

Framework for MPLS-TE to GMPLS Migration

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

The migration from Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) to Generalized MPLS (GMPLS) is the process of evolving an MPLS-TE control plane to a GMPLS control plane. An appropriate migration strategy will be selected based on various factors including the service provider's network deployment plan, customer demand, and operational policy.

This document presents several migration models and strategies for migrating from MPLS-TE to GMPLS. In the course of migration, MPLS-TE and GMPLS devices, or networks, may coexist that may require interworking between MPLS-TE and GMPLS protocols. Aspects of the required interworking are discussed as it will influence the choice of a migration strategy. This framework document provides a migration toolkit to aid the operator in selection of an appropriate strategy.

This framework document also lists a set of solutions that may aid in interworking, and highlights a set of potential issues.

Table of Contents

1. Introduction	3
2. Conventions Used in This Document	3
3. Motivations for Migration	4
4. MPLS to GMPLS Migration Models	5
4.1. Island Model	5
4.1.1. Balanced Islands	6
4.1.2. Unbalanced Islands	6
4.2. Integrated Model	7
4.3. Phased Model	8
5. Migration Strategies and Toolkit	8
5.1. Migration Toolkit	9
5.1.1. Layered Networks	9
5.1.2. Routing Interworking	11
5.1.3. Signaling Interworking	12
5.1.4. Path Computation Element	13
6. Manageability Considerations	13
6.1. Control of Function and Policy	13
6.2. Information and Data Models	14
6.3. Liveness Detection and Monitoring	14
6.4. Verifying Correct Operation	14
6.5. Requirements on Other Protocols and Functional Components	14
6.6. Impact on Network Operation	15
6.7. Other Considerations	15
7. Security Considerations	15
8. Acknowledgements	16
9. References	16
9.1. Normative References	16
9.2. Informative References	17
10. Contributors' Addresses	17

1. Introduction

Multiprotocol Label Switching Traffic Engineering (MPLS-TE) to Generalized MPLS (GMPLS) migration is the process of evolving an MPLS-TE-based control plane to a GMPLS-based control plane. The network under consideration for migration is, therefore, a packet-switching network.

There are several motivations for such migration, mainly the desire to take advantage of new features and functions added to the GMPLS protocols, which are not present in MPLS-TE for packet networks. Additionally, before migrating a packet-switching network from MPLS-TE to GMPLS, one may choose to first migrate a lower-layer network with no control plane (e.g., controlled by a management plane) to using a GMPLS control plane. This may lead to the desire for MPLS-TE/GMPLS (transport network) interworking to provide enhanced TE support and facilitate the later migration of the packet-switching network.

Although an appropriate migration strategy will be selected based on various factors including the service provider's network deployment plan, customer demand, deployed network equipments, operational policy, etc., the transition mechanisms used should also provide consistent operation of newly introduced GMPLS networks, while minimizing the impact on the operation of existing MPLS-TE networks.

This document describes several migration strategies and the interworking scenarios that arise during migration. It also examines the implications for network deployments and for protocol usage. As the GMPLS signaling and routing protocols are different from the MPLS-TE control protocols, interworking mechanisms between MPLS-TE and GMPLS networks, or network elements, may be needed to compensate for the differences.

Note that MPLS-TE and GMPLS protocols can coexist as "ships in the night" without any interworking issues.

2. Conventions Used in This Document

This is not a requirements document, nevertheless 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 [RFC2119] in order to clarify the recommendations that are made.

In the rest of this document, the term "GMPLS" includes both packet switching capable (PSC) and non-PSC. Otherwise, the term "PSC GMPLS" or "non-PSC GMPLS" is used explicitly.

In general, the term "MPLS" is used to indicate MPLS traffic engineering (MPLS-TE) only ([RFC3209], [RFC3630], and [RFC3784]) and excludes other MPLS protocols, such as the Label Distribution Protocol (LDP). TE functionalities of MPLS could be migrated to GMPLS, but non-TE functionalities could not. If non-TE MPLS is intended, it is indicated explicitly.

The reader is assumed to be familiar with the terminology introduced in [RFC3945].

