

Internet Engineering Task Force (IETF)
Request for Comments: 6001
Updates: 4202, 4203, 4206, 4874, 4974, 5307
Category: Standards Track
ISSN: 2070-1721

D. Papadimitriou
M. Vigoureux
Alcatel-Lucent
K. Shiomoto
NTT
D. Brungard
ATT
JL. Le Roux
France Telecom
October 2010

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

Abstract

There are specific requirements for the support of networks comprising Label Switching Routers (LSRs) participating in 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. It covers the elements of a single GMPLS control plane instance controlling multiple Label Switched Path (LSP) regions or layers within a single Traffic Engineering (TE) domain.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6001>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	3
1.1. Conventions Used in This Document	4
2. Summary of the Requirements and Evaluation	4
3. Interface Adjustment Capability Descriptor (IACD)	5
3.1. Overview	5
3.2. Interface Adjustment Capability Descriptor (IACD)	6
4. Multi-Region Signaling	9
4.1. XRO Subobjects	10
5. Virtual TE Link	12
5.1. Edge-to-Edge Association	13
5.2. Soft Forwarding Adjacency (Soft FA)	16
6. Backward Compatibility	18
7. Security Considerations	18
8. IANA Considerations	18
8.1. RSVP	18
8.2. OSPF	20
8.3. IS-IS	20
9. References	20
9.1. Normative References	20
9.2. Informative References	22
Acknowledgments.....	23
Contributors	23

1. Introduction

Generalized Multi-Protocol Label Switching (GMPLS) [RFC3945] extends MPLS to handle multiple switching technologies: packet switching (PSC), Layer 2 switching (L2SC), Time-Division Multiplexing (TDM) Switching, 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 LSP 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 participating in 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 TE domain is defined as group of Label Switching Routers (LSRs) that enforces a common TE policy. A Control Plane (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 MLN/MRN unified control. The rationales and requirements for Multi-Layer/Region networks are set forth in [RFC5212]. These requirements are evaluated against GMPLS protocols in [RFC5339] 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.

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].

In addition, the reader is assumed to be familiar with [RFC3945], [RFC3471], [RFC4201], [RFC4202], [RFC4203], [RFC4206], and [RFC5307].

2. Summary of the Requirements and Evaluation

As identified in [RFC5339], most MLN/MRN requirements rely on mechanisms and procedures (such as local procedures and policies, or specific TE mechanisms and algorithms) that are outside the scope of the GMPLS protocols, and thus do not require any GMPLS protocol extensions.

Four areas for extensions of GMPLS protocols and procedures have been identified in [RFC5339]:

- o GMPLS routing extensions for the advertisement of the internal adjustment capability of hybrid nodes. See Section 3.2.2 of [RFC5339].

- o GMPLS signaling extensions for constrained multi-region signaling (Switching Capability inclusion/exclusion). See Section 3.2.1 of [RFC5339]. An additional eXclude Route Object (XRO) Label subobject is also defined since it was absent from [RFC4874].
- o GMPLS signaling extensions for the setup/deletion of virtual TE links (as well as exact trigger for its actual provisioning). See Section 3.1.1.2 of [RFC5339].
- o GMPLS routing and signaling extensions for graceful TE link deletion. See Section 3.1.1.3 of [RFC5339].

The first three requirements are addressed in Sections 3, 4, and 5 of this document, respectively. The fourth requirement is addressed in [RFC5710] with additional context provided by [RFC5817].

3. Interface Adjustment Capability Descriptor (IACD)

In the MRN context, nodes that have at least one interface that supports more than one switching capability are called hybrid nodes [RFC5212]. The logical composition of a hybrid node contains at least two distinct switching elements that are interconnected by "internal links" to provide adjustment between the supported switching capabilities. These internal links have finite capacities that MUST be taken into account when computing the path of a multi-region TE-LSP. The advertisement of the internal adjustment 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 multiple switching capabilities, such as PSC and TDM or PSC and LSC, all under the control of a single GMPLS instance.

These nodes, hosting multiple Interface Switching Capabilities (ISