MPLS TRANSPORT PROFILE (MPLS-TP)

A Set of Enhancements to the Rich MPLS Toolkit

Table of Contents

Table of Figures

Executive Summary

The rise of bandwidth-hungry applications such as IPTV and mobile video, coupled with the pressure to minimize the cost per bit and maximize the value per bit, is forcing carriers to transition their transport networks from circuit-based technologies to packet-based technologies. MPLS has been a widely successful connection-oriented packet transport technology for more than a decade. However, it requires a few enhancements to provide functionality and manageability that is equivalent to the current circuit-switched transport networks. The set of these enhancements is called MPLS Transport Profile (MPLS-TP) by the standards bodies that are helping to develop it. MPLS-TP reuses most of the existing protocols from the rich MPLS/GMPLS (generalized MPLS) suite, and then adds a few enhancements, most notably in the area of Operation, Administration, and Management (OAM). The MPLS-TP enhancements will increase the applicability of MPLS overall, allowing it to serve both the transport (access and core) and the services networks.

Juniper Networks is committed to and actively driving the advancement of MPLS and MPLS-TP in the Internet Engineering Task Force (IETF) and the International Telecommunication Union (ITU). Juniper already implements most of the existing MPLS/GMPLS technology that MPLS-TP includes, and will continue to add IETF standardized MPLS-TP enhancements to its rich MPLS technology portfolio, with an emphasis on extending the functionality in areas of OAM for measuring delay and loss.

Introduction

The world is moving towards packet-based transport networks, primarily because all of the applications and services that use these networks are packet-based, and a packet-based transport network is best suited for carrying packets. However, a large number of current transport networks have been built using circuit-switched time-division multiplexing (TDM) technologies such as T1/E1 and SONET/SDH. Before these networks can be migrated to packet-based technologies such as IP/MPLS, the packet-based technologies need to be enhanced to provide equivalency with legacy networks, both in terms of functionality and manageability.

MPLS, which has been a widely successful, connection oriented, packet transport technology, has been production-hardened in thousands of networks worldwide and is ideally suited for packet-based transport networks.. A significant number of service providers have already moved their core networks to MPLS, and many would like to converge their next-generation access, aggregation, and core networks to MPLS as well. However, for this migration to take place, a few enhancements to the MPLS protocol suite are needed to provide functionality and manageability that matches the current circuit-switched transport networks.

GMPLS was an evolutionary technology which was developed after the success of MPLS in packet-switched networks. Traditional MPLS was designed to carry Layer 3 IP traffic by establishing IP-based paths and associating these paths with arbitrarily assigned labels. GMPLS extends the functionality of MPLS to encompass label switching for TDM (e.g., SONET/SDH, PDH), wavelength (lambdas), and spatial switching (e.g., incoming port or fiber to outgoing port or fiber) devices.

What is MPLS-TP

MPLS-TP is a profile of MPLS for transport networks. It takes a subset of MPLS/GMPLS protocol suite and adds a few extensions to address transport network requirements. These enhancements extend the already rich MPLS/GMPLS protocol suite such that it will be able to serve both transport and services networks. The standardization effort for MPLS-TP is being pursued jointly by IETF (working groups MPLS, CCAMP, and PWE) and ITU-T (group SG15). Based on their agreement, IETF will define the necessary extensions to the protocols, and ITU-T will define the requirements and will work with IETF on the enhancements.

Note that MPLS-TP refers to the entire set of enhancements, just like MPLS refers to a suite of protocols. Therefore, when an implementation needs to be checked for compliancy, it should be tested against individual RFCs or drafts that are being developed by the MPLS-TP effort.

MPLS and MPLS-TP Components

As mentioned previously, MPLS refers to a suite of protocols, and MPLS-TP refers to a set of compatible enhancements to the MPLS protocol suite. These protocols and new enhancements can be separated into the following categories:

- Network Architecture—Covers the definition of various functions and the interactions among them.
- Data Plane—Covers the protocols and mechanisms that are used to forward the data packets. This can further be divided into the following subcategories:
  - Framing, forwarding, encapsulation
  - OAM
  - Resiliency (protection and restoration)
- Control Plane—Covers the protocols and mechanisms used to set up the label-switched paths (LSPs) that are used to forward the data packets.
- Management Plane—Covers the protocols and mechanisms that are used to manage the network.

