics on request from a management system.

7.2. Performance Measurement Instrumentation

7.2.1. Measurement Frequency

- 1) For performance measurement mechanisms that support both proactive and on-demand modes, the MPLS-TP NE MUST support the capability to be configured to operate on-demand or proactively.

7.2.2. Measurement Scope

On measurement of packet loss and loss ratio:

- 1) For bidirectional (both co-routed and associated) point-to-point (P2P) connections
 - a) on-demand measurement of single-ended packet loss and loss ratio measurement is REQUIRED;
 - b) proactive measurement of packet loss and loss ratio measurement for each direction is REQUIRED.
- 2) For unidirectional (P2P and point-to-multipoint (P2MP)) connection, proactive measurement of packet loss and loss ratio is REQUIRED.

On Delay measurement:

- 3) For a unidirectional (P2P and P2MP) connection, on-demand measurement of delay measurement is REQUIRED.
- 4) For a co-routed bidirectional (P2P) connection, on-demand measurement of one-way and two-way delay is REQUIRED.
- 5) For an associated bidirectional (P2P) connection, on-demand measurement of one-way delay is REQUIRED.

8. Security Management Requirements

- 1) The MPLS-TP NE MUST support secure management and control planes.

8.1. Management Communication Channel Security

- 1) Secure communication channels MUST be supported for all network traffic and protocols used to support management functions. This MUST include, at least, protocols used for configuration, monitoring, configuration backup, logging, time synchronization, authentication, and routing.
- 2) The MCC MUST support application protocols that provide confidentiality and data-integrity protection.
- 3) The MPLS-TP NE MUST support the following:
 - a) Use of open cryptographic algorithms (see RFC 3871 [4]).

- b) Authentication - allow management connectivity only from authenticated entities.
- c) Authorization - allow management activity originated by an authorized entity, using (for example) an Access Control List (ACL).
- d) Port Access Control - allow management activity received on an authorized (management) port.

8.2. Signaling Communication Channel Security

Security requirements for the SCC are driven by considerations similar to MCC requirements described in Section 8.1.

Security Requirements for the control plane are out of scope for this document and are expected to be defined in the appropriate control plane specifications.

- 1) Management of control plane security MUST be defined in the appropriate control plane specifications.

8.3. Distributed Denial of Service

A denial-of-service (DoS) attack is an attack that tries to prevent a target from performing an assigned task, or providing its intended service(s), through any means. A Distributed DoS (DDoS) can multiply attack severity (possibly by an arbitrary amount) by using multiple (potentially compromised) systems to act as topologically (and potentially geographically) distributed attack sources. It is possible to lessen the impact and potential for DoS and DDoS by using secure protocols, turning off unnecessary processes, logging and monitoring, and ingress filtering. RFC 4732 [26] provides background on DoS in the context of the Internet.

- 1) An MPLS-TP NE MUST support secure management protocols and SHOULD do so in a manner that reduces potential impact of a DoS attack.
- 2) An MPLS-TP NE SHOULD support additional mechanisms that mitigate a DoS (or DDoS) attack against the management component while allowing the NE to continue to meet its primary functions.

9. Security Considerations

Section 8 includes a set of security requirements that apply to MPLS-TP network management.

- 1) Solutions MUST provide mechanisms to prevent unauthorized and/or unauthenticated access to management capabilities and private information by network elements, systems, or users.

Performance of diagnostic functions and path characterization involves extracting a significant amount of information about network construction that the network operator might consider private.

10. Acknowledgments

The authors/editors gratefully acknowledge the thoughtful review, comments, and explanations provided by Adrian Farrel, Alexander Vainshtein, Andrea Maria Mazzini, Ben Niven-Jenkins, Bernd Zeuner, Dan Romascanu, Daniele Ceccarelli, Diego Caviglia, Dieter Beller, He Jia, Leo Xiao, Maarten Vissers, Neil Harrison, Rolf Winter, Yoav Cohen, and Yu Liang.

11. References

11.1. Normative References

- [1] ITU-T Recommendation G.7710/Y.1701, "Common equipment management function requirements", July, 2007.
- [2] 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.
- [3] 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.
- [4] Jones, G., Ed., "Operational Security Requirements for Large Internet Service Provider (ISP) IP Network Infrastructure", RFC 3871, September 2004.
- [5] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [6] ITU-T Recommendation G.7712/Y.1703, "Architecture and specification of data communication network", June 2008.

- [7] 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.
- [8] 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.
- [9] Mansfield, S. Ed., Gray, E., Ed., and K. Lam, Ed., "Network Management Framework for MPLS-based Transport Networks", RFC 5950, September 2010.

12.2. Informative References

- [10] Beller, D. and A. Farrel, "An In-Band Data Communication Network For the MPLS Transport Profile", RFC 5718, January 2010.
- [11] Chisholm, S. and D. Romascanu, "Alarm Management Information Base (MIB)", RFC 3877, September 2004.
- [12] ITU-T Recommendation M.20, "Maintenance philosophy for telecommunication networks", October 1992.
- [13] Telcordia, "Network Maintenance: Network Element and Transport Surveillance Messages" (GR-833-CORE), Issue 5, August 2004.
- [14] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, June 2009.
- [15] Harrington, D., "Guidelines for Considering Operations and Management of New Protocols and Protocol Extensions", RFC 5706, November 2009.
- [16] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", Work in Progress, July 2010.
- [17] Presuhn, R., Ed., "Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3416, December 2002.
- [18] OMG Document formal/04-03-12, "The Common Object Request Broker: Architecture and Specification", Revision 3.0.3. March 12, 2004.

