Internet Engineering Task Force (IETF)
Request for Comments: 5950
Category: Informational
ISSN: 2070-1721

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

Network Management Framework for MPLS-based Transport Networks

Abstract

This document provides the network management framework for the Transport Profile for Multi-Protocol Label Switching (MPLS-TP).

This framework relies on the management terminology from the ITU-T to describe the management architecture that could be used for an MPLS-TP management network.

The management of the MPLS-TP network could be based on multi-tiered distributed management systems. This document provides a description of the network and element management architectures that could be applied and also describes heuristics associated with fault, configuration, and performance aspects of the management system.

This document is a product of a joint Internet Engineering Task Force (IETF) / International Telecommunication Union Telecommunication Standardization Sector (ITU-T) effort to include an MPLS Transport Profile within the IETF MPLS and PWE3 architectures to support the capabilities and functionalities of a packet transport network.

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/rfc5950>.

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	3
1.1. Terminology	3
2. Management Architecture	5
2.1. Network Management Architecture	5
2.2. Element Management Architecture	6
2.3. Standard Management Interfaces	10
2.4. Management- and Control-Specific Terminology	11
2.5. Management Channel	11
3. Fault Management	13
3.1. Supervision	13
3.2. Validation	13
3.3. Alarm Handling	13
4. Configuration Management	13
4.1. LSP Ownership Handover	14
5. Performance Management	15
6. Acknowledgements	15
7. Security Considerations	16
8. References	16
8.1. Normative References	16
8.2. Informative References	17

1. Introduction

This document provides the network management framework for the Transport Profile for Multi-Protocol Label Switching (MPLS-TP). Requirements for network management in an MPLS-TP network are documented in "Network Management Requirements for MPLS-based Transport Networks" [3], and this document explains how network elements and networks that support MPLS-TP can be managed using solutions that satisfy those requirements. The relationship between Operations, Administration, and Maintenance (OAM), management, and other framework documents is described in the MPLS-TP framework [4] document.

This document is a product of a joint Internet Engineering Task Force (IETF) / International Telecommunication Union Telecommunication Standardization Sector (ITU-T) effort to include an MPLS Transport Profile within the IETF MPLS and PWE3 architectures to support the capabilities and functionalities of a packet transport network.

1.1. Terminology

This framework relies on the management terminology from the ITU-T to describe the management architecture that could be used for an MPLS-TP management network. The terminology listed below are taken from/based on the definitions found in ITU-T G.7710 [6], ITU-T G.7712 [7], and ITU-T M.3013 [13].

- o Communication Channel (CCh): A logical channel between network elements (NEs) that can be used in (for example) management plane applications or control plane applications. For MPLS-TP, the physical channel supporting the CCh is the MPLS-TP Management Communication Channel (MCC).
- o Data Communication Network (DCN): A network that supports Layer 1 (physical), Layer 2 (data-link), and Layer 3 (network) 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.). See ITU-T G.7712 [7].
- o Equipment Management Function (EMF): The management functions within an NE. See ITU-T G.7710 [6].
- o Local Craft Terminal (LCT): An out-of-band device that connects to an NE for management purposes. See ITU-T G.7710 [6].

- o Label Switched Path (LSP): An MPLS-TP LSP is an LSP that uses a subset of the capabilities of an MPLS LSP in order to meet the requirements of an MPLS transport network as described in the MPLS-TP framework [4].
- o Management Application Function (MAF): An application process that participates in system management. See ITU-T G.7710 [6].
- o Management Communication Channel (MCC): A CCh dedicated for management plane communications. See ITU-T G.7712 [7].
- o Message Communication Function (MCF): The communications process that performs functions such as information interchange and relay. See ITU-T M.3013 [13].
- o Management Communication Network (MCN): A DCN supporting management plane communication is referred to as a Management Communication Network (MCN). See ITU-T G.7712 [7].
- o MPLS-TP NE: A network element (NE) that supports MPLS-TP functions. Another term that is used for a network element is node. In terms of this document, the term node is equivalent to NE.
- o MPLS-TP network: A network in which MPLS-TP NEs are deployed.
- o Network Element Function (NEF): The set of functions necessary to manage a network element. See ITU-T M.3010 [11].
- o Operations, Administration, and Maintenance (OAM): For the MPLS-TP effort the term OAM means the set of tools that consist of "operation" activities that are undertaken to keep the network up and running, "administration" activities that keep track of resources in the network and how they are used, and "maintenance" activities that facilitate repairs and upgrades. For a complete expansion of the acronym, see "The OAM Acronym Soup" [15].
- o Operations System (OS): A system that performs the functions that support processing of information related to operations, administration, maintenance, and provisioning (OAM&P) (see "The OAM Acronym Soup" [15]) for the networks, including surveillance and testing functions to support customer access maintenance. See ITU-T M.3010 [11].
- o Signaling Communication Network (SCN): A DCN supporting control plane communication is referred to as a Signaling Communication Network (SCN). See ITU-T G.7712 [7].