A list of protocols and mechanisms in each of these categories is provided in Figure 1. The figure also highlights the set of enhancements that are being pursued by MPLS-TP. The protocol and mechanisms highlighted in blue are being added to the MPLS/GMPLS protocol suite as part of the MPLS-TP effort. In Figure 1, the protocols and mechanisms highlighted in red might not be needed for the transport networks and are, therefore, being made optional. Note that these mechanisms will remain as part of the MPLS/GMPLS protocol suite. It is IETF's guidance to vendors that these mechanisms do not need to be supported on the platforms that are being targeted towards transport networks.

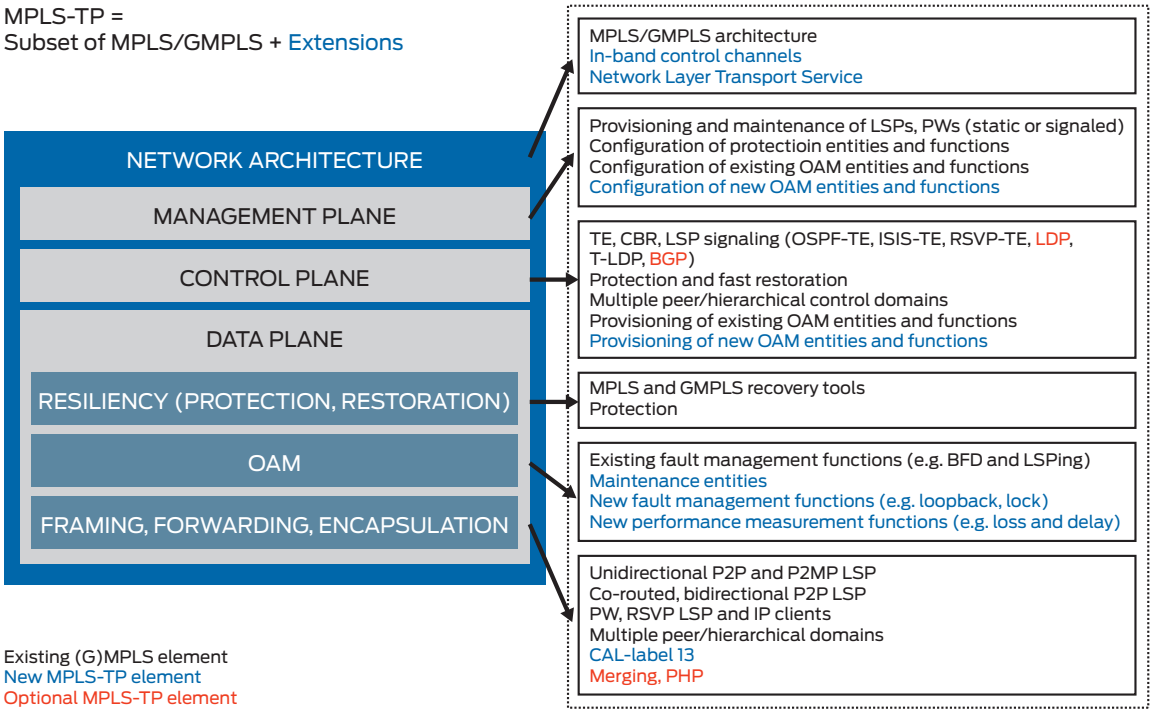

Figure 1: Components of MPLS and MPLS-TP

The following sections discuss the protocols and mechanisms that are relevant to MPLS-TP.

Operation, Administration, and Management (OAM)

This is the key focus area of the MPLS-TP and by far the most needed one by MPLS in general. Legacy transport networks use extensive and well established tools to monitor and manage transport networks, as providing and enforcing service-level agreements (SLAs) is a critical requirement for these networks. Note that the OAM enhancements that are being added to the MPLS protocol suite are required for transport networks but will prove to be extremely valuable for other types of MPLS networks as well.