- [19] Caviglia, D., Bramanti, D., Li, D., and D. McDysan, "Requirements for the Conversion between Permanent Connections and Switched Connections in a Generalized Multiprotocol Label Switching (GMPLS) Network", RFC 5493, April 2009.
- [20] Caviglia, D., Ceccarelli, D., Bramanti, D., Li, D., and S. Bardalai, "RSVP-TE Signaling Extension for LSP Handover from the Management Plane to the Control Plane in a GMPLS-Enabled Transport Network", RFC 5852, April 2010.
- [21] ITU-T Recommendation G.806, "Characteristics of transport equipment - Description methodology and generic functionality", January, 2009.
- [22] ITU-T Recommendation Y.1731, "OAM functions and mechanisms for Ethernet based networks", February, 2008.
- [23] ITU-T Recommendation G.8601, "Architecture of service management in multi bearer, multi carrier environment", June 2006.
- [24] Lam, H., Huynh, A., and D. Perkins, "Alarm Reporting Control Management Information Base (MIB)", RFC 3878, September 2004.
- [25] ITU-T Recommendation Y.1563, "Ethernet frame transfer and availability performance", January 2009.
- [26] Handley, M., Ed., Rescorla, E., Ed., and IAB, "Internet Denial-of-Service Considerations", RFC 4732, December 2006.

Appendix A. Communication Channel (CCh) Examples

A CCh can be realized in a number of ways.

1. The CCh can be provided by a link in a physically distinct network, that is, a link that is not part of the transport network that is being managed. For example, the nodes in the transport network can be interconnected in two distinct physical networks: the transport network and the DCN.

This is a "physically distinct out-of-band CCh".

2. The CCh can be provided by a link in the transport network that is terminated at the ends of the DCC and that is capable of encapsulating and terminating packets of the management protocols. For example, in MPLS-TP, a single-hop LSP might be established between two adjacent nodes, and that LSP might be capable of carrying IP traffic. Management traffic can then be inserted into the link in an LSP parallel to the LSPs that carry user traffic.

This is a "physically shared out-of-band CCh."

3. The CCh can be supported as its native protocol on the interface alongside the transported traffic. For example, if an interface is capable of sending and receiving both MPLS-TP and IP, the IP-based management traffic can be sent as native IP packets on the interface.

This is a "shared interface out-of-band CCh".

4. The CCh can use overhead bytes available on a transport connection. For example, in TDM networks there are overhead bytes associated with a data channel, and these can be used to provide a CCh. It is important to note that the use of overhead bytes does not reduce the capacity of the associated data channel.

This is an "overhead-based CCh".

This alternative is not available in MPLS-TP because there is no overhead available.

5. The CCh can be provided by a dedicated channel associated with the data link. For example, the generic associated label (GAL) [14] can be used to label DCC traffic being exchanged on a data link between adjacent transport nodes, potentially in the absence of any data LSP between those nodes.

This is a "data link associated CCh".

It is very similar to case 2, and by its nature can only span a single hop in the transport network.

6. The CCh can be provided by a dedicated channel associated with a data channel. For example, in MPLS-TP, the GAL [14] can be imposed under the top label in the label stack for an MPLS-TP LSP to create a channel associated with the LSP that can carry management traffic. This CCh requires the receiver to be capable of demultiplexing management traffic from user traffic carried on the same LSP by use of the GAL.

This is a "data channel associated CCh".

7. The CCh can be provided by mixing the management traffic with the user traffic such that is indistinguishable on the link without deep-packet inspection. In MPLS-TP, this could arise if there is a data-carrying LSP between two nodes, and management traffic is inserted into that LSP. This approach requires that the termination point of the LSP be able to demultiplex the management and user traffic. This might be possible in MPLS-TP if the MPLS-TP LSP is carrying IP user traffic.

This is an "in-band CCh".

These realizations can be categorized as:

- A. Out-of-fiber, out-of-band (types 1 and 2)
- B. In-fiber, out-of-band (types 2, 3, 4, and 5)
- C. In-band (types 6 and 7)

The MCN and SCN are logically separate networks and can be realized by the same DCN or as separate networks. In practice, that means that, between any pair of nodes, the MCC and SCC can be the same link or separate links.

It is also important to note that the MCN and SCN do not need to be categorised as in-band, out-of-band, etc. This definition only applies to the individual links, and it is possible for some nodes to be connected in the MCN or SCN by one type of link, and other nodes by other types of link. Furthermore, a pair of adjacent nodes can be connected by multiple links of different types.

Lastly, note that the division of DCN traffic between links between a pair of adjacent nodes is purely an implementation choice. Parallel links can be deployed for DCN resilience or load sharing. Links can be designated for specific use. For example, so that some links

carry management traffic and some carry control plane traffic, or so that some links carry signaling protocol traffic while others carry routing protocol traffic.

It is important to note that the DCN can be a routed network with forwarding capabilities, but that this is not a requirement. The ability to support forwarding of management or control traffic within the DCN can substantially simplify the topology of the DCN and improve its resilience, but does increase the complexity of operating the DCN.

See also RFC 3877 [11], ITU-T M.20 [12], and Telcordia document GR-833-CORE [13] for further information.

Contributor's Address

Adrian Farrel
Old Dog Consulting
EMail: adrian@olddog.co.uk

Authors' Addresses

Eric Gray
Ericsson
900 Chelmsford Street
Lowell, MA, 01851
Phone: +1 978 275 7470
EMail: Eric.Gray@Ericsson.com

Scott Mansfield
Ericsson
250 Holger Way
San Jose CA, 95134
+1 724 931 9316
EMail: Scott.Mansfield@Ericsson.com

Hing-Kam (Kam) Lam
Alcatel-Lucent
600-700 Mountain Ave
Murray Hill, NJ, 07974
Phone: +1 908 582 0672
EMail: Kam.Lam@Alcatel-Lucent.com