3. Motivations for Migration

Motivations for migration will vary for different service providers. This section is presented to provide background so that the migration discussions may be seen in context. Sections 4 and 5 provide examples to illustrate the migration models and processes.

Migration of an MPLS-capable Label Switching Router (LSR) to include GMPLS capabilities may be performed for one or more reasons, including, not exhaustively:

- o To add all GMPLS PSC features to an existing MPLS network (upgrade MPLS LSRs).
- o To add specific GMPLS PSC features and operate them within an MPLS network (e.g., [RFC4872] and [RFC4873]).
- o To integrate a new GMPLS PSC network with an existing MPLS network (without upgrading any of the MPLS LSRs).
- o To allow existing MPLS LSRs to interoperate with new non-MPLS LSRs supporting only GMPLS PSC and/or non-PSC features.
- o To integrate multiple control networks, e.g., managed by separate administrative organizations, and which independently utilize MPLS or GMPLS.
- o To build integrated PSC and non-PSC networks. The non-PSC networks are controlled by GMPLS.

The objective of migration from MPLS to GMPLS is that all LSRs, and the entire network, support GMPLS protocols. During this process, various interim situations may exist, giving rise to the interworking situations described in this document. The interim situations may exist for considerable periods of time, but the ultimate objective is not to preserve these situations. For the purposes of this document, they should be considered as temporary and transitory.

4. MPLS to GMPLS Migration Models

Three reference migration models are described below. Multiple migration models may coexist in the same network.

4.1. Island Model

In the island model, "islands" of network nodes operating one protocol exist within a "sea" of nodes using the other protocol.

For example, consider an island of GMPLS-capable nodes (PSC) that is introduced into a legacy MPLS network. Such an island might be composed of newly added GMPLS nodes, or it might arise from the upgrade of existing nodes that previously operated MPLS protocols.

The opposite is also quite possible. That is, there is a possibility that an island happens to be MPLS-capable within a GMPLS sea. Such a situation might arise in the later stages of migration, when all but a few islands of MPLS-capable nodes have been upgraded to GMPLS.

It is also possible that a lower-layer, manually-provisioned network (for example, a Time Division Multiplexing (TDM) network) is constructed under an MPLS PSC network. During the process of migrating both networks to GMPLS, the lower-layer network might be migrated first. This would appear as a GMPLS island within an MPLS sea.

Lastly, it is possible to consider individual nodes as islands. That is, it would be possible to upgrade or insert an individual GMPLS-capable node within an MPLS network, and to treat that GMPLS node as an island.

Over time, collections of MPLS devices are replaced or upgraded to create new GMPLS islands or to extend existing ones, and distinct GMPLS islands may be joined together until the whole network is GMPLS-capable.

From a migration/interworking point of view, we need to examine how these islands are positioned and how Label Switched Paths (LSPs) connect between the islands.

Four categories of interworking scenarios are considered: (1) MPLS-GMPLS-MPLS, (2) GMPLS-MPLS-GMPLS, (3) MPLS-GMPLS, and (4) GMPLS-MPLS. In case 1, the interworking behavior is examined based on whether the GMPLS islands are PSC or non-PSC.

Figure 1 shows an example of the island model for MPLS-GMPLS-MPLS interworking. The model consists of a transit GMPLS island in an MPLS sea. The nodes at the boundary of the GMPLS island (G1, G2, G5, and G6) are referred to as "island border nodes". If the GMPLS island was non-PSC, all nodes except the island border nodes in the GMPLS-based transit island (G3 and G4) would be non-PSC devices, i.e., optical equipment (TDM, Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC)).




Figure 1: Example of the island model for MPLS-GMPLS-MPLS interworking

4.1.1. Balanced Islands

In the MPLS-GMPLS-MPLS and GMPLS-MPLS-GMPLS cases, LSPs start and end using the same protocols. Possible strategies include:

- tunneling the signaling across the island network using LSP nesting or stitching [RFC5150] (the latter is only for GMPLS-PSC)
- protocol interworking or mapping (both are only for GMPLS-PSC)

4.1.2. Unbalanced Islands

As previously discussed, there are two island interworking models that support bordering islands. GMPLS(PSC)-MPLS and MPLS-GMPLS(PSC) island cases are likely to arise where the migration strategy is not based on a core infrastructure, but has edge nodes (ingress or egress) located in islands of different capabilities.

