

Network Working Group
Request for Comments: 5339
Category: Informational

JL. Le Roux, Ed.
France Telecom
D. Papadimitriou, Ed.
Alcatel-Lucent
September 2008

Evaluation of Existing GMPLS Protocols
against Multi-Layer and Multi-Region Networks (MLN/MRN)

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

This document provides an evaluation of Generalized Multiprotocol Label Switching (GMPLS) protocols and mechanisms against the requirements for Multi-Layer Networks (MLNs) and Multi-Region Networks (MRNs). In addition, this document identifies areas where additional protocol extensions or procedures are needed to satisfy these requirements, and provides guidelines for potential extensions.

Table of Contents

1. Introduction	3
1.1. Conventions Used in This Document	4
2. MLN/MRN Requirements Overview	4
3. Analysis	5
3.1. Aspects of Multi-Layer Networks	5
3.1.1. Support for Virtual Network Topology Reconfiguration	5
3.1.1.1. Control of FA-LSPs Setup/Release	5
3.1.1.2. Virtual TE Links	6
3.1.1.3. Traffic Disruption Minimization during FA Release	8
3.1.1.4. Stability	8
3.1.2. Support for FA-LSP Attribute Inheritance	9
3.1.3. FA-LSP Connectivity Verification	9
3.1.4. Scalability	9
3.1.5. Operations and Management of the MLN/MRN	10
3.1.5.1. MIB Modules	10
3.1.5.2. OAM	11
3.2. Specific Aspects of Multi-Region Networks	12
3.2.1. Support for Multi-Region Signaling	12
3.2.2. Advertisement of Adjustment Capacities	13
4. Evaluation Conclusion	16
4.1. Traceability of Requirements	16
5. Security Considerations	20
6. Acknowledgments	20
7. References	21
7.1. Normative References	21
7.2. Informative References	21
8. Contributors' Addresses	23

1. Introduction

Generalized MPLS (GMPLS) extends MPLS to handle multiple switching technologies: packet switching, layer-2 switching, TDM (Time Division Multiplexing) switching, wavelength switching, and fiber switching (see [RFC3945]). The Interface Switching Capability (ISC) concept is introduced for these switching technologies and is designated as follows: PSC (Packet Switch Capable), L2SC (Layer-2 Switch Capable), TDM capable, LSC (Lambda Switch Capable), and FSC (Fiber Switch Capable). The representation, in a GMPLS control plane, of a switching technology domain is referred to as a region [RFC4206]. A switching type describes the ability of a node to forward data of a particular data plane technology, and uniquely identifies a network region.

A data plane switching layer describes a data plane switching granularity level. For example, LSC, TDM VC-11 and TDM VC-4-64c are three different layers. [RFC5212] defines a Multi-Layer Network (MLN) to be a Traffic Engineering (TE) domain comprising multiple data plane switching layers either of the same ISC (e.g., TDM) or different ISC (e.g., TDM and PSC) and controlled by a single GMPLS control plane instance. [RFC5212] further defines a particular case of MLNs. A Multi-Region Network (MRN) is defined as a TE domain supporting at least two different switching types (e.g., PSC and TDM), either hosted on the same device or on different ones, and under the control of a single GMPLS control plane instance.

The objectives of this document are to evaluate existing GMPLS mechanisms and protocols ([RFC3945], [RFC4202], [RFC3471], [RFC3473]) against the requirements for MLNs and MRNs, defined in [RFC5212]. From this evaluation, we identify several areas where additional protocol extensions and modifications are required in order to meet these requirements, and we provide guidelines for potential extensions.

A summary of MLN/MRN requirements is provided in Section 2. Then Section 3 evaluates whether current GMPLS protocols and mechanisms meet each of these requirements. When the requirements are not met by existing protocols, the document identifies whether the required mechanisms could rely on GMPLS protocols and procedure extensions, or whether it is entirely out of the scope of GMPLS protocols.

Note that this document specifically addresses GMPLS control plane functionality for MLN/MRN in the context of a single administrative control plane partition. Partitions of the control plane where separate layers are under distinct administrative control are for future study.

This document uses terminologies defined in [RFC3945], [RFC4206], and [RFC5212].

1.1. Conventions Used in This Document

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].

2. MLN/MRN Requirements Overview

Section 5 of [RFC5212] lists a set of functional requirements for Multi-Layer/Region Networks (MLN/MRN). These requirements are summarized below, and a mapping with sub-sections of [RFC5212] is provided.