- o Signaling Communication Channel (SCC): A CCh dedicated for control plane communications. The SCC may be used for GMPLS/ASON signaling and/or other control plane messages (e.g., routing messages). See ITU-T G.7712 [7].

2. Management Architecture

The management of the MPLS-TP network could be based on a multi-tiered distributed management systems, for example as described in ITU-T M.3010 [11] and ITU-T M.3060/Y.2401 [12]. Each tier provides a predefined level of network management capabilities. The lowest tier of this organization model includes the MPLS-TP network element that provides the transport service and the Operations System (OS) at the Element Management Level. The Management Application Function (MAF) within the NEs and OSs provides the management support. The MAF at each entity can include agents only, managers only, or both agents and managers. The MAF that includes managers is capable of managing an agent included in other MAF.

The management communication to peer NEs and/or OSs is provided via the Message Communication Function (MCF) within each entity (e.g., NE and OS). The user can access the management of the MPLS-TP transport network via a Local Craft Terminal (LCT) attached to the NE or via a Work Station (WS) attached to the OS.

2.1. Network Management Architecture

A transport Management Network (MN) may consist of several transport-technology-specific Management Networks. Management network partitioning (Figure 1) below (based on ITU-T G.7710 [6]) shows the management network partitioning. Notation used in G.7710 for a transport-technology-specific MN is x.MN, where x is the transport-specific technology. An MPLS-TP-specific MN is abbreviated as MT.MN. Where there is no ambiguity, we will use "MN" for an MPLS-TP-specific MN. In the figure below, O.MSN is equivalent to an OTN management Subnetwork.

Management Network Partitioning

Figure 1

The management of the MPLS-TP network is separable from the management of the other technology-specific networks, and it operates independently of any particular client- or server-layer management plane.

An MPLS-TP Management Network (MT.MN) could be partitioned into MPLS-TP Management SubNetworks ("MT.MSN" or "MPLS-TP MSN", or just "MSN" where usage is unambiguous) for consideration of scalability (e.g., geographic or load balancing) or administration (e.g., operation or ownership).

The MPLS-TP MSN could be connected to other parts of the MN through one or more LCTs and/or OSSs. The Message Communication Function (MCF) of an MPLS-TP NE initiates/terminates, routes, or otherwise processes management messages over CChs or via an external interface.

Multiple addressable MPLS-TP NEs could be present at a single physical location (i.e., site or office). The inter-site communications link between the MPLS-TP NEs will normally be provided by the CChs. Within a particular site, the NEs could communicate via an intra-site CCh or via a LAN.

2.2. Element Management Architecture

The Equipment Management Function (EMF) of an MPLS-TP NE provides the means through which a management system manages the NE.

The EMF interacts with the NE's transport functions by exchanging Management Information (MI) across the Management Point (MP) Reference Points. The EMF may contain a number of functions that

provide a data reduction mechanism on the information received across the MP Reference Points.

The EMF includes functions such as Date and Time, FCAPS (Fault, Configuration, Accounting, Performance, and Security) management, and Control Plane functions. The EMF provides event message processing, data storage, and logging. The management Agent, a component of the EMF, converts internal management information (MI signals) into Management Application messages and vice versa. The Agent responds to Management Application messages from the Message Communication Function (MCF) by performing the appropriate operations on (for example) the Managed Objects in a Management Information Base (MIB), as necessary. The MCF contains communications functions related to the world outside of the NE (i.e., Date and Time source, Management Plane, Control Plane, Local Craft Terminal, and Local Alarms).

The Date and Time functions keep track of the NE's date/time, which is used by the FCAPS management functions to e.g., time stamp event reports.

Below are diagrams that illustrate the components of the Equipment Management Function (EMF) of a Network Element (NE). The high-level decomposition of the Network Element Function (NEF) picture (Figure 2) provides the breakdown of the NEF, then the EMF picture (Figure 3) provides the details of Equipment Management Function, and finally the Message Communication Function (MCF) picture (Figure 4) details the MCF.

