Network Working Group   D. Papadimitriou (Alcatel)
Internet Draft   M. Vigoureux (Alcatel)
Expiration Date: December 2007   K. Shiomoto (NTT)
                                                             D. Brungard (ATT)
                                           J.L. Le Roux (France Telecom)

July 2007

Generalized Multi-Protocol Label Switching (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)

draft-papadimitriou-ccamp-gmpls-mrn-extensions-04.txt

Abstract

There are requirements for the support of networks ccomprising LSRs with different data plane switching layers controlled by a single Generalized Multi Protocol Label Switching (GMPLS) control plane

instance, referred to as GMPLS Multi-Layer Networks/Multi-Region Networks (MLN/MRN). This document defines extensions to GMPLS routing and signaling protocols so as to support the operation of GMPLS Multi-Layer/Multi-Region Networks.

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 RFC-2119.

In addition the reader is assumed to be familiar with the concepts developed in [RFC3945], [RFC3471], and [RFC4202] as well as [RFC4206] and [RFC4201].

1. Introduction

Generalized Multi-Protocol Label Switching (GMPLS) [RFC 3945] extends MPLS to handle multiple switching technologies: packet switching (PSC), layer-two switching (L2SC), TDM switching (TDM), wavelength switching (LSC) and fiber switching (FSC). A GMPLS switching type (PSC, TDM, etc.) describes the ability of a node to forward data of a particular data plane technology, and uniquely identifies a control plane region. LSP Regions are defined in [RFC4206]. A network comprised of multiple switching types (e.g. PSC and TDM) controlled by a single GMPLS control plane instance is called a Multi-Region Network (MRN).

A data plane layer is a collection of network resources capable of terminating and/or switching data traffic of a particular format. For example, LSC, TDM VC-11 and TDM VC-4-64c represent three different layers. A network comprising transport nodes with different data plane switching layers controlled by a single GMPLS control plane instance is called a Multi-Layer Network (MLN).

The applicability of GMPLS to multiple switching technologies provides the unified control and operations for both LSP provisioning and recovery. This document covers the elements of a single GMPLS control plane instance controlling multiple layers within a given TE domain. A CP instance can serve one, two or more layers. Other possible approaches such as having multiple CP instances serving disjoint sets of layers are outside the scope of this document.

The next sections provide the procedural aspects in terms of routing and signaling for such environments as well as the extensions required to instrument GMPLS to provide the capabilities for MLM/MRN unified control. The rationales and requirements for Multi-Layer/Region networks are set forth in [MLN-REQ]. These requirements are evaluated against GMPLS protocols in [MLN-EVAL] and several

areas where GMPLS protocol extensions are required are identified. This document defines GMPLS routing and signaling extensions so as to cover GMPLS MLN/MRN requirements.

2. Summary of the Requirements and Evaluation

As identified in [MRN-EVAL] most of MLN/MRN requirements rely on mechanisms and procedures that are outside the scope of the GMPLS protocols, and thus do not require any GMPLS protocol extensions. They rely on local procedures and policies, and on specific TE mechanisms and algorithms, which are outside the scope of GMPLS protocols.

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

- GMPLS routing extension for the advertisement of the internal adaptation capability of hybrid nodes.

- GMPLS signaling extension for constrained multi-region signaling (SC inclusion/exclusion).

- GMPLS signaling extension for the setup/deletion of Virtual TE-links (as well as exact trigger for its actual provisioning).

- GMPLS routing and signaling extension for graceful TE-link deletion (covered in [GR-TELINK]).

The first three requirements are addressed in Sections 3, 4 and 5, respectively, of this document. The fourth one is addressed in [GR-TELINK].

3. Interface adaptation capability descriptor (IACD)

In the MRN context, nodes supporting more than one switching capability on at least one interface are called Hybrid nodes. Hybrid nodes contain at least two distinct switching elements that are interconnected by internal links to provide adaptation 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 internal adaptation capability is required as it provides critical information when performing multi-region path computation.

3.1 Overview

In an MRN environment, some LSRs could contain, under the control of a single GMPLS instance, multiple switching capabilities such as PSC and TDM or PSC and Lambda Switching Capability (LSC).

These nodes, hosting multiple Interface Switching Capabilities (ISC), just like other nodes (hosting a single Interface Switching Capability) are required to hold and advertise resource information on link states and topology. They also may have to consider certain portions of internal node resources to terminate hierarchical label switched paths (LSPs), since circuit switch capable units such as TDMs, LSCs, and FSCs require rigid resources. For example, a node with PSC+LSC hierarchical switching capability can switch a Lambda LSP but may not be able to can never terminate the Lambda LSP if there is no unused adaptation capability between the LSC and the PSC switching capabilities.

Another example occurs when L2SC (Ethernet) switching can be adapted in LAPS X.86 and GFP for instance before reaching the TDM switching matrix. Similar circumstances can occur, if a switching fabric that supports both PSC and L2SC functionalities is assembled with LSC interfaces enabling "lambda" encoding. In the switching fabric, some interfaces can terminate Lambda LSPs and perform frame (or cell) switching whilst other interfaces can terminate Lambda LSPs and perform packet switching.

Therefore, within multi-region networks, the advertisement of the so-called adaptation capability to terminate LSPs (not the interface capability since the latter can be inferred from the bandwidth available for each switching capability) provides critical information to take into account when performing multi-region path computation. This concept enables a node to discriminate the remote nodes (and thus allows their selection during path computation) with respect to their adaptation capability e.g. to terminate LSPs at the PSC or LSC level.

Hence, we introduce the idea of discriminating the (internal) adaptation capability from the (interface) switching capability by considering an interface adaptation capability descriptor.

A more detailed problem statement can be found in [MLN-EVAL].

3.2 Interface Adaptation Capability Descriptor (IACD) Format

The interface switching capability descriptor (IACD) provides the information for the forwarding/switching) capability only.

3.2.1 OSPF-TE

In OSPF, the IACD sub-TLV is defined as a sub-TLV of the Link TLV (see [RFC 3630]), with type TBD. The IACD sub-TLV format is defined as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Switching Cap |   Encoding    | Switching Cap |   Encoding    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 0              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 1              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 2              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 3              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 4              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 5              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 6              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 7              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        Adaptation Capability-specific information             |
|                  (variable)                                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Where:

- first Switching Capability (SC) field (byte 1): lower switching capability (as defined for the existing ISC sub-TLV)
- first Encoding field (byte 2): as defined for the existing ISC sub-TLV
- second SC value (byte 3): upper switching capability (new)
- second encoding value (byte 4): set to the encoding of the available adaptation pool and to 0xFF when the corresponding SC value has no access to the wire (i.e. there is no ISC sub-TLV for this upper switching capability)

Multiple IACD sub-TLVs may be present within a given TE Link TLV and the bandwidth simply provides an indication of resources still available to perform insertion/ extraction for a given adaptation (pool concept).

3.2.2 ISIS-TE

In IS-IS, the IACD sub-TLV is a sub-TLV of the Extended IS Reachability TLV (see [RFC 3784]) with type TBD. The IACD sub-TLV format is defined as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Switching Cap |   Encoding    | Switching Cap |   Encoding    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 0              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 1              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 2              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 3              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 4              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 5              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 6              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  Max LSP Bandwidth at priority 7              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        Adaptation Capability-specific information             |
|                  (variable)                                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Where the fields have the same processing and interpretation rules as for Section 3.2.1

Multiple IACD sub-TLVs may be present within a given extended IS reachability TLV and the bandwidth simply provides an indication of resources still available to perform insertion/ extraction for a given adaptation (pool concept).

4. Multi-Region Signaling

Section 8.2 of [RFC4206] specifies that when a region boundary node receives a Path message, the node determines whether it is at the edge of an LSP region with respect to the ERO carried in the message. If the node is at the edge of a region, it must then determine the other edge of the region with respect to the ERO, using the IGP database. The node then extracts from the ERO the subsequence of hops from itself to the other end of the region.