Here is the list of requirements that apply to MLN (and thus to MRN):

- Support for robust Virtual Network Topology (VNT) reconfiguration. This implies the following requirements:
 - Optimal control of Forwarding Adjacency Label Switched Path (FA-LSP) setup and release (Section 5.8.1 of [RFC5212]);
 - Support for virtual TE links (Section 5.8.2 of [RFC5212]);
 - Minimization of traffic disruption during FA-LSP release (Section 5.5 of [RFC5212]);
 - Stability (Section 5.4 of [RFC5212]);
- Support for FA-LSP attribute inheritance (Section 5.6 of [RFC5212]);
- Support for FA-LSP data plane connectivity verification (Section 5.9 of [RFC5212]);
- MLN Scalability (Section 5.3 of [RFC5212]);
- MLN Operations and Management (OAM) (Section 5.10 of [RFC5212]);

Here is the list of requirements that apply to MRN only:

- Support for Multi-Region signaling (Section 5.7 of [RFC5212]);
- Advertisement of the adjustment capacity (Section 5.2 of [RFC5212]);

3. Analysis

3.1. Aspects of Multi-Layer Networks

3.1.1. Support for Virtual Network Topology Reconfiguration

A set of lower-layer FA-LSPs provides a Virtual Network Topology (VNT) to the upper-layer [RFC5212]. By reconfiguring the VNT (FA-LSP setup/release) according to traffic demands between source and destination node pairs within a layer, network performance factors (such as maximum link utilization and residual capacity of the network) can be optimized. Such optimal VNT reconfiguration implies several mechanisms that are analyzed in the following sections.

Note that the VNT approach is just one possible approach to performing inter-layer Traffic Engineering.

3.1.1.1. Control of FA-LSPs Setup/Release

In a Multi-Layer Network, FA-LSPs are created, modified, and released periodically according to the change of incoming traffic demands from the upper layer.

This implies a TE mechanism that takes into account the demands matrix, the TE topology, and potentially the current VNT, in order to compute and setup a new VNT.

Several functional building blocks are required to support such a TE mechanism:

- Discovery of TE topology and available resources.
- Collection of upper-layer traffic demands.
- Policing and scheduling of VNT resources with regard to traffic demands and usage (that is, decision to setup/release FA-LSPs). The functional component in charge of this function is called a VNT Manager (VNTM) [PCE-INTER].
- VNT Path Computation according to TE topology, potentially taking into account the old (existing) VNT in order to minimize changes. The functional component in charge of VNT computation may be distributed on network elements or may be performed on an external element (such as a Path Computation Element (PCE), [RFC4655]).
- FA-LSP setup/release.

GMPLS routing protocols provide TE topology discovery. GMPLS signaling protocols allow setting up/releasing FA-LSPs.

VNTM functions (resources policing/scheduling, decision to setup/release FA-LSPs, FA-LSP configuration) are out of the scope of GMPLS protocols. Such functionalities can be achieved directly on layer-border Label Switching Routers (LSRs), or through one or more external tools. When an external tool is used, an interface is required between the VNTM and the network elements so as to setup/release FA-LSPs. This could use standard management interfaces such as [RFC4802].

The set of traffic demands of the upper layer is required for the VNT Manager to take decisions to setup/release FA-LSPs. Such traffic demands include satisfied demands, for which one or more upper-layer LSP have been successfully setup, as well as unsatisfied demands and future demands, for which no upper layer LSP has been setup yet. The collection of such information is beyond the scope of GMPLS protocols. Note that it may be partially inferred from parameters carried in GMPLS signaling or advertised in GMPLS routing.

Finally, the computation of FA-LSPs that form the VNT can be performed directly on layer-border LSRs or on an external element (such as a Path Computation Element (PCE), [RFC4655]), and this is independent of the location of the VNTM.

Hence, to summarize, no GMPLS protocol extensions are required to control FA-LSP setup/release.

3.1.1.2. Virtual TE Links

A virtual TE link is a TE link between two upper layer nodes that is not actually associated with a fully provisioned FA-LSP in a lower layer. A virtual TE link represents the potentiality to setup an FA-LSP in the lower layer to support the TE link that has been advertised. A virtual TE link is advertised as any TE link, following the rules in [RFC4206] defined for fully provisioned TE links. In particular, the flooding scope of a virtual TE link is within an IGP area, as is the case for any TE link.

If an upper-layer LSP attempts (through a signaling message) to make use of a virtual TE link, the underlying FA-LSP is immediately signaled and provisioned (provided there are available resources in the lower layer) in the process known as triggered signaling.

The use of virtual TE links has two main advantages:

- Flexibility: allows the computation of an LSP path using TE links without needing to take into account the actual provisioning status of the corresponding FA-LSP in the lower layer;
- Stability: allows stability of TE links in the upper layer, while avoiding wastage of bandwidth in the lower layer, as data plane connections are not established until they are actually needed.