In this case, an LSP starts or ends in a GMPLS (PSC) island and correspondingly ends or starts in an MPLS island. This mode of operation can only be addressed using protocol interworking or

mapping. Figure 2 shows the reference model for this migration scenario. Head-end and tail-end LSRs are in distinct control plane clouds.

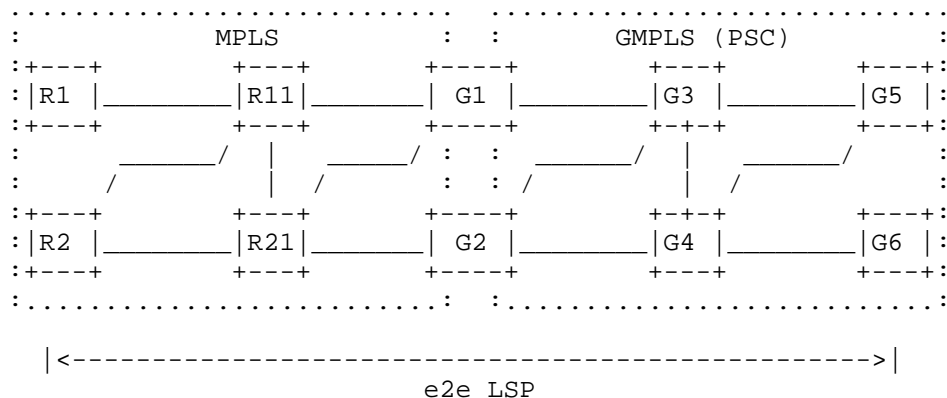


Figure 2: GMPLS-MPLS interworking model

It is important to underline that this scenario is also impacted by the directionality of the LSP, and the direction in which the LSP is established.

4.2. Integrated Model

The second migration model involves a more integrated migration strategy. New devices that are capable of operating both MPLS and GMPLS protocols are introduced into the MPLS network.

In the integrated model, there are two types of nodes present during migration:

- those that support MPLS only (legacy nodes); and
- those that support MPLS and GMPLS.

In this model, as existing MPLS devices are upgraded to support both MPLS and GMPLS, the network continues to operate with an MPLS control plane, but some LSRs are also capable of operating with a GMPLS control plane. So, LSPs are provisioned using MPLS protocols where one end point of a service is a legacy MPLS node and/or where the selected path between end points traverses a legacy node that is not GMPLS-capable. But where the service can be provided using only GMPLS-capable nodes [RFC5073], it may be routed accordingly and can achieve a higher level of functionality by utilizing GMPLS features.

Once all devices in the network are GMPLS-capable, the MPLS-specific protocol elements may be turned off, and no new devices need to support these protocol elements.

In this model, the questions to be addressed concern the coexistence of the two protocol sets within the network. Actual interworking is not a concern.

4.3. Phased Model

The phased model introduces GMPLS features and protocol elements into an MPLS network one by one. For example, some objects or sub-objects (such as the Explicit Route Object (ERO) label sub-object, [RFC3473]) might be introduced into the signaling used by LSRs that are otherwise MPLS-capable. This would produce a kind of hybrid LSR.

This approach may appear simpler to implement as one is able to quickly and easily pick up new key functions without needing to upgrade the whole protocol implementation. It is most likely to be used where there is a desire to rapidly implement a particular function within a network without the necessity to install and test the full GMPLS function.

Interoperability concerns though are exacerbated by this migration model, unless all LSRs in the network are updated simultaneously and there is a clear understanding of which subset of features are to be included in the hybrid LSRs. Interworking between a hybrid LSR and an unchanged MPLS LSR would put the hybrid LSR in the role of a GMPLS LSR, as described in the previous sections, and puts the unchanged LSR in the role of an MPLS LSR. The potential for different hybrids within the network will complicate matters considerably. This model is, therefore, only appropriate for use when the set of new features to be deployed is well known and limited, and where there is a clear understanding of and agreement on this set of features by the network operators of the ISP(s) involved as well as all vendors whose equipment will be involved in the migration.