The node then compares the subsequence of hops with all existing FA-LSPs originated by the node:

- if a match is found, that FA-LSP has enough unreserved bandwidth for the LSP being signaled, and the PID of the FA-LSP is compatible with the PID of the LSP being signaled, the node uses that FA-LSP as follows. The Path message for the original LSP is sent to the egress of the FA-LSP. The PHOP in the message is the address of the node at the head-end of the FA-LSP. Before sending the Path message, the ERO in that message is adjusted by removing the subsequence of the ERO that lies in the FA-LSP, and replacing it with just the end point of the FA-LSP.
- if no existing FA-LSP is found, the node sets up a new FA-LSP. That is, it initiates a new LSP setup just for the FA-LSP.

Applying this procedure, in a MRN environment MAY lead to setup one-hop FA-LSPs between each node. Therefore, considering that the path computation is able to take into account richness of information with regard to the SC available on given nodes belonging to the path, it is consistent to provide enough signaling information to indicate the SC to be used and on over which link. Particularly, in case a TE link has multiple SC advertised as part of its ISCD sub-TLVs, an ERO does not allow selecting a particular SC.

Limiting modifications to existing RSVP-TE procedures [RFC3473] and referenced, this document defines a new sub-object of the eXclude Route Object [XRO], called Switching Capability sub-object. This sub-object enables (when desired) the explicit identification of (at least one) switching capability to be excluded from the resource selection process described here above.

Including this sub-object as part of the XRO that explicitly indicates which SCs have to be excluded (before initiating the procedure described here above) over a specified TE link solves the ambiguous choice among SCs that are potentially used along a given path and give the possibility to optimize resource usage on a multi-region basis. Note that implicit SC inclusion is easily supported by explicitly excluding other SCs (e.g. to include LSC, it is required to exclude PSC, L2SC, TDM and FSC).

Note: usage of the EXRS is under investigation.

4.1 SC Subobject Encoding

The contents of an EXCLUDE_ROUTE object defined in [XRO] are a series of variable-length data items called subobjects. This document defines the SC subobject of the XRO (Type TBD), its encoding and processing.

Subobject Type TBD: Switching Capability

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|L|    Type     |     Length    |   Attribute   | Switching Cap |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

L

0 indicates that the attribute specified MUST be excluded
1 indicates that the attribute specified SHOULD be avoided

Attribute

0 reserved value

1 indicates that the specified SC should be excluded or avoided with respect to the preceding numbered (Type 1 or Type 2) or unnumbered interface (Type) subobject

Switching Cap (8-bits)

Switching Capability value to be excluded.

This sub-object must follow the set of numbered or unnumbered interface sub-objects to which this sub-object refers. In case, of loose hop ERO subobject, the XRO sub-object must precede the loose-hop sub-object identifying the tail-end node/interface of the traversed region(s).

Furthermore, it is expected, when label sub-object are following numbered or unnumbered interface sub-objects, that the label value is compliant with the SC capability to be explicitly excluded.

5. Virtual TE link

Two techniques can be used for the setup operation and maintenance of Virtual TE links. The corresponding GMPLS protocols extensions are described in this section.

5.1 Edge-to-edge Association

This approach that does not require state maintenance on transit LSRs relies on extensions to the GMPLS RSVP-TE Call procedure ([GMPLS-CALL]).

This technique consists of exchanging identification and TE attributes information directly between TE link end points. These TE link end-points correspond to the LSP head-end and tail-end points of of the LSPs that will be established. The end-points MUST belong to

the same (LSP) region through the establishment of a call between terminating LSRs.

Once the call is established the resulting association populates the local TEDB and the resulting TE link is advertised as any other TE link. The latter can then be used to attract traffic. Once an upper layer/lower region LSP makes use of this TE link. A set of one or more LSPs must be initially established before the FA LSP can be used for nesting the incoming LSP.

In order to distinguish usage of such call from a classical call (as defined e.g. in [RFC4139]), a CALL ATTRIBUTES object is introduced.

5.1.1 CALL_ATTRIBUTES Object

The CALL_ATTRIBUTES object is used to signal attributes required in support of a call, or to indicate the nature or use of a call. It is built on the LSP-ATTRIBUTES object defined in [RFC4420].

The CALL_ATTRIBUTES object class is 201 (TBD by IANA) of the form 11bbbbbb. This C-Num value (see [RFC2205], Section 3.10) ensures that LSRs that do not recognize the object pass it on transparently.

One C-Type is defined, C-Type = 1 for CALL Attributes. This object is optional and may be placed on Notify messages to convey additional information about the desired attributes of the call.

5.1.2 Processing

Specifically, if an egress (or intermediate) LSR does not support the object, it forwards it unexamined and unchanged. This facilitates the exchange of attributes across legacy networks that do not support this new object.

The CALL_ATTRIBUTES object may be used to report call operational state on a Notify message.

```
   CALL_ATTRIBUTES class = 201, C-Type = 1

    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                                                               |
   //                       Attributes TLVs                       //
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The Attributes TLVs are encoded as described in Section 3.

5.1.3 Attributes TLVs

Attributes carried by the CALL_ATTRIBUTES object are encoded within TLVs. One or more TLVs may be present in each object.

There are no ordering rules for TLVs, and no interpretation should be placed on the order in which TLVs are received.

Each TLV is encoded as follows.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Type              |           Length              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
//                            Value                            //
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Type

The identifier of the TLV.

Length

The length of the Value field in bytes. Thus, if no Value field is present the Length field contains the value zero. Each Value field must be zero padded at the end to take it up to a four byte boundary -- the padding is not included in the length so that a one byte value would be encoded in an eight byte TLV with Length field set to one.

Value

The data for the TLV padded as described above.

TLV Type 1 indicates the Attributes Flags TLV. Other TLV types may be defined in the future with type values assigned by IANA (see Section 11.2). The Attributes Flags TLV may be present in a CALL_ATTRIBUTES object.

The Attribute Flags TLV value field is an array of units of 32 flags numbered from the most significant bit as bit zero. The Length field for this TLV is therefore always a multiple of 4 bytes, regardless of the number of bits carried and no padding is required.

Unassigned bits are considered as reserved and MUST be set to zero on transmission by the originator of the object. Bits not contained in

the TLV MUST be assumed to be set to zero. If the TLV is absent either because it is not contained in the CALL_ATTRIBUTES object or because this object is itself absent, all processing MUST be performed as though the bits were present and set to zero. That is to say, assigned bits that are not present either because the TLV is deliberately foreshortened or because the TLV is not included MUST be treated as though they are present and are set to zero.

5.1.4 Call inheritance Flag

This document introduces a specific flag (MSB position bit 0) of the Attributes Flags TLV, to indicate that the association initiated between the end-points belonging to a call results into a (virtual) TE link advertisement.

The Call inheritance flag MUST be set to 1 in order to indicate that the established association is to be translated into a TE link advertisement. The value of this flag is by default set to 1. Setting this flag to 0 results in a hidden TE link or in deleting the corresponding TE link advertisement (by setting the corresponding Opaque LSA Age to MaxAge).

The notify message used for establishing the association is defined as per [GMPLS-CALL]. Additionally, the notify message must carry an LSP_TUNNEL_INTERFACE_ID Object, that allows identifying unnumbered FA-LSPs ([RFC3477], [RFC4206], [RFC4206-bis]) and numbered FA-LSPs ([RFC4206], [RFC4206-bis]).

5.2. Soft-FA approach

The Soft Forwarding Adjacency (Soft FA) approach consists of setting up the FA LSP at the control plane level without actually committing resources in the data plane. This means that the corresponding LSP exists only in the control plane domain.

Once such FA is established the corresponding TE link can be advertized following the procedures described in [RFC 4206].

5.2.1 Pre-planned LSP Flag

The LSP ATTRIBUTES object and Attributes Flags TLV are defined in [RFC4420]. The present document defines a new flag, the pre-planned LSP Flag, in the existing Attributes Flags TLV (numbered as Type 1).

The position of this flag is TBD in accordance with IANA assignment. This flag, part of the LSP_REQUIRED ATTRIBUTE object, follows processing of [RFC4420] for that object. That is, LSRs that do not recognize the object reject the LSP setup effectively saying that

they do not support the attributes requested. Indeed, the newly defined attribute requires examination at all transit LSRs.

The pre-planned LSP Flag can take one of the following values:

o) When set to 0 this means that the LSP should be fully provisioned. Absence of this flag (hence corresponding TLV) is therefore compliant with the signaling message processing per [RFC3473])

o) When set to 1 this means that the LSP should be provisioned in the control plane only.

If an LSP is established with the pre-planned Flag set to 1, no resources are committed at the data plane level. The operation of committing data plane resources occurs by re-signaling the same LSP with the pre-planned Flag set to 0. It is RECOMMENDED that no other modifications are made to other RSVP objects during this operation. That is each intermediate node, processing a Flag transiting from 1 to 0 shall only be concerned with the commitment of data plane resources and no other modification of the LSP properties and/or attributes.

If an LSP is established with the pre-planned Flag set to 0, it MAY be re-signaled by setting the Flag to 1.

5.2.2 Path Provisioned LSPs

There is a difference in between an LSP that is established with 0 bandwidth (path provisioning) and an LSP that is established with a certain bandwidth value not committed at the data plane level (i.e. pre-planned LSP).

However, the former is currently not possible using the GMPLS protocol suite (following technology specific SENDER_TSPEC/FLOWSPEC definition). Indeed, Traffic Parameters such as those defined in [RFC 4606] do not support setup of 0 bandwidth LSPs.

Mechanisms for provisioning (pre-planned or not) LSP with 0 bandwidth will be described in next release of this document.

6. Backward compatibility

New objects and procedures defined in this document are running within a given TE domain. The latter is expected to run in the context of a consistent TE policy.

In such TE domains, we distinguish between edge LSRs and intermediate LSrs. Edge LSRs must be able to process Call Attribute as defined in section 5.1 if this is method selected or creating edge-to-edge

associations. In that domain, intermediate LSRs are by definition transparent to the Call processing.

In case the Soft FA method is used for the creation of Virtual TE links, edge and intermediate LSRs must support processing of the LSP ATTRIBUTE object per Section 5.2.

7. Security Considerations

In its current version, this memo does not introduce new security consideration from the ones already detailed in the GMPLS protocol suite.

8. References

8.1 Normative References

[GMPLS-CALL]D.Papadimitriou and A.Farrel, "Generalized MPLS (GMPLS) RSVP-TE Signaling Extensions in support of Calls," Work in progress, draft-ietf-ccamp-gmpls-rsvp-te-call-04.txt, January 2007.