Virtual TE links are setup/deleted/modified dynamically, according to the change of the (forecast) traffic demand, operator's policies for capacity utilization, and the available resources in the lower layer.

The support of virtual TE links requires two main building blocks:

- A TE mechanism for dynamic modification of virtual TE link topology;
- A signaling mechanism for the dynamic setup and deletion of virtual TE links. Setting up a virtual TE link requires a signaling mechanism that allows an end-to-end association between virtual TE link end points with the purpose of exchanging link identifiers as well as some TE parameters.

The TE mechanism responsible for triggering/policing dynamic modification of virtual TE links is out of the scope of GMPLS protocols.

Current GMPLS signaling does not allow setting up and releasing virtual TE links. Hence, GMPLS signaling must be extended to support virtual TE links.

We can distinguish two options for setting up virtual TE links:

- The Soft FA approach consists of setting up the FA-LSP in the control plane without actually activating cross connections in the data plane. On the one hand, this requires state maintenance on all transit LSRs (N square issue), but on the other hand, this may allow for some admission control. Indeed, when a Soft FA is activated, the resources may no longer be available for use by other Soft FAs that have common links. These Soft FA will be dynamically released, and corresponding virtual TE links will be deleted. The Soft FA LSPs may be setup using procedures similar to those described in [RFC4872] for setting up secondary LSPs.

- The remote association approach simply consists of exchanging virtual TE link IDs and parameters directly between TE link end points. This does not require state maintenance on transit LSRs, but reduces admission control capabilities. Such an association between virtual TE link end points may rely on extensions to the Resource Reservation Protocol - Traffic Engineering (RSVP-TE) Automatically Switched Optical Network (ASON) call procedure [RFC4974].

Note that the support of virtual TE links does not require any GMPLS routing extension.

3.1.1.3. Traffic Disruption Minimization during FA Release

Before deleting a given FA-LSP, all nested LSPs have to be rerouted and removed from the FA-LSP to avoid traffic disruption. The mechanisms required here are similar to those required for graceful deletion of a TE link. A Graceful TE link deletion mechanism allows for the deletion of a TE link without disrupting traffic of TE-LSPs that were using the TE link.

Hence, GMPLS routing and/or signaling extensions are required to support graceful deletion of TE links. This may utilize the procedures described in [GR-SHUT]: a transit LSR notifies a head-end LSR that a TE link along the path of an LSP is going to be torn down, and also withdraws the bandwidth on the TE link so that it is not used for new LSPs.

3.1.1.4. Stability

The stability of upper-layer LSP may be impaired if the VNT undergoes frequent changes. In this context, robustness of the VNT is defined as the capability to smooth the impact of these changes and avoid their subsequent propagation.

Guaranteeing VNT stability is out of the scope of GMPLS protocols and relies entirely on the capability of the TE and VNT management algorithms to minimize routing perturbations. This requires that the algorithms take into account the old VNT when computing a new VNT, and try to minimize the perturbation.

Note that a full mesh of lower-layer LSPs may be created between every pair of border nodes between the upper and lower layers. The merit of a full mesh of lower-layer LSPs is that it provides stability to the upper-layer routing. That is, the forwarding table used in the upper layer is not impacted if the VNT undergoes changes. Further, there is always full reachability and immediate access to bandwidth to support LSPs in the upper layer. But it also has

significant drawbacks, since it requires the maintenance of n^2 RSVP-TE sessions (where n is the number of border nodes), which may be quite CPU- and memory-consuming (scalability impact). Also, this may lead to significant bandwidth wastage. Note that the use of virtual TE links solves the bandwidth wastage issue, and may reduce the control plane overload.

3.1.2. Support for FA-LSP Attribute Inheritance

When an FA TE Link is advertised, its parameters are inherited from the parameters of the FA-LSP, and specific inheritance rules are applied.

This relies on local procedures and policies and is out of the scope of GMPLS protocols. Note that this requires that both head-end and tail-end of the FA-LSP are driven by same policies.

3.1.3. FA-LSP Connectivity Verification

Once fully provisioned, FA-LSP liveness may be achieved by verifying its data plane connectivity.

FA-LSP connectivity verification relies on technology-specific mechanisms (e.g., for SDH using G.707 and G.783; for MPLS using Bidirectional Forwarding Detection (BFD); etc.) as for any other LSP. Hence, this requirement is out of the scope of GMPLS protocols.

The GMPLS protocols should provide mechanisms for the coordination of data link verification in the upper-layer network where data links are lower-layer LSPs.

- o GMPLS signaling allows an LSP to be put into 'test' mode [RFC3473].
- o The Link Management Protocol [RFC4204] is a targeted protocol and can be run end-to-end across lower-layer LSPs.
- o Coordination of testing procedures in different layers is an operational matter.

3.1.4. Scalability

As discussed in [RFC5212]), MRN/MLN routing mechanisms must be designed to scale well with an increase of any of the following:

- Number of nodes
- Number of TE links (including FA-LSPs)
- Number of LSPs
- Number of regions and layers
- Number of Interface Switching Capability Descriptors (ISCDs) per TE link.

GMPLS routing provides the necessary advertisement functions and is based on IETF-designed IGPs. These are known to scale relatively well with the number of nodes and links. Where there are multiple regions or layers, there are two possibilities.

1. If a single routing instance distributes information about multiple network layers, the effect is no more than to increase the number of nodes and links in the network.
2. If the MLN is fully integrated (i.e., constructed from hybrid nodes), there is an increase in the number of nodes and links (as just mentioned), and also a potential increase in the amount of ISCD information advertised per link. This is a relatively small amount of information (e.g., 36 bytes in OSPF [RFC4203]) per switching type, and each interface is unlikely to have more than two or three switching types.

The number of LSPs in a lower layer that are advertised as TE links may impact the scaling of the routing protocol. A full mesh of FA-LSPs in the lower layer would lead to n^2 TE links, where n is the number of layer-border LSRs. This must be taken into consideration in the VNT management process. This is an operational matter beyond the scope of GMPLS protocols.

Since it requires the maintenance of n^2 RSVP-TE sessions (which may be quite CPU- and memory-consuming), a full mesh of LSPs in the lower layer may impact the scalability of GMPLS signaling. The use of virtual TE links may reduce the control plane overload (see Section 3.1.1.2).

3.1.5. Operations and Management of the MLN/MRN

[RFC5212] identifies various requirements for effective management and operation of the MLN. Some features already exist within the GMPLS protocol set, some more are under development, and some requirements are not currently addressed and will need new development work in order to support them.

3.1.5.1. MIB Modules

MIB modules have been developed to model and control GMPLS switches [RFC4803] and to control and report on the operation of the signaling protocol [RFC4802]. These may be successfully used to manage the operation of a single instance of the control plane protocols that operate across multiple layers.

[RFC4220] provides a MIB module for managing TE links, and this may be particularly useful in the context of the MLN because LSPs in the lower layers are made available as TE links in the higher layer.

The traffic engineering database provides a repository for all information about the existence and current status of TE links within a network. This information is typically flooded by the routing protocol operating within the network, and is used when LSP routes are computed. [TED-MIB] provides a way to inspect the TED to view the TE links at the different layers of the MLN.

As observed in [RFC5212], although it would be possible to manage the MLN using only the existing MIB modules, a further MIB module could be produced to coordinate the management of separate network layers in order to construct a single MLN entity. Such a MIB module would effectively link together entries in the MIB modules already referenced.

3.1.5.2. OAM

At the time of writing, the development of OAM tools for GMPLS networks is at an early stage. GMPLS OAM requirements are addressed in [GMPLS-OAM].

In general, the lower layer network technologies contain their own technology-specific OAM processes (for example, SDH/SONET, Ethernet, and MPLS). In these cases, it is not necessary to develop additional OAM processes, but GMPLS procedures may be desirable to coordinate the operation and configuration of these OAM processes.

[ETH-OAM] describes some early ideas for this function, but more work is required to generalize the technique to be applicable to all technologies and to MLN. In particular, an OAM function operating within a server layer must be controllable from the client layer, and client layer control plane mechanisms must map and enable OAM in the server layer.

Where a GMPLS-controlled technology does not contain its own OAM procedures, this is usually because the technology cannot support in-band OAM (for example, Wavelength Division Multiplexing (WDM) networks). In these cases, there is very little that a control plane can add to the OAM function since the presence of a control plane cannot make any difference to the physical characteristics of the data plane. However, the existing GMPLS protocol suite does provide a set of tools that can help to verify the data plane through the control plane. These tools are equally applicable to network technologies that do contain their own OAM.

- Route recording is available through the GMPLS signaling protocol [RFC3473], making it possible to check the route reported by the control plane against the expected route. This mechanism also includes the ability to record and report the interfaces and labels used for the LSP at each hop of its path.
- The status of TE links is flooded by the GMPLS routing protocols [RFC4203] and [RFC4205] making it possible to detect changes in the available resources in the network as an LSP is set up.
- The GMPLS signaling protocol [RFC3473] provides a technique to place an LSP into a "test" mode so that end-to-end characteristics (such as power levels) may be sampled and modified.
- The Link Management Protocol [RFC4204] provides a mechanism for fault isolation on an LSP.
- GMPLS signaling [RFC3473] provides a Notify message that can be used to report faults and issues across the network. The message includes scaling features to allow one message to report the failure of multiple LSPs.
- Extensions to GMPLS signaling [RFC4783] enable alarm information to be collected and distributed along the path of an LSP for more easy coordination and correlation.

3.2. Specific Aspects of Multi-Region Networks

3.2.1. Support for Multi-Region Signaling

There are actually several cases where a transit node could choose between multiple Switching Capabilities (SCs) to be used for a lower-region FA-LSP:

- Explicit Route Object (ERO) expansion with loose hops: The transit node has to expand the path, and may have to select among a set of lower-region SCs.
- Multi-SC TE link: When the ERO of an FA LSP, included in the ERO of an upper-region LSP, comprises a multi-SC TE link, the region border node has to select among these SCs.

Existing GMPLS signaling procedures do not allow solving this ambiguous choice of the SC that may be used along a given path.

Hence, an extension to GMPLS signaling has to be defined to indicate the SC(s) that can be used and the SC(s) that cannot be used along the path.

3.2.2. Advertisement of Adjustment Capacities

In the MRN context, nodes supporting more than one switching capability on at least one interface are called hybrid nodes [RFC5212]. Conceptually, hybrid nodes can be viewed as containing at least two distinct switching elements interconnected by internal links that provide adjustment between the supported switching capabilities. These internal links have finite capacities and must be taken into account when computing the path of a multi-region TE-LSP. The advertisement of the adjustment capacities is required, as it provides critical information when performing multi-region path computation.

The term "adjustment capacity" refers to the property of a hybrid node to interconnect different switching capabilities it provides through its external interfaces [RFC5212]. This information allows path computation to select an end-to-end multi-region path that includes links of different switching capabilities that are joined by LSRs that can adapt the signal between the links.

Figure 1a below shows an example of a hybrid node. The hybrid node has two switching elements (matrices), which support TDM and PSC switching, respectively. The node has two PSC and TDM ports (Port1 and Port2, respectively). It also has an internal link connecting the two switching elements.

The two switching elements are internally interconnected in such a way that it is possible to terminate some of the resources of the TDM Port2; also, they can provide adjustment of PSC traffic that is received/sent over the internal PSC interface (#b). Two ways are possible to set up PSC LSPs (Port1 or Port2). Available resources advertisement (e.g., Unreserved and Min/Max LSP Bandwidth) should cover both ways.

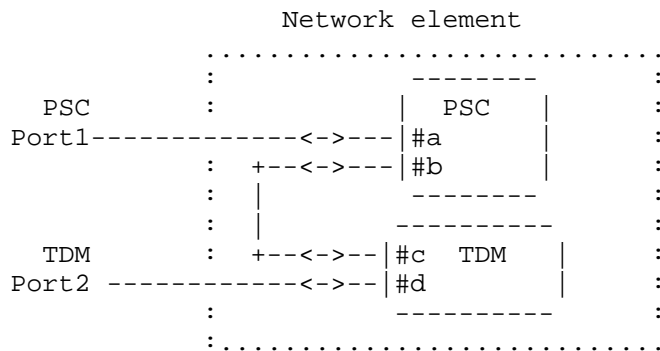


Figure 1a. Hybrid node.

Port1 and Port2 can be grouped together thanks to internal Dense Wavelength Division Multiplexing (DWDM), to result in a single interface: Link1. This is illustrated in Figure 1b below.




Figure 1b. Hybrid node.

Let's assume that all interfaces are STM16 (with VC4-16c capable as Max LSP bandwidth). After setting up several PSC LSPs via port #a and setting up and terminating several TDM LSPs via port #d and port #b, a capacity of only 155 Mb is still available on port #b. However, a 622 Mb capacity remains on port #a, and VC4-5c capacity remains on port #d.

When computing the path for a new VC4-4c TDM LSP, one must know that this node cannot terminate this LSP, as there is only a 155 Mb capacity still available for TDM-PSC adjustment. Hence, the TDM-PSC adjustment capacity must be advertised.

With current GMPLS routing [RFC4202], this advertisement is possible if link bundling is not used and if two TE links are advertised for Link1.

We would have the following TE link advertisements:

- TE link 1 (Port1):
 - ISCD sub-TLV: PSC with Max LSP bandwidth = 622 Mb
 - Unreserved bandwidth = 622 Mb.

TE link 2 (Port2):

- ISCD #1 sub-TLV: TDM with Max LSP bandwidth = VC4-4c,
- ISCD #2 sub-TLV: PSC with Max LSP bandwidth = 155 Mb,
- Unreserved bandwidth (equivalent): 777 Mb.

The ISCD #2 in TE link 2 actually represents the TDM-PSC adjustment capacity.

However, if for obvious scalability reasons, link bundling is done, then the adjustment capacity information is lost with current GMPLS routing, as we have the following TE link advertisement:

TE link 1 (Port1 + Port2):

- ISCD #1 sub-TLV: TDM with Max LSP bandwidth = VC4-4c,
- ISCD #2 sub-TLV: PSC with Max LSP bandwidth = 622 Mb,
- Unreserved bandwidth (equivalent): 1399 Mb.

With such a TE link advertisement, an element computing the path of a VC4-4c LSP cannot know that this LSP cannot be terminated on the node.

Thus, current GMPLS routing can support the advertisement of the adjustment capacities, but this precludes performing link bundling and thus faces significant scalability limitations.

Hence, GMPLS routing must be extended to meet this requirement. This could rely on the advertisement of the adjustment capacities as a new TE link attribute (that would complement the Interface Switching Capability Descriptor TE link attribute).

Note: Multiple ISCDs MAY be associated with a single switching capability. This can be performed to provide (e.g., for TDM interfaces) the Min/Max LSP Bandwidth associated to each layer (or set of layers) for that switching capability. For example, an interface associated to TDM switching capability and supporting VC-12 and VC-4 switching can be associated to one ISCD sub-TLV or two ISCD sub-TLVs. In the first case, the Min LSP Bandwidth is set to VC-12 and the Max LSP Bandwidth to VC-4. In the second case, the Min LSP Bandwidth is set to VC-12 and the Max LSP Bandwidth to VC-12, in the first ISCD sub-TLV; and the Min LSP Bandwidth is set to VC-4 and the Max LSP Bandwidth to VC-4, in the second ISCD sub-TLV. Hence, in the first case, as long as the Min LSP Bandwidth is set to VC-12 (and not VC-4), and in the second case, as long as the first ISCD sub-TLV is advertised, there is sufficient capacity across that interface to setup a VC-12 LSP.

4. Evaluation Conclusion

Most of the required MLN/MRN functions will rely on mechanisms and procedures that are out of the scope of the GMPLS protocols, and thus do not require any GMPLS protocol extensions. They will rely on local procedures and policies, and on specific TE mechanisms and algorithms.

As regards Virtual Network Topology (VNT) computation and reconfiguration, specific TE mechanisms need to be defined, but these mechanisms are out of the scope of GMPLS protocols.

Six areas for extensions of GMPLS protocols and procedures have been identified:

- GMPLS signaling extension for the setup/deletion of the virtual TE links;
- GMPLS signaling extension for graceful TE link deletion;
- GMPLS signaling extension for constrained multi-region signaling (SC inclusion/exclusion);
- GMPLS routing extension for the advertisement of the adjustment capacities of hybrid nodes.
- A MIB module for coordination of other MIB modules being operated in separate layers.
- GMPLS signaling extensions for the control and configuration of technology-specific OAM processes.

4.1. Traceability of Requirements

This section provides a brief cross-reference to the requirements set out in [RFC5212] so that it is possible to verify that all of the requirements listed in that document have been examined in this document.

- Path computation mechanism should be able to compute paths and handle topologies consisting of any combination of (simplex) nodes ([RFC5212], Section 5.1).
 - o Path computation mechanisms are beyond the scope of protocol specifications, and out of scope for this document.

- A hybrid node should maintain resources on its internal links ([RFC5212], Section 5.2).
 - o This is an implementation requirement and is beyond the scope of protocol specifications, and it is out of scope for this document.
- Path computation mechanisms should be prepared to use the availability of termination/adjustment resources as a constraint in path computation ([RFC5212], Section 5.2).
 - o Path computation mechanisms are beyond the scope of protocol specifications, and out of scope for this document.
- The advertisement of a node's ability to terminate lower-region LSPs and to forward traffic in the upper-region (adjustment capability) is required ([RFC5212], Section 5.2).
 - o See Section 3.2.2 of this document.
- The path computation mechanism should support the coexistence of upper-layer links directly connected to upper-layer switching elements, and upper-layer links connected through internal links between upper-layer and lower-layer switching elements ([RFC5212], Section 5.2).
 - o Path computation mechanisms are beyond the scope of protocol specifications, and out of scope for this document.
- MRN/MLN routing mechanisms must be designed to scale well with an increase of any of the following:
 - Number of nodes
 - Number of TE links (including FA-LSPs)
 - Number of LSPs
 - Number of regions and layers
 - Number of ISCDs per TE link.([RFC5212], Section 5.3).
 - o See Section 3.1.4 of this document.
- Design of the routing protocols must not prevent TE information filtering based on ISCDs ([RFC5212], Section 5.3).
 - o All advertised information carries the ISCD, and so a receiving node may filter as required.
- The path computation mechanism and the signaling protocol should be able to operate on partial TE information, ([RFC5212], Section 5.3).
 - o Path computation mechanisms are beyond the scope of protocol specifications, and out of scope for this document.