The OAM functions being added as part of MPLS-TP (see Table 1 below) are fault detection (e.g., connectivity check, connectivity/path verification), fault localization (e.g., loopback, lock), and performance monitoring (e.g., delay and loss measurement). Note that the existing MPLS tools such as Bidirectional Forwarding Detection (BFD), LSP ping, and LSP trace are being extended to support these new OAM functions. The following table describes the role of these new OAM functions and the tools that are being used to enable them.

Table 1: OAM Enhancements in MPLS-TPs

OAM FUNCTION	OAM SUB-FUNCTION	GOAL/PURPOSE	TOOLS USED
Fault management (detection and localization)	Continuity Check (CC)	Provides rapid, proactive identification of faults.	Extended BFD (proactive) Extended LSP ping (reactive)
	Connectivity Verification (CV)	Allows on-demand localization of the fault after detection by CC.	Extended BFD (proactive) Extended LSP ping (reactive)
	Loopback	Allows an operator to put an LSP in loopback mode (i.e., the packets on that path are forwarded back to the originator). This is useful for testing or measurements.	In-band message in G-Ach (or) Extended LSP ping
	Lock	Allows an operator to put a path out of service. On a locked LSP, only test or OAM traffic can be sent. An LSP needs to be locked before putting it in the loopback mode for testing or measurements.	In-band message in G-ACh (or) Extended LSP ping
	Remote Defect Indication (RDI)	Used by the endpoints to communicate defect notifications.	Extended BFD
Performance monitoring	Delay measurement	Allows measurement of delay in forwarding the packets on an established path between two endpoints.	New DM tool
	Loss measurement	Allows measurement of loss on an established path between two endpoints.	New LM tool
	Throughput measurement	Allows measurement of throughput on an established path between two endpoints.	New LM tool
	Delay variation measurement	Allows measurement of variation in delay on an established path between two endpoints.	New DM tool

Since MPLS-TP is designed to work in devices where IP routing is not supported, these OAM functions need to operate without any IP layer functionalities. In order to make that possible, the framing, forwarding, and encapsulation component of the MPLS protocol suite is being enhanced with Generic Associated Channel (G-ACh) and G-ACh Label (GAL) to carry the OAM packets without any reliance on IP. Also, the OAM packets need to traverse the same path as the data packets. To support this requirement, the network architecture component of the MPLS protocol suite is being enhanced to support the in-band control channels.

G-ACh and GAL

In order to ensure congruency between the OAM packets and the data path, the OAM packets use in-band control channels. The idea of tagging the packets with an additional header was first introduced in the context of MPLS pseudowires, via the ACH [RFC 4485]. The ACH indicates that the tagged packet must be processed by an appropriate OAM function. This idea was generalized to a generic ACh (G-ACh) as part of the MPLS-TP effort and now applies to LSPs and segments as well. So, G-ACh is simply a header in the packet that provides the demultiplexor function for OAM packets for appropriate handling.

Note that the existence of ACh was negotiated when the pseudowire was set up, which is not feasible if static provisioning is used. This problem has been solved by using one of the reserved labels for this purpose. RFC 5586 identifies the reserved value 13 as a G-ACh label (GAL), thus providing the necessary tagging. Use of GAL for tagging OAM packets also enables easy extraction of the OAM packets at either a midpoint or an endpoint of an LSP or a pseudowire.

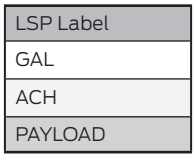

Figure 2: Packet headers for OAM packets in MPLS-TP

Control Plane (Static or Dynamic)

The control plane mechanism is responsible for the setup of LSPs (dynamically or statically) across a MPLS network. The MPLS protocol suite supports a robust and mature dynamic control plane with protocols such as OSPF-TE, IS-IS-TE, RSVP-TE, LDP, and BGP. The current transport networks, however, have been using a static control plane, i.e., the circuits are statically provisioned by an intelligent network management system (NMS). Dynamic control plane is optional with MPLS-TP.

Static provisioning in MPLS has been supported by many vendors, including Juniper, for quite some time. The static control plane may have applicability in scenarios where some equipment, especially equipment used at the edges of the network, does not support a dynamic control plane, or in which static configuration is preferred for security reasons. The NMS-driven control plane also allows operators to manage the packet-based network in the same way that they have historically been managing the circuit-switched network.

Even though the use of dynamic control plane is optional in MPLS-TP, a dynamic control plane has its own advantages, in particular with regards to scaling. It also provides advanced protection functions (for example, schemes such as LSP tail-end protection). Therefore, operators that are comfortable with the dynamic control plane can and are encouraged to use GMPLS and T-LDP to set up the LSPs and the pseudowires respectively in the context of MPLS-TP. The dynamic control plane, coupled with Juniper Networks® Junos® SDK, can deliver maximum flexibility and embedded logic for path computation, routing, and application or location-aware routing decisions.

Resiliency (Protection and Restoration)

MPLS has a rich set of protection and restoration mechanisms such as LSP fast reroute, pseudowire redundancy, and path protection. MPLS-TP work enhances the resiliency mechanism of MPLS by adding support for OAM-triggered protection (i.e., allowing an operator to trigger the LSP to a secondary path) and optimizing protection in ring topologies. Ring topologies are important, as circuit networks are typically built as interconnected rings, and it is expected that many initial deployments of MPLS-TP will consist of replacing the circuit-switched nodes with MPLS-TP packet switching nodes. Though MPLS fast reroute works in ring topologies, it does so in an inefficient way. Various optimizations and schemes (such as wrapping and steering) have been developed as part of MPLS-TP work to provide efficient protection in ring topologies.

Applicability and Deployment Options

MPLS-TP enhancements are primarily applicable to the access and aggregation networks, where the majority of the migration from circuit-switched networks to packet-based networks is currently occurring, and where higher scale and lower cost is required. Juniper believes that the OAM enhancements to the MPLS protocol suite, however, will be extremely valuable to all MPLS networks, especially in the MPLS-based core networks. These OAM enhancements will allow service providers to have better visibility into their existing MPLS-based core networks, which will allow further optimization. The new OAM capabilities will also help the wholesale business by improving the tools required to measure and enforce strict SLAs. Juniper, therefore, is prioritizing the implementation of these OAM enhancements, such as the enhancements to BFD and LSP ping.

Figure 3 illustrates how IP/MPLS and MPLS-TP can be deployed together and are very complementary in nature.

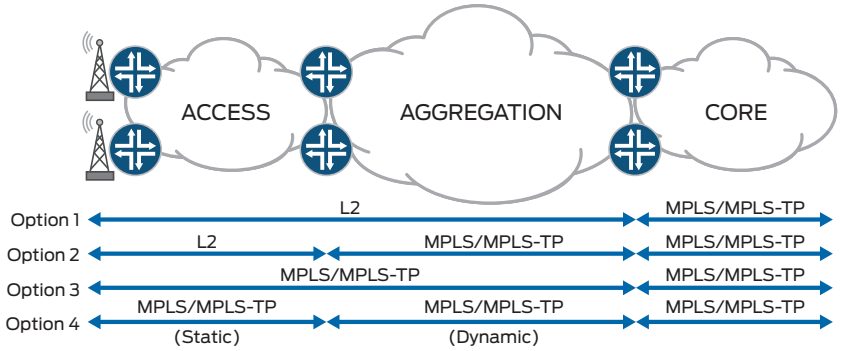

Figure 3: MPLS and MPLS-TP Deployment Options

Standardization History and Status

The effort to enhance MPLS for transport networks first started in ITU-T under the name T-MPLS, even though IETF has been the standardization body that developed the original MPLS technology. ITU-T soon recognized the harmfulness of uncoordinated protocol development [RFC 5704]. ITU-T and IETF jointly decided that standardization of an MPLS-based transport network should be progressed through the IETF standards process [Joint Working Team Report – RFC 5317]. ITU-T recognized the IETF as the design authority for MPLS and agreed to terminate all work on T-MPLS. A joint team called MPLS Interoperability Design Team (MEAD) was formed to develop the MPLS-TP standard within IETF soon after. The scope of this joint activity is to develop a set of mutually agreed upon requirements, and to develop protocols to satisfy those requirements by following the IETF standards process. The joint requirements have been captured in the IETF documents listed in the Bibliographic Citation section of this paper.