High-Level Decomposition of NEF

Figure 2

Equipment Management Function

Figure 3

Message Communication Function

Figure 4

2.3. Standard Management Interfaces

The "Network Management Requirements for MPLS-based Transport Networks" document [3] places no restriction on which management interface is to be used for managing an MPLS-TP network. It is possible to provision and manage an end-to-end connection across a network where some segments are created/managed/deleted, for example by NETCONF or SNMP and other segments by CORBA interfaces. Use of any network management interface for one management-related purpose does not preclude use of another network management interface for

other management-related purposes, or the same purpose at another time. The protocol(s) to be supported are at the discretion of the operator.

2.4. Management- and Control-Specific Terminology

Data Communication Network (DCN) is the common term for the network used to transport Management and Signaling information between: management systems and network elements, management systems to other management systems, and network elements to other network elements. The Management Communications Network (MCN) is the part of the DCN that supports the transport of Management information for the Management Plane. The Signaling Communications Network (SCN) is the part of the DCN that supports transport of signaling information for the Control Plane. As shown in , the communication channel terminology picture (Figure 5) each technology has its own terminology that is used for the channels that support the transfer of management and control plane information. For MPLS-TP, the management plane uses the Management Communication Channel (MCC), and the control plane uses the Signaling Communication Channel (SCC).

2.5. Management Channel

The Communication Channel (CCh) provides a logical channel between NEs for transferring Management and/or Signaling information. Note that some technologies provide separate communication channels for Management (MCC) and Signaling (SCC).

MPLS-TP NEs communicate via the DCN. The DCN connects NEs with management systems, NEs with NEs, and management systems with management systems.

Common Terminology

Management Terminology

Control Terminology

Communication Channel Terminology

Figure 5

3. Fault Management

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. Fault management provides the mechanisms to detect, verify, isolate, notify, and recover from the fault.

3.1. Supervision

ITU-T G.7710 [6] lists five basic categories of supervision that provide the functionality necessary to detect, verify, and notify a fault. The categories are: Transmission Supervision, Quality of Service Supervision, Processing Supervision, Hardware Supervision, and Environment Supervision. Each of the categories provides a set of recommendations to ensure that the fault management process is fulfilled.

3.2. Validation

ITU-T G.7710 [6] describes a fault cause as a limited interruption of the required function. It is not reasonable for every fault cause to be reported to maintenance personnel. The validation process is used to turn fault causes (events) into failures (alarms).

3.3. Alarm Handling

Within an element management system, it is important to consider mechanisms to support severity assignment, alarm reporting control, and logging.

4. Configuration Management

Configuration management provides the mechanisms to:

- o provision the MPLS-TP services
- o set up security for the MPLS-TP services and MPLS-TP network elements
- o provide the destination for fault notifications and performance parameters
- o configure and control OAM

Also associated with configuration management are hardware and software provisioning and inventory reporting.

4.1. LSP Ownership Handover

MPLS-TP networks can be managed not only by Network Management Systems (i.e., Management Plane (MP)), but also by Control Plane (CP) protocols. The utilization of the control plane is not a mandatory requirement (see MPLS-TP Requirements [2]), but it is often used by network operators in order to make network configuration and Label Switched Path (LSP) recovery both faster and simpler.

In networks where both CP and MP are provided, an LSP could be created by either (CP or MP). The entity creating an LSP owns the data plane resources comprising that LSP. Only the owner of an LSP is typically able to modify/delete it. This results in a need for interaction between the MP and CP to allow either to manage all the resources of a network.

Network operators might prefer to have full control of the network resources during the set-up phase and then allow the network to be automatically maintained by the Control Plane. This can be achieved by creating LSPs via the Management Plane and subsequently transferring LSP ownership to the Control Plane. This is referred to as "ownership handover" RFC 5493 [10]. MP to CP ownership handover is then considered a requirement where a Control Plane is in use that supports it. The converse (CP to MP ownership handover) is a feature that is recommended -- but not required -- for (G)MPLS networks because it has only minor applications (for example, moving LSPs from one path to another as a maintenance operation).

The LSP handover procedure has already been standardized for GMPLS networks, where the signaling protocol used is RSVP-TE (RFC 3209 [1]). The utilization of RSVP-TE enhancements are defined in [5].

MP and CP interworking also includes the exchange of information that is either requested by the MP, or a notification by the CP as a consequence of a request from the MP or an automatic action (for example, a failure occurs or an operation is performed). The CP is asked to notify the MP in a reliable manner about the status of the operations it performs and to provide a mechanism to monitor the status of Control Plane objects (e.g., TE Link status, available resources), and to log operations related to Control Plane LSP. Logging is one of the most critical aspects because the MP always needs to have an accurate history and status of each LSP and all Data Plane resources involved in it.

