

Interworking Requirements to Support Operation of MPLS-TE
over GMPLS Networks

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

Operation of a Multiprotocol Label Switching (MPLS) traffic engineering (TE) network as a client network to a Generalized MPLS (GMPLS) network has enhanced operational capabilities compared to those provided by a coexistent protocol model (i.e., operation of MPLS-TE over an independently managed transport layer).

The GMPLS network may be a packet or a non-packet network, and may itself be a multi-layer network supporting both packet and non-packet technologies. An MPLS-TE Label Switched Path (LSP) originates and terminates on an MPLS Label Switching Router (LSR). The GMPLS network provides transparent transport for the end-to-end MPLS-TE LSP.

This document describes a framework and Service Provider requirements for operating MPLS-TE networks over GMPLS networks.

Table of Contents

1. Introduction	3
1.1. Terminology	4
2. Reference Model	4
3. Detailed Requirements	5
3.1. End-to-End Signaling	5
3.2. Triggered Establishment of GMPLS LSPs	5
3.3. Diverse Paths for End-to-End MPLS-TE LSPs	6
3.4. Advertisement of MPLS-TE Information via the GMPLS Network	6
3.5. Selective Advertisement of MPLS-TE Information via a Border Node	6
3.6. Interworking of MPLS-TE and GMPLS Protection	7
3.7. Independent Failure Recovery and Reoptimization	7
3.8. Complexity and Risks	7
3.9. Scalability Considerations	7
3.10. Performance Considerations	8
3.11. Management Considerations	8
4. Security Considerations	8
5. Recommended Solution Architecture	9
5.1. Use of Contiguous, Hierarchical, and Stitched LSPs	10
5.2. MPLS-TE Control Plane Connectivity	10
5.3. Fast Reroute Protection	10
5.4. GMPLS LSP Advertisement	11
5.5. GMPLS Deployment Considerations	11
6. Acknowledgments	11
7. References	11
7.1. Normative References	11
7.2. Informative References	12
8. Contributors' Addresses	13

1. Introduction

Multiprotocol Label Switching traffic engineering (MPLS-TE) networks are often deployed over transport networks such that the transport networks provide connectivity between the Label Switching Routers (LSRs) in the MPLS-TE network. Increasingly, these transport networks are operated using a Generalized Multiprotocol Label Switching (GMPLS) control plane. Label Switched Paths (LSPs) in the GMPLS network provide connectivity as virtual data links advertised as TE links in the MPLS-TE network.

GMPLS protocols were developed as extensions to MPLS-TE protocols. MPLS-TE is limited to the control of packet switching networks, but GMPLS can also control technologies at layers one and two.

The GMPLS network may be managed by an operator as a separate network (as it may have been when it was under management plane control before the use of GMPLS as a control plane), but optimizations of management and operation may be achieved by coordinating the use of the MPLS-TE and GMPLS networks and operating the two networks with a close client/server relationship.

GMPLS LSP setup may be triggered by the signaling of MPLS-TE LSPs in the MPLS-TE network so that the GMPLS network is reactive to the needs of the MPLS-TE network. The triggering process can be under the control of operator policies without needing direct intervention by an operator.

The client/server configuration just described can also apply in migration scenarios for MPLS-TE packet switching networks that are being migrated to be under GMPLS control. [RFC5145] describes a migration scenario called the Island Model. In this scenario, groups of nodes (islands) are migrated from the MPLS-TE protocols to the GMPLS protocols and operate entirely surrounded by MPLS-TE nodes (the sea). This scenario can be effectively managed as a client/server network relationship using the framework described in this document.

In order to correctly manage the dynamic interaction between the MPLS and GMPLS networks, it is necessary to understand the operational requirements and the control that the operator can impose. Although this problem is very similar to the multi-layer networks described in [MLN-REQ], it must be noted that those networks operate GMPLS protocols in both the client and server networks, which facilitates smoother interworking. Where the client network uses MPLS-TE protocols over the GMPLS server network, there is a need to study the interworking of the two protocol sets.

This document examines the protocol requirements for protocol interworking to operate an MPLS-TE network as a client network over a GMPLS server network, and provides a framework for such operations.

1.1. Terminology

Although this Informational document is not a protocol specification, 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 [RFC2119] for clarity of exposure of the requirements.

2. Reference Model

The reference model used in this document is shown in Figure 1. It can easily be seen that the interworking between MPLS-TE and GMPLS protocols must occur on a node and not on a link. Nodes on the interface between the MPLS-TE and GMPLS networks must be responsible for handling both protocol sets and for providing any protocol interworking that is required. We call these nodes Border Routers.