The IETF process is documented in draft-ietf-mpls-tp-process. IETF has made significant progress on the MPLS-TP work. The requirements and frameworks have been published as RFCs already, and most of the important drafts are in implementable form. Juniper is actively contributing to both ITU-T and IETF, as well as driving the direction of MPLS-TP and MPLS within the IETF MEAD team. Appendix A identifies the status of each of the elements of the MPLS-TP profile standards.

Misconceptions about MPLS-TP

MPLS-TP has been the subject of heated debate and is at the center of much standardization activity. As a result, there are a lot of misconceptions about MPLS-TP. This section summarizes some of the most popular misconceptions, along with the reasons for their perpetuation.

The first false claim is that MPLS-TP is a new technology and it is not part of the MPLS umbrella. In fact, as shown in Figure 1, MPLS-TP is a subset of MPLS with a few extensions. A related claim is that the extensions introduced by MPLS-TP are not applicable to MPLS. In fact, these extensions, such as the transport-like OAM functions, are meant to apply generally to MPLS and will make MPLS protocol suite broader so that it is applicable to both transport and service networks. Yet another misconception is that MPLS-TP requires substantial changes to MPLS. One of the design goals of MPLS-TP is to keep the MPLS architecture intact and reuse as many of the existing components of MPLS as possible.

Another claim often put forth is that MPLS-TP requires static provisioning as transport networks relay heavily on static provisioning. As can be seen in the MPLS-TP requirements documents, MPLS-TP supports both static and dynamic control planes.

Yet another claim that is often made is that MPLS-TP requires forklift upgrades from the existing hardware. Even though this might be true for some vendors, Juniper strongly believes in delivering investment protection. All Juniper products are powered by one Juniper Networks Junos® operating system. Once support for MPLS-TP enhancements is added into Junos OS, it will be available on all relevant hardware platforms.

Conclusion

MPLS-TP is a set of enhancements to the already rich MPLS protocol suite. The current MPLS suite has successfully served packet-based networks for more than a decade. The MPLS-TP enhancements will increase the scope of MPLS overall, allowing it to serve both the transport and the services networks.

The biggest and most important enhancements that are being developed under the MPLS-TP effort are OAM related (e.g., fault management and performance monitoring). These OAM enhancements will prove to be very valuable for the existing MPLS networks, as they will allow operators to improve the efficiency and effectiveness of their networks by enabling full end-to-end integration with the existing and the next-generation MPLS networks.

Juniper Networks is committed to addressing the requirements of the packet-based transport networks, and we are actively driving the advancement of MPLS and MPLS-TP in the standards bodies. Juniper Networks Junos operating system already supports most of the existing MPLS/GMPLS technology included in MPLS-TP, and will add IETF standardized MPLS-TP enhancements to its rich MPLS technology portfolio, with an emphasis on extending the functionality in areas of OAM for measuring delay and loss.

Juniper is also working on enhancing MPLS outside the scope of MPLS-TP, so that MPLS can scale to the order of hundreds of thousands of nodes, a scale that is required for the next-generation access and aggregation networks. This work falls under the umbrella of seamless MPLS, which will be very vital to the next-generation networks.