5. Migration Strategies and Toolkit

An appropriate migration strategy is selected by a network operator based on factors including the service provider's network deployment plan, customer demand, existing network equipment, operational policy, support from its vendors, etc.

For PSC networks, the migration strategy involves the selection between the models described in the previous section. The choice will depend upon the final objective (full GMPLS capability, partial

upgrade to include specific GMPLS features, or no change to existing IP/MPLS networks), and upon the immediate objectives (full, phased, or staged upgrade).

For PSC networks serviced by non-PSC networks, two basic migration strategies can be considered. In the first strategy, the non-PSC network is made GMPLS-capable, first, and then the PSC network is migrated to GMPLS. This might arise when, in order to expand the network capacity, GMPLS-based non-PSC sub-networks are introduced into the legacy MPLS-based networks. Subsequently, the legacy MPLS-based PSC network is migrated to be GMPLS-capable, as described in the previous paragraph. Finally, the entire network, including both PSC and non-PSC nodes, may be controlled by GMPLS.

The second strategy is to migrate the PSC network to GMPLS first, and then enable GMPLS within the non-PSC network. The PSC network is migrated as described before, and when the entire PSC network is completely converted to GMPLS, GMPLS-based non-PSC devices and networks may be introduced without any issues of interworking between MPLS and GMPLS.

These migration strategies and the migration models described in the previous section are not necessarily mutually exclusive. Mixtures of all strategies and models could be applied. The migration models and strategies selected will give rise to one or more of the interworking cases described in the following section.

5.1. Migration Toolkit

As described in the previous sections, an essential part of a migration and deployment strategy is how the MPLS and GMPLS or hybrid LSRs interwork. This section sets out some of the alternatives for achieving interworking between MPLS and GMPLS, and it identifies some of the issues that need to be addressed. This document does not describe solutions to these issues.

Note that it is possible to consider upgrading the routing and signaling capabilities of LSRs from MPLS to GMPLS separately.

5.1.1. Layered Networks

In the balanced island model, LSP tunnels [RFC4206] are a solution to carry the end-to-end LSPs across islands of incompatible nodes. Network layering is often used to separate domains of different data plane technology. It can also be used to separate domains of different control plane technology (such as MPLS and GMPLS protocols), and the solutions developed for multiple data plane

technologies can be usefully applied to this situation [RFC3945], [RFC4206], and [RFC4726]. [MLN-REQ] gives a discussion of the requirements for multi-layered networks.

The GMPLS architecture [RFC3945] identifies three architectural models for supporting multi-layer GMPLS networks, and these models may be applied to the separation of MPLS and GMPLS control plane islands.

- In the peer model, both MPLS and GMPLS nodes run the same routing instance, and routing advertisements from within islands of one level of protocol support are distributed to the whole network. This is achievable only, as described in Section 5.1.2, either by direct distribution or by mapping of parameters.

Signaling in the peer model may result in contiguous LSPs, stitched LSPs [RFC5150] (only for GMPLS PSC), or nested LSPs. If the network islands are non-PSC, then the techniques of [MLN-REQ] may be applied, and these techniques may be extrapolated to networks where all nodes are PSC, but where there is a difference in signaling protocols.

- The overlay model preserves strict separation of routing information between network layers. This is suitable for the balanced island model, and there is no requirement to handle routing interworking. Even though the overlay model preserves separation of signaling information between network layers, there may be some interaction in signaling between network layers.

The overlay model requires the establishment of control plane connectivity for the higher layer across the lower layer.

- The augmented model allows limited routing exchange from the lower-layer network to the higher-layer network. Generally speaking, this assumes that the border nodes provide some form of filtering, mapping, or aggregation of routing information advertised from the lower-layer network. This architectural model can also be used for balanced island model migrations. Signaling interworking is required as described for the peer model.
- The border peer architecture model is defined in [RFC5146]. This is a modification of the augmented model where the layer border routers have visibility into both layers, but no routing information is otherwise exchanged between routing protocol instances. This architectural model is particularly suited to the MPLS-GMPLS-MPLS island model for PSC and non-PSC GMPLS islands. Signaling interworking is required as described for the peer model.