- Protocol mechanisms must be provided to enable creation, deletion, and modification of LSPs triggered through operational actions ([RFC5212], Section 5.4).
 - o Such mechanisms are standard in GMPLS signaling [RFC3473].
- Protocol mechanisms should be provided to enable similar functions triggered by adjacent layers ([RFC5212], Section 5.4).
 - o Such mechanisms are standard in GMPLS signaling [RFC3473].
- Protocol mechanisms may be provided to enable adaptation to changes such as traffic demand, topology, and network failures. Routing robustness should be traded with adaptability of those changes ([RFC5212], Section 5.4).
 - o See Section 3.1.1 of this document.
- Reconfiguration of the VNT must be as non-disruptive as possible and must be under the control of policy configured by the operator ([RFC5212], Section 5.5).
 - o See Section 3.1.1.3 of this document
- Parameters of a TE link in an upper layer should be inherited from the parameters of the lower-layer LSP that provides the TE link, based on polices configured by the operator ([RFC5212], Section 5.6).
 - o See Section 3.1.2 of this document.
- The upper-layer signaling request may contain an ERO that includes only hops in the upper layer ([RFC5212], Section 5.7).
 - o Standard for GMPLS signaling [RFC3473]. See also Section 3.2.1.
- The upper-layer signaling request may contain an ERO specifying the lower layer FA-LSP route ([RFC5212], Section 5.7).
 - o Standard for GMPLS signaling [RFC3473]. See also Section 3.2.1.
- As part of the re-optimization of the MLN, it must be possible to reroute a lower-layer FA-LSP while keeping interface identifiers of the corresponding TE links unchanged and causing only minimal disruption to higher-layer traffic ([RFC5212], Section 5.8.1).
 - o See Section 3.1.1.3.
- The solution must include measures to protect against network destabilization caused by the rapid setup and tear-down of lower-layer LSPs, as traffic demand varies near a threshold ([RFC5212], Sections 5.8.1 and 5.8.2).
 - o See Section 3.1.1.4.

- Signaling of lower-layer LSPs should include a mechanism to rapidly advertise the LSP as a TE link in the upper layer, and to coordinate into which routing instances the TE link should be advertised ([RFC5212], Section 5.8.1).
 - o This is provided by [RFC4206] and enhanced by [HIER-BIS]. See also Section 3.1.1.2.
- If an upper-layer LSP is set up making use of a virtual TE link, the underlying LSP must immediately be signaled in the lower layer ([RFC5212], Section 5.8.2).
 - o See Section 3.1.1.2.
- The solution should provide operations to facilitate the build-up of virtual TE links, taking into account the forecast upper-layer traffic demand, and available resource in the lower layer ([RFC5212], Section 5.8.2).
 - o See Section 3.1.1.2 of this document.
- The GMPLS protocols should provide mechanisms for the coordination of data link verification in the upper-layer network where data links are lower layer LSPs ([RFC5212], Section 5.9).
 - o See Section 3.1.3 of this document.
- Multi-layer protocol solutions should be manageable through MIB modules ([RFC5212], Section 5.10).
 - o See Section 3.1.5.1.
- Choices about how to coordinate errors and alarms, and how to operate OAM across administrative and layer boundaries must be left open for the operator ([RFC5212], Section 5.10).
 - o This is an implementation matter, subject to operational policies.
- It must be possible to enable end-to-end OAM on an upper-layer LSP. This function appears to the ingress LSP as normal LSP-based OAM [GMPLS-OAM], but at layer boundaries, depending on the technique used to span the lower layers, client-layer OAM operations may need to be mapped to server-layer OAM operations ([RFC5212], Section 5.10).
 - o See Section 3.1.5.2.
- Client-layer control plane mechanisms must map and enable OAM in the server layer ([RFC5212], Section 5.10).
 - o See Section 3.1.5.2.
- OAM operation enabled for an LSP in a client layer must operate for that LSP along its entire length ([RFC5212], Section 5.10).
 - o See Section 3.1.5.2.

- OAM function operating within a server layer must be controllable from the client layer. Such control should be subject to policy at the layer boundary ([RFC5212], Section 5.10).
 - o This is an implementation matter.
- The status of a server layer LSP must be available to the client layer. This information should be configurable to be automatically notified to the client layer at the layer boundary, and should be subject to policy ([RFC5212], Section 5.10).
 - o This is an implementation matter.
- Implementations may use standardized techniques (such as MIB modules) to convey status information between layers.
 - o This is an implementation matter.

5. Security Considerations

[RFC5212] sets out the security requirements for operating a MLN or MRN. These requirements are, in general, no different from the security requirements for operating any GMPLS network. As such, the GMPLS protocols already provide adequate security features. An evaluation of the security features for GMPLS networks may be found in [MPLS-SEC], and where issues or further work is identified by that document, new security features or procedures for the GMPLS protocols will need to be developed.

[RFC5212] also identifies that where the separate layers of a MLN/MRN are operated as different administrative domains, additional security considerations may be given to the mechanisms for allowing inter-layer LSP setup. However, this document is explicitly limited to the case where all layers under GMPLS control are part of the same administrative domain.

Lastly, as noted in [RFC5212], it is expected that solution documents will include a full analysis of the security issues that any protocol extensions introduce.

6. Acknowledgments

We would like to thank Julien Meuric, Igor Bryskin, and Adrian Farrel for their useful comments.

Thanks also to Question 14 of Study Group 15 of the ITU-T for their thoughtful review.

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.
- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4202] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.
- [RFC5212] Shiimoto, K., Papadimitriou, D., Le Roux, JL., Vigoureux, M., and D. Brungard, "Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", RFC 5212, July 2008.

7.2. Informative References

- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC4203] Kompella, K., Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.
- [RFC4204] Lang, J., Ed., "Link Management Protocol (LMP)", RFC 4204, October 2005.
- [RFC4205] Kompella, K., Ed., and Y. Rekhter, Ed., "Intermediate System to Intermediate System (IS-IS) Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4205, 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.

- [RFC4220] Dubuc, M., Nadeau, T., and J. Lang, "Traffic Engineering Link Management Information Base", RFC 4220, November 2005.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC4783] Berger, L., Ed., "GMPLS - Communication of Alarm Information", RFC 4783, December 2006.
- [RFC4802] Nadeau, T., Ed., and A. Farrel, Ed., "Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering Management Information Base", RFC 4802, February 2007.
- [RFC4803] Nadeau, T., Ed., and A. Farrel, Ed., "Generalized Multiprotocol Label Switching (GMPLS) Label Switching Router (LSR) Management Information Base", RFC 4803, February 2007.
- [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.
- [RFC4974] Papadimitriou, D. and A. Farrel, "Generalized MPLS (GMPLS) RSVP-TE Signaling Extensions in Support of Calls", RFC 4974, August 2007.
- [ETH-OAM] Takacs, A., Gero, B., and D. Mohan, "GMPLS RSVP-TE Extensions to Control Ethernet OAM", Work in Progress, July 2008.
- [GMPLS-OAM] Nadeau, T., Otani, T. Brungard, D., and A. Farrel, "OAM Requirements for Generalized Multi-Protocol Label Switching (GMPLS) Networks", Work in Progress, October 2007.
- [GR-SHUT] Ali, Z., Zamfir, A., and J. Newton, "Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks", Work in Progress, July 2008.
- [HIER-BIS] Shiimoto, K., Rabbat, R., Ayyangar, A., Farrel, A., and Z. Ali, "Procedures for Dynamically Signaled Hierarchical Label Switched Paths", Work in Progress, February 2008.
- [MPLS-SEC] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", Work in Progress, July 2008.

- [PCE-INTER] Oki, E., Le Roux, J-L., and A. Farrel, "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", Work in Progress, June 2008.
- [TED-MIB] Miyazawa, M., Otani, T., Nadeau, T., and K. Kunaki, "Traffic Engineering Database Management Information Base in support of MPLS-TE/GMPLS", Work in Progress, July 2008.

8. Contributors' Addresses

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

Eiji Oki
NTT
3-9-11 Midori-Cho
Musashino, Tokyo 180-8585, Japan
EMail: oki.eiji@lab.ntt.co.jp

Kohei Shiomoto
NTT
3-9-11 Midori-Cho
Musashino, Tokyo 180-8585, Japan
EMail: shiomoto.kohei@lab.ntt.co.jp

M. Vigoureux
Alcatel-Lucent France
Route de Villejust
91620 Nozay
FRANCE
EMail: martin.vigoureux@alcatel-lucent.fr

Editors' Addresses

Jean-Louis Le Roux
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex, France
EMail: jeanlouis.leroux@orange-ftgroup.com

Dimitri Papadimitriou
Alcatel-Lucent
Francis Wellensplein 1,
B-2018 Antwerpen, Belgium
EMail: dimitri.papadimitriou@alcatel-lucent.be

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.