Figure 1. Reference model of MPLS-TE/GMPLS interworking

MPLS-TE network connectivity is provided through a GMPLS LSP which is created between Border Routers. End-to-end connectivity between MPLS LSRs in the client MPLS-TE networks is provided by an MPLS-TE LSP that is carried across the MPLS-TE network by the GMPLS LSP using hierarchical LSP techniques [RFC4206], LSP stitching segments

[RFC5150], or a contiguous LSP. LSP stitching segments and contiguous LSPs are only available where the GMPLS network is a packet switching network.

3. Detailed Requirements

This section describes detailed requirements for MPLS-TE/GMPLS interworking in support of the reference model shown in Figure 1.

The functional requirements for GMPLS-MPLS interworking described in this section must be met by any device participating in the interworking. This may include routers, servers, network management devices, path computation elements, etc.

3.1. End-to-End Signaling

The solution MUST be able to preserve MPLS signaling information signaled within the MPLS-TE client network at the start of the MPLS-TE LSP and deliver it on the other side of the GMPLS server network for use within the MPLS-TE client network at the end of the MPLS-TE LSP. This may require protocol mapping (and re-mapping), protocol tunneling, or the use of remote protocol adjacencies.

3.2. Triggered Establishment of GMPLS LSPs

The solution MUST provide the ability to establish end-to-end MPLS-TE LSPs over a GMPLS server network. It SHOULD be possible for GMPLS LSPs across the core network to be set up between Border Routers triggered by the signaling of MPLS-TE LSPs in the client network, and in this case, policy controls MUST be made available at the border routers so that the operator of the GMPLS network can manage how core network resources are utilized. GMPLS LSPs MAY also be pre-established as the result of management plane control.

Note that multiple GMPLS LSPs may be set up between a given pair of Border Routers in support of connectivity in the MPLS client network. If these LSPs are advertised as TE links in the client network, the use of link bundling [RFC4201] can reduce any scaling concerns associated with the advertisements.

The application of the Path Computation Element (PCE) [RFC4655] in the context of an inter-layer network [PCE-INT] may be considered to determine an end-to-end LSP with triggered GMPLS segment or tunnel.

3.3. Diverse Paths for End-to-End MPLS-TE LSPs

The solution SHOULD provide the ability to establish end-to-end MPLS-TE LSPs having diverse paths for protection of the LSP traffic. This means that MPLS-TE LSPs SHOULD be kept diverse both within the client MPLS-TE network and as they cross the server GMPLS network. This means that there SHOULD be a mechanism to request the provision of diverse GMPLS LSPs between a pair of Border Routers to provide protection of the GMPLS span, but also that there SHOULD be a way to keep GMPLS LSPs between different Border Routers disjoint.

3.4. Advertisement of MPLS-TE Information via the GMPLS Network

The solution SHOULD provide the ability to exchange advertisements of TE information between MPLS-TE client networks across the GMPLS server network.

The advertisement of TE information from within an MPLS-TE client network to all LSRs in the client network enables a head-end LSR to compute an optimal path for an LSP to a tail-end LSR that is reached over the GMPLS server network.

Where there is more than one client MPLS-TE network, the TE information from separate MPLS-TE networks MUST be kept private, confidential and secure.

3.5. Selective Advertisement of MPLS-TE Information via a Border Node

The solution SHOULD provide the ability to distribute TE reachability information from the GMPLS server network to MPLS-TE networks selectively. This information is useful for the LSRs in the MPLS-TE networks to compute paths that cross the GMPLS server network and to select the correct Border Routers to provide connectivity.

The solution MUST NOT distribute TE information from within a non-PSC (Packet Switch Capable) GMPLS server network to any client MPLS-TE network as that information may cause confusion and selection of inappropriate paths.

The use of PCE [RFC4655] may provide a solution for non-PSC GMPLS networks supporting PSC MPLS networks.

3.6. Interworking of MPLS-TE and GMPLS Protection

If an MPLS-TE LSP is protected using MPLS Fast Reroute (FRR) [RFC4090], then similar protection MUST be provided over the GMPLS island. Operator and policy controls SHOULD be made available at the Border Router to determine how suitable protection is provided in the GMPLS island.

3.7. Independent Failure Recovery and Reoptimization

The solution SHOULD provide failure recovery and reoptimization in the GMPLS server network without impacting the MPLS-TE client network and vice versa. That is, it SHOULD be possible to recover from a fault within the GMPLS island or to reoptimize the path across the GMPLS island without requiring signaling activity within the MPLS-TE client network. Similarly, it SHOULD be possible to perform recovery or reoptimization within the MPLS-TE client network without requiring signaling activity within the GMPLS server networks.

If a failure in the GMPLS server network can not be repaired transparently, some kind of notification of the failure SHOULD be transmitted to MPLS-TE network.

3.8. Complexity and Risks

The solution SHOULD NOT introduce unnecessary complexity to the current operating network to such a degree that it would affect the stability and diminish the benefits of deploying such a solution in service provider networks.

3.9. Scalability Considerations

The solution MUST scale well with consideration to at least the following metrics.

- The number of GMPLS-capable nodes (i.e., the size of the GMPLS server network).
- The number of MPLS-TE-capable nodes (i.e., the size of the MPLS-TE client network).
- The number of MPLS-TE client networks.
- The number of GMPLS LSPs.
- The number of MPLS-TE LSPs.

3.10. Performance Considerations

The solution SHOULD be evaluated with regard to the following criteria.

- Failure and restoration time.
- Impact and scalability of the control plane due to added overheads.
- Impact and scalability of the data/forwarding plane due to added overheads.

3.11. Management Considerations

Manageability of the deployment of an MPLS-TE client network over GMPLS server network MUST address the following considerations.

- Need for coordination of MIB modules used for control plane management and monitoring in the client and server networks.
- Need for diagnostic tools that can discover and isolate faults across the border between the MPLS-TE client and GMPLS server networks.

4. Security Considerations

Border routers in the model described in this document are present on administrative domain boundaries. That is, the administrative boundary does not lie on a link as it might in the inter-Autonomous-System (inter-AS) case seen in IP networks. Thus, many security concerns for the inter-domain exchange of control plane messages do not arise in this model -- the border router participates fully in both the MPLS and the GMPLS network and must participate in the security procedures of both networks. Security considerations for MPLS-TE and GMPLS protocols are discussed in [SECURITY].

However, policy considerations at the border routers are very important and may be considered to form part of the security of the networks. In particular, the server network (the GMPLS network) may wish to protect itself from behavior in the client network (such as frequent demands to set up and tear down server LSPs) by appropriate policies implemented at the border routers. It should be observed that, because the border routers form part of both networks, they are trusted in both networks, and policies configured (whether locally or centrally) for use by a border router are expected to be observed.

Nevertheless, authentication and access controls for operators will be particularly important at border routers. Operators of the client

MPLS-TE network MUST NOT be allowed to configure the server GMPLS network (including setting server network policies), and operators of the server GMPLS network MUST NOT be able to configure the client MPLS-TE network. Obviously, it SHOULD be possible to grant an operator privileges in both networks. It may also be desirable to give operators of one network access to (for example) status information about the other network.

Mechanisms for authenticating operators and providing access controls are not part of the responsibilities of the GMPLS protocol set, and will depend on the management plane protocols and techniques implemented.

5. Recommended Solution Architecture

The recommended solution architecture to meet the requirements set out in Section 3 is known as the Border Peer Model. This architecture is a variant of the Augmented Model described in [RFC3945]. The remainder of this document presents an overview of this architecture.

In the Augmented Model, routing information from the lower layer (server) network is filtered at the interface to the higher layer (client) network and a subset of the information is distributed within the higher layer network.

In the Border Peer Model, the interface between the client and server networks is the Border Router. This router has visibility of the routing information in the server network yet also participates as a peer in the client network. Thus, the Border Router has full visibility into both networks. However, the Border Router does not distribute server routing information into the client network, nor does it distribute client routing information into the server network.

The Border Peer Model may also be contrasted with the Overlay Model [RFC3945]. In this model there is a protocol request/response interface (the user network interface (UNI)) between the client and server networks. [RFC4208] shows how this interface may be supported by GMPLS protocols operated between client edge and server edge routers while retaining the routing information within the server network. That is, in the Overlay Model there is no exchange of routing or reachability information between client and server networks, and no network element has visibility into both client and server networks. The Border Peer Model can be viewed as placing the UNI within the Border Router thus giving the Border Router peer capabilities in both the client and server network.

5.1. Use of Contiguous, Hierarchical, and Stitched LSPs

All three LSP types MAY be supported in the Border Peer Model, but contiguous LSPs are the hardest to support because they require protocol mapping between the MPLS-TE client network and the GMPLS server network. Such protocol mapping can be achieved currently since MPLS-TE signaling protocols are a subset of GMPLS, but this mechanism is not future-proofed.

Contiguous and stitched LSPs can only be supported where the GMPLS server network has the same switching type (that is, packet switching) as the MPLS-TE network. Requirements for independent failure recovery within the GMPLS island require the use of loose path reoptimization techniques [RFC4736] and end-to-end make-before-break [RFC3209], which will not provide rapid recovery.

For these reasons, the use of hierarchical LSPs across the server network is RECOMMENDED for the Border Peer Model, but see the discussion of Fast Reroute protection in Section 5.3.