5.1.2. Routing Interworking

Migration strategies may necessitate some interworking between MPLS and GMPLS routing protocols. GMPLS extends the TE information advertised by the IGPs to include non-PSC information and extended PSC information. Because the GMPLS information is provided as additional TLVs that are carried along with the MPLS information, MPLS LSRs are able to "see" all GMPLS LSRs as though they were MPLS PSC LSRs. They will also see other GMPLS information, but will ignore it, flooding it transparently across the MPLS network for use by other GMPLS LSRs.

- Routing separation is achieved in the overlay and border peer models. This is convenient since only the border nodes need to be aware of the different protocol variants, and no mapping is required. It is suitable to the MPLS-GMPLS-MPLS and GMPLS-MPLS-GMPLS island migration models.
- Direct distribution involves the flooding of MPLS routing information into a GMPLS network, and GMPLS routing information into an MPLS network. The border nodes make no attempt to filter the information. This mode of operation relies on the fact that MPLS routers will ignore, but continue to flood, GMPLS routing information that they do not understand. The presence of additional GMPLS routing information will not interfere with the way that MPLS LSRs select routes. Although this is not a problem in a PSC-only network, it could cause problems in a peer architecture network that includes non-PSC nodes, as the MPLS nodes are not capable of determining the switching types of the other LSRs and will attempt to signal end-to-end LSPs assuming all LSRs to be PSC. This fact would require island border nodes to take triggered action to set up tunnels across islands of different switching capabilities.

GMPLS LSRs might be impacted by the absence of GMPLS-specific information in advertisements initiated by MPLS LSRs. Specific procedures might be required to ensure consistent behavior by GMPLS nodes. If this issue is addressed, then direct distribution can be used in all migration models (except the overlay and border peer architectural models where the problem does not arise).

- Protocol mapping converts routing advertisements so that they can be received in one protocol and transmitted in the other. For example, a GMPLS routing advertisement could have all of its GMPLS-specific information removed and could be flooded as an MPLS advertisement. This mode of interworking would require careful standardization of the correct behavior especially where an MPLS advertisement requires default values of GMPLS-specific fields to

be generated before the advertisement can be flooded further. There is also considerable risk of confusion in closely meshed networks where many LSRs have MPLS- and GMPLS-capable interfaces. This option for routing interworking during migration is NOT RECOMMENDED for any migration model. Note that converting GMPLS-specific sub-TLVs to MPLS-specific ones but not stripping the GMPLS-specific ones is considered a variant of the proposed solution in the previous bullet (unknown sub-TLVs should be ignored [RFC3630] but must continue to be flooded).

- Ships in the night refers to a mode of operation where both MPLS and GMPLS routing protocol variants are operated in the same network at the same time as separate routing protocol instances. The two instances are independent and are used to create routing adjacencies between LSRs of the same type. This mode of operation may be appropriate to the integrated migration model.

5.1.3. Signaling Interworking

Signaling protocols are used to establish LSPs and are the principal concern for interworking during migration. Issues of compatibility arise because of differences in the encodings and codepoints used by MPLS and GMPLS signaling, but also because of differences in functionality provided by MPLS and GMPLS.

- Tunneling and stitching [RFC5150] (GMPLS-PSC case) mechanisms provide the potential to avoid direct protocol interworking during migration in the island model because protocol elements are transported transparently across migration islands without being inspected. However, care may be needed to achieve functional mapping in these modes of operation since one set of features may need to be supported across a network designed to support a different set of features. In general, this is easily achieved for the MPLS-GMPLS-MPLS model, but may be hard to achieve in the GMPLS-MPLS-GMPLS model, for example, when end-to-end bidirectional LSPs are requested, since the MPLS island does not support bidirectional LSPs.

Note that tunneling and stitching are not available in unbalanced island models because in these cases, the LSP end points use different protocols.

- Protocol mapping is the conversion of signaling messages between MPLS and GMPLS. This mechanism requires careful documentation of the protocol fields and how they are mapped. This is relatively straightforward in the MPLS-GMPLS unbalanced island model for LSPs signaled in the MPLS-GMPLS direction. However, it may be more complex for LSPs signaled in the opposite direction, and this will

lead to considerable complications for providing GMPLS services over the MPLS island and for terminating those services at an egress LSR that is not GMPLS-capable. Further, in balanced island models, and in particular where there are multiple small (individual node) islands, the repeated conversion of signaling parameters may lead to loss of information (and functionality) or mis-requests.

- Ships in the night could be used in the integrated migration model to allow MPLS-capable LSRs to establish LSPs using MPLS signaling protocols and GMPLS LSRs to establish LSPs using GMPLS signaling protocols. LSRs that can handle both sets of protocols could work with both types of LSRs, and no conversion of protocols would be needed.

5.1.4. Path Computation Element

The Path Computation Element (PCE) [RFC4655] may provide an additional tool to aid MPLS to GMPLS migration. If a layered network approach (Section 5.1.1) is used, PCEs may be used to facilitate the computation of paths for LSPs in the different layers [PCE-INT].

6. Manageability Considerations

Attention should be given during migration planning to how the network will be managed during and after migration. For example, will the LSRs of different protocol capabilities be managed separately or as one management domain? For example, in the Island Model, it is possible to consider managing islands of one capability separately from the surrounding sea. In the case of islands that have different switching capabilities, it is possible that the islands already have separate management in place before the migration: the resultant migrated network may seek to merge the management or to preserve the separation.

6.1. Control of Function and Policy

The most critical control functionality to be applied is at the moment of changeover between different levels of protocol support. Such a change may be made without service halt or during a period of network maintenance.

Where island boundaries exist, it must be possible to manage the relationships between protocols and to indicate which interfaces support which protocols on a border LSR. Further, island borders are a natural place to apply policy, and management should allow configuration of such policies.

6.2. Information and Data Models

No special information or data models are required to support migration, but note that migration in the control plane implies migration from MPLS management tools to GMPLS management tools. During migration, therefore, it may be necessary for LSRs and management applications to support both MPLS and GMPLS management data.

The GMPLS MIB modules are designed to allow support of the MPLS protocols, and they are built on the MPLS MIB modules through extensions and augmentations. This may make it possible to migrate management applications ahead of the LSRs that they manage.

6.3. Liveness Detection and Monitoring

Migration will not impose additional issues for Operations, Administration, and Management (OAM) above those that already exist for inter-domain OAM and for OAM across multiple switching capabilities.

Note, however, that if a flat PSC MPLS network is migrated using the island model, and is treated as a layered network using tunnels to connect across GMPLS islands, then requirements for a multi-layer OAM technique may be introduced into what was previously defined in the flat OAM problem-space. The OAM framework of MPLS/GMPLS interworking will need further consideration.

6.4. Verifying Correct Operation

The concerns for verifying correct operation (and in particular, correct connectivity) are the same as for liveness detection and monitoring. Specifically, the process of migration may introduce tunneling or stitching [RFC5150] into what was previously a flat network.

6.5. Requirements on Other Protocols and Functional Components

No particular requirements are introduced on other protocols. As it has been observed, the management components may need to migrate in step with the control plane components, but this does not impact the management protocols, just the data that they carry.

It should also be observed that providing signaling and routing connectivity across a migration island in support of a layered architecture may require the use of protocol tunnels (such as Generic

Routing Encapsulation (GRE)) between island border nodes. Such tunnels may impose additional configuration requirements at the border nodes.

6.6. Impact on Network Operation

The process of migration is likely to have significant impact on network operation while migration is in progress. The main objective of migration planning should be to reduce the impact on network operation and on the services perceived by the network users.

To this end, planners should consider reducing the number of migration steps that they perform and minimizing the number of migration islands that are created.

A network manager may prefer the island model especially when migration will extend over a significant operational period because it allows the different network islands to be administered as separate management domains. This is particularly the case in the overlay, augmented network and border peer models where the details of the protocol islands remain hidden from the surrounding LSRs.