Bibliographic Citation

1. “IETF MPLS-TP Document Process,” Draft draft-ietf-mpls-tp-process-05.
2. “Joint Working Team (JWT) Report on MPLS Architectural Considerations for a Transport Profile,” RFC 5317.
3. “Uncoordinated Protocol Development Considered Harmful,” RFC 5704.
4. “EXP field renamed to Traffic Class field,” RFC 5462.
5. “MPLS-TP Requirements,” RFC 5654.
6. “An Inband Data Communication Network for the MPLS-TP,” RFC 5718.
7. “Requirements for OAM in MPLS Transport Networks,” RFC 5860.
8. “MPLS TP Network Management Framework,” RFC 5950.
9. “Network Management Requirements for MPLS-TP,” RFC 5951.
10. “MPLS Transport Profile Data Plane Architecture,” RFC 5960.
11. “A Framework for MPLS in Transport Networks,” RFC 5921.
12. “MPLS Generic Associated Channel,” RFC 5586.
13. “Architecture of MPLS-TP Layer Network,” ITU-T Recommendation G.8110.1.
14. “Interfaces for the MPLS-TP Hierarchy,” ITU-T Recommendation G.8112.
15. “Characteristics of MPLS-TP Network Equipment Functional Blocks,” ITU-T Recommendation G.8121.
16. “MPLS-TP Linear Protection,” ITU-T Recommendation G.8131.
17. “MPLS-TP Ring Protection,” ITU-T Recommendation G.8132.
18. Juniper Networks, white paper, “Seamless MPLS,” 2010.

Appendix A: Current MPLS-TP Working Group (WG) Document Status

DOCUMENT TYPE	DOCUMENT NAME	IETF STATUS	COMPLIANCY REQUIREMENT
RFCs	**RFC 5317**: JWT Report on MPLS Architectural Considerations for a Transport Profile	Published	Compliance not needed for requirements, framework, or architecture RFCs.
	RFC 5462: EXP Field Renamed to Traffic Class Field		
	RFC 5654: MPLS-TP Requirements		
	RFC 5718: An Inband Data Communication Network for the MPLS-TP		
	RFC 5860: Requirements for OAM in MPLS Transport Networks		
	RFC 5950: MPLS-TP Network Management Framework		
	RFC 5951: Network Management Requirements for MPLS-TP		
	RFC 5960: MPLS Transport Profile Data Plane Architecture		
	RFC 5921: A Framework for MPLS in Transport Networks		
	RFC 5586: MPLS Generic Associated Channel	Published	
WG-Drafts*	**draft-ietf-mpls-tp-survive-fwk**: Multiprotocol Label Switching Transport Profile Survivability Framework	RFC Ed's Queue	Compliance not needed for framework RFCs.
	draft-ietf-mpls-tp-oam-framework: OAM Framework for MPLS-TP	RFC Ed's Queue	
WG-Drafts*	**draft-ietf-mpls-lsp-ping-mpls-tp-oam-conf**: Configuration of Proactive MPLS-TP OAM Functions Using LSP Ping	Active Draft	
	draft-ietf-mpls-tp-cc-cv-rdi: Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for MPLS-TP	Active Draft	
	draft-ietf-mpls-tp-fault: MPLS Fault Management OAM	Active Draft	
	draft-ietf-mpls-tp-li-lb: MPLS-TP Lock Instruct and Loopback Functions	Active Draft	
	draft-ietf-mpls-tp-linear-protection: MPLS-TP Linear Protection	Active Draft	
	draft-ietf-mpls-tp-loss-delay: Packet Loss and Delay Measurement for the MPLS-TP	Active Draft	
	draft-ietf-mpls-tp-lsp-ping-bfd-procedures: LSP-Ping and BFD Encapsulation over ACH	Active Draft	
	draft-ietf-mpls-tp-on-demand-cv: MPLS On-demand Connectivity Verification and Route Tracing	Active Draft	
	draft-ietf-mpls-tp-identifiers: MPLS-TP Identifiers	Active Draft	
Individual Drafts*	Around 35 individual drafts in MPLS, CCAMP, PWE3 WG	Not accepted by WG yet	Compliance needs to be considered only after the drafts are accepted by WGs.

* The individual drafts are ideas submitted to IETF by individuals and carry no endorsement from IETF. On the other hand, the WG drafts are the drafts that a WG believes are highly relevant and, therefore, have been accepted by WG within the IETF consensus process. The WG drafts are owned by the WG/IETF and not by the authors.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at **www.juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

2000406-001-EN Apr 2011

Printed on recycled paper