5.2. MPLS-TE Control Plane Connectivity

Control plane connectivity between MPLS-TE LSRs connected by a GMPLS island in the Border Peer Model MAY be provided by the control channels of the GMPLS network. If this is done, a tunneling mechanism (such as GRE [RFC2784]) SHOULD be used to ensure that MPLS-TE information is not consumed by the GMPLS LSRs. But care is required to avoid swamping the control plane of the GMPLS network with MPLS-TE control plane (particularly routing) messages.

In order to ensure scalability, control plane messages for the MPLS-TE client network MAY be carried between Border Routers in a single hop MPLS-TE LSP routed through the data plane of the GMPLS server network.

5.3. Fast Reroute Protection

If the GMPLS network is packet switching, Fast Reroute protection can be offered on all hops of a contiguous LSP. If the GMPLS network is packet switching then all hops of a hierarchical GMPLS LSP or GMPLS stitching segment can be protected using Fast Reroute. If the end-to-end MPLS-TE LSP requests Fast Reroute protection, the GMPLS packet switching network SHOULD provide such protection.

However, note that it is not possible to provide FRR node protection of the upstream Border Router without careful consideration of available paths, and protection of the downstream Border Router is not possible where hierarchical LSPs or stitching segments are used.

Note further that Fast Reroute is not available in non-packet technologies. However, other protection techniques are supported by GMPLS for non-packet networks and are likely to provide similar levels of protection.

The limitations of FRR need careful consideration by the operator and may lead to the decision to provide end-to-end protection for the MPLS-TE LSP.

5.4. GMPLS LSP Advertisement

In the Border Peer Model, the LSPs established by the Border Routers in the GMPLS server network SHOULD be advertised in the MPLS-TE client network as real or virtual links. In case real links are advertised into the MPLS-TE client network, the Border Routers in the MPLS-TE client network MAY establish IGP neighbors. The Border Routers MAY automatically advertise the GMPLS LSPs when establishing them.

5.5. GMPLS Deployment Considerations

The Border Peer Model does not require the existing MPLS-TE client network to be GMPLS aware and does not affect the operation and management of the existing MPLS-TE client network. Only border routers need to be upgraded with the GMPLS functionality. In this fashion, the Border Peer Model renders itself for incremental deployment of the GMPLS server network, without requiring reconfiguration of existing areas/ASs, changing operation of IGP and BGP or software upgrade of the existing MPLS-TE client network.

6. Acknowledgments

The author would like to express thanks to Raymond Zhang, Adrian Farrel, and Deborah Brungard for their helpful and useful comments and feedback.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3209] 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.

- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.
- [RFC4201] Kompella, K., Rekhter, Y., and L. Berger, "Link Bundling in MPLS Traffic Engineering (TE)", RFC 4201, October 2005.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.
- [RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.
- [RFC5150] Ayyangar, A., Kompella, K., Vasseur, JP., and A. Farrel, "Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)", RFC 5150, February 2008.

7.2. Informative References

- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC4736] Vasseur, JP., Ed., Ikejiri, Y., and R. Zhang, "Reoptimization of Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Loosely Routed Label Switched Path (LSP)", RFC 4736, November 2006.
- [RFC5145] Shiomoto, K., Ed., "Framework for MPLS-TE to GMPLS Migration", RFC 5145, March 2008.

- [MLN-REQ] Shiomoto, K., Papadimitriou, D., Le Roux, J.L., Vigoureux, M., and D. Brungard, "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", Work in Progress, January 2008.
- [PCE-INT] Oki, E., Le Roux, J-L., and A. Farrel, "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering," Work in Progress, January 2008.
- [SECURITY] Fang, L., "Security Framework for MPLS and GMPLS Networks", Work in Progress, November 2007.

8. Contributors' Addresses

Tomohiro Otani
KDDI R&D Laboratories, Inc.
2-1-15 Ohara Kamifukuoka
Saitama, 356-8502, Japan

Phone: +81-49-278-7357
EMail: otani@kddilabs.jp

Shuichi Okamoto
NICT JGN II Tsukuba Reserach Center
1-8-1, Otemachi Chiyoda-ku,
Tokyo, 100-0004, Japan

Phone: +81-3-5200-2117
EMail: okamoto-s@nict.go.jp

Kazuhiro Fujihara
NTT Communications Corporation
Tokyo Opera City Tower 3-20-2 Nishi Shinjuku, Shinjuku-ku
Tokyo 163-1421, Japan

EMail: kazuhiro.fujihara@ntt.com

Yuichi Ikejiri
NTT Communications Corporation
Tokyo Opera City Tower 3-20-2 Nishi Shinjuku, Shinjuku-ku
Tokyo 163-1421, Japan

EMail: y.ikejiri@ntt.com

Editor's Address

Kenji Kumaki
KDDI Corporation
Garden Air Tower
Iidabashi, Chiyoda-ku,
Tokyo, 102-8460, JAPAN

EMail: ke-kumaki@kddi.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