5. Performance Management

Performance statistics could overwhelm a Management Network, so it is important to provide flexible instrumentation that enables control over the amount of performance data to be collected. Mechanisms for limiting the quantity of information collected are well known and deployed in IETF standards (see RFC 2819 (RMON) [8] and RFC 4502 (RMON2) [9]). The details of the performance data collected (including loss and delay measurement data) are found in the "Network Management Requirements for MPLS-based Transport Networks" document [3].

A distinction is made between performance data that is collected on-demand and data that is collected proactively. The definitions of on-demand and proactive measurement are provided for OAM in the "Network Management Requirements for MPLS-based Transport Networks" document [3].

On-demand measurement provides the operator with the ability to do performance measurement for maintenance purpose, such as diagnosis or to provide detailed verification of proactive measurement. It is used typically on specific LSP service instances for a limited time, thus limiting its impact on network performance under normal operations. Therefore, on-demand measurement does not result in scaling issues.

Proactive measurement is used continuously over time after being configured with periodicity and storage information. Data collected from proactive measurement are usually used for verifying the performance of the service. Proactive performance monitoring has the potential to overwhelm both the process of collecting performance data at a network element (for some arbitrary number of service instances traversing the NE), and the process of reporting this information to the OS. As a consequence of these considerations, operators would typically limit the services to which proactive performance measurement would be applied to a very selective subset of the services being provided and would limit the reporting of this information to statistical summaries (as opposed to raw or detailed performance statistics).

6. Acknowledgements

The authors/editors gratefully acknowledge the thoughtful review, comments and explanations provided by Diego Caviglia, Bernd Zeuner and Dan Romascanu.

7. Security Considerations

The ability for the authorized network operator to access EMF interfaces (Section 2.3) when needed is critical to proper operation. Therefore, the EMF interfaces need to be protected from denial-of-service conditions or attack. The EMF interfaces that use or access private information should be protected from eavesdropping, mis-configuration, and/or mal-configuration by unauthorized 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 considers private.

Section 4.3 of the "Security Framework for MPLS and GMPLS Networks" document [14] provides a description of the attacks on the Operation and Management Plane and also discusses the background necessary to understand security practices in Internet Service Provider environments. The security practices described are applicable to MPLS-TP environments.

8. References

8.1. Normative References

- [1] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [2] Niven-Jenkins, B., Brungard, D., Betts, M., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, September 2009.
- [3] Lam, K., Mansfield, S., and E. Gray, "Network Management Requirements for MPLS-based Transport Networks", RFC 5951, September 2010.
- [4] Bocci, M., Bryant, S., Frost, D., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, July 2010.
- [5] 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.
- [6] International Telecommunication Union, "Common equipment management function requirements", ITU-T Recommendation G.7710/Y.1701, July 2007.

- [7] International Telecommunication Union, "Architecture and specification of data communication network", ITU-T Recommendation G.7712/Y.1703, June 2008.

8.2. Informative References

- [8] Waldbusser, S., "Remote Network Monitoring Management Information Base", STD 59, RFC 2819, May 2000.
- [9] Waldbusser, S., "Remote Network Monitoring Management Information Base Version 2", RFC 4502, May 2006.
- [10] 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.
- [11] International Telecommunication Union, "Principles for a telecommunication management network", ITU-T Recommendation M.3010, April 2005.
- [12] International Telecommunication Union, "Principles for the Management of Next Generation Networks", ITU-T Recommendation M.3060/Y.2401, March 2006.
- [13] International Telecommunication Union, "Considerations for a telecommunication management network", ITU-T Recommendation M.3013, February 2000.
- [14] Fang, L., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
- [15] Andersson, L., Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "The OAM Acronym Soup", Work in progress, June 2010.

Authors' Addresses

Scott Mansfield (editor)
Ericsson
300 Holger Way
San Jose, CA 95134
US

Phone: +1 724 931 9316
Email: scott.mansfield@ericsson.com

Eric Gray (editor)
Ericsson
900 Chelmsford Street
Lowell, MA 01851
US

Phone: +1 978 275 7470
Email: eric.gray@ericsson.com

Hing-Kam Lam (editor)
Alcatel-Lucent
600-700 Mountain Ave
Murray Hill, NJ 07974
US

Phone: +1 908 582 0672
Email: Kam.Lam@alcatel-lucent.com