6.7. Other Considerations

A migration strategy may also imply moving an MPLS state to a GMPLS state for an in-service LSP. This may arise once all of the LSRs along the path of the LSP have been updated to be both MPLS- and GMPLS-capable. Signaling mechanisms to achieve the replacement of an MPLS LSP with a GMPLS LSP without disrupting traffic exist through make-before-break procedures [RFC3209] and [RFC3473], and should be carefully managed under operator control.

7. Security Considerations

Security and confidentiality is often applied (and attacked) at administrative boundaries. Some of the models described in this document introduce such boundaries, for example, between MPLS and GMPLS islands. These boundaries offer the possibility of applying or modifying the security as when crossing an IGP area or Autonomous System (AS) boundary, even though these island boundaries might lie within an IGP area or AS.

No changes are proposed to the security procedures built into MPLS and GMPLS signaling and routing. GMPLS signaling and routing inherit their security mechanisms from MPLS signaling and routing without any changes. Hence, there will be no additional issues with security in interworking scenarios. Further, since the MPLS and GMPLS signaling and routing security is provided on a hop-by-hop basis, and since all

signaling and routing exchanges described in this document for use between any pair of LSRs are based on either MPLS or GMPLS, there are no changes necessary to the security procedures.

8. Acknowledgements

The authors are grateful to Daisaku Shimazaki for discussion during the initial work on this document. The authors are grateful to Dean Cheng and Adrian Farrel for their valuable comments.

9. References

9.1. 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.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.
- [RFC3784] Smit, H. and T. Li, "Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)", RFC 3784, June 2004.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4872] Lang, J., Ed., Rekhter, Y., Ed., and D. Papadimitriou, Ed., "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, May 2007.
- [RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, May 2007.
- [RFC5073] Vasseur, J., Ed., and J. Le Roux, Ed., "IGP Routing Protocol Extensions for Discovery of Traffic Engineering Node Capabilities", RFC 5073, December 2007.

9.2. Informative References

- [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.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC4726] Farrel, A., Vasseur, J.-P., and A. Ayyangar, "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering", RFC 4726, November 2006.
- [RFC5150] Ayyangar, A., Kompella, A., Vasseur, JP., and A. Farrel, "Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering", RFC 5150, February 2008.
- [RFC5146] Kumaki, K., Ed., "Interworking Requirements to Support Operation of MPLS-TE over GMPLS Networks", RFC 5146, 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.

10. Contributors' Addresses

Dimitri Papadimitriou
Alcatel
Francis Wellensplein 1,
B-2018 Antwerpen, Belgium
Phone: +32 3 240 8491
EMail: dimitri.papadimitriou@alcatel-lucent.be

Jean-Louis Le Roux
France Telecom
av Pierre Marzin 22300
Lannion, France
Phone: +33 2 96 05 30 20
EMail: jeanlouis.leroux@orange-ftgroup.com

Deborah Brungard
AT&T
Rm. D1-3C22 - 200 S. Laurel Ave.
Middletown, NJ 07748, USA
Phone: +1 732 420 1573
EMail: dbrungard@att.com

Zafar Ali
Cisco Systems, Inc.
EMail: zali@cisco.com

Kenji Kumaki
KDDI Corporation
Garden Air Tower
Iidabashi, Chiyoda-ku,
Tokyo 102-8460, JAPAN
Phone: +81-3-6678-3103
EMail: ke-kumaki@kddi.com

Eiji Oki
NTT
Midori 3-9-11
Musashino, Tokyo 180-8585, Japan
Phone: +81 422 59 3441
EMail: oki.eiji@lab.ntt.co.jp

Ichiro Inoue
NTT
Midori 3-9-11
Musashino, Tokyo 180-8585, Japan
Phone: +81 422 59 3441
EMail: inoue.ichiro@lab.ntt.co.jp

Tomohiro Otani
KDDI Laboratories
EMail: otani@kddilabs.jp

Editor's Address

Kohei Shiomoto
NTT
Midori 3-9-11
Musashino, Tokyo 180-8585, Japan
Phone: +81 422 59 4402
EMail: shiomoto.kohei@lab.ntt.co.jp

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.