[L2SC-LSP] D.Papadimitriou, et al., "Generalized MPLS Signaling for Layer-2 Label Switched Paths (LSP)", Work in Progress, draft-papadimitriou-ccamp-gmpls-l2sc-lsp-03.txt.

[MRN-EVAL] J.-L. Leroux et al., "Evaluation of existing GMPLS Protocols against Multi Region and Multi Layer Networks (MRN/MLN)", Work in Progress, draft-ietf-ccamp-gmpls-mln-eval-03.txt.

[MRN-REQ] K.Shiomoto et al., "Requirements for GMPLS-based multi-region and multi-layer networks (MRN/MLN)", Work in Progress, draft-ietf-ccamp-gmpls-mrn-reqs-03.txt.

[RFC2119] Bradner, S., 'Key words for use in RFCs to indicate requirements levels', RFC 2119, March 1997.

[RFC2370] R.Coltun, "The OSPF Opaque LSA Option", RFC 2370, July 1998.

[RFC3471] L.Berger et al., "Generalized Multi-Protocol Label Switching (GMPLS) - Signaling Functional Description", RFC 3471, January 2003.

[RFC3630] D.Katz et al., "Traffic Engineering (TE) Extensions to OSPF Version 2," RFC 3630, September 2003.

[RFC3667] Bradner, S., "IETF Rights in Contributions", BCP 78,

RFC 3667, February 2004.

[RFC3668] Bradner, S., Ed., 'Intellectual Property Rights in IETF Technology', BCP 79, RFC 3668, February 2004.

[RFC4201] K.Kompella, et al., "Link Bundling in MPLS Traffic Engineering", RFC 4201, October 2005.

[RFC4202] K.Kompella (Editor), Y. Rekhter (Editor) et al. "Routing Extensions in Support of Generalized MPLS", RFC 4202, October 2005.

[RFC4206] K.Kompella and Y.Rekhter, "LSP Hierarchy with Generalized MPLS TE", RFC 4206, October 2005.

[RFC4206-bis] Shimoto et al. " Procedures for Dynamically Signaled Hierarchical Label Switched Paths ", draft-ietf-ccamp-lsp-hierarchy-bis, work in progress.

[RFC4420] A.Farrel et al., "Encoding of Attributes for Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Establishment Using Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4420, February 2006.

[RFC4428] D.Papadimitriou et al. "Analysis of Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery Mechanisms (including Protection and Restoration)", RFC 4428, March 2006.

[XRO] C.Y.Lee et al. "Exclude Routes - Extension to RSVP-TE," Work in progress, draft-ietf-ccamp-rsvp-te-exclude-route-05.txt, August 2005.

8.2 Informative References

[MAMLTE] K.Shiomoto et al., "Multi-area multi-layer traffic engineering using hierarchical LSPs in GMPLS networks", Work in Progress, draft-shiomoto-multiarea-te-01.txt.

[MLRT] W.Imajuku et al., "Multilayer routing using multilayer switch capable LSRs, Work in Progress, draft-imajuku-ml-routing-02.txt.

Acknowledgments

The authors would like to thank Mr. Wataru Imajuku for the discussions on adaptation between regions [MLRT].

Author's Addresses

Dimitri Papadimitriou
Alcatel-Lucent Bell
Copernicuslaan 50
B-2018 Antwerpen, Belgium
Phone : +32 3 240 8491
E-mail: dimitri.papadimitriou@alcatel-lucent.be

Martin Vigoureux
Alcatel-Lucent France
Route de Villejust
91620 Nozay, France
Tel : +33 1 30 77 26 69
Email: martin.vigoureux@alcatel-lucent.fr

Kohei Shiomoto
NTT Network Service Systems Laboratories
3-9-11 Midori-cho
Musashino-shi, Tokyo 180-8585, Japan
Phone: +81 422 59 4402
E-mail: shiomoto.kohei@lab.ntt.co.jp

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

Jean-Louis Le Roux
FTRD/DAC/LAN
Avenue Pierre Marzin
22300 Lannion, France
Phone: +33 (0)2 96 05 30 20
E-mail:jean-louis.leroux@rd.francetelecom.com

Contributors

Eiji Oki
NTT Network Service Systems Laboratories
3-9-11 Midori-cho
Musashino-shi, Tokyo 180-8585, Japan
Phone : +81 422 59 3441
Email: oki.eiji@lab.ntt.co.jp

Ichiro Inoue
NTT Network Service Systems Laboratories
3-9-11 Midori-cho
Musashino-shi, Tokyo 180-8585, Japan

Phone : +81 422 59 6076
Email: ichiro.inoue@lab.ntt.co.jp

Emmanuel Dotaro
Alcatel-Lucent France
Route de Villejust
91620 Nozay, France
Phone : +33 1 6963 4723
Email: emmanuel.dotaro@alcatel-lucent.fr

Gert Grammel
Alcatel-Lucent SEL
Lorenzstrasse, 10
70435 Stuttgart, Germany
Email: gert.grammel@alcatel-lucent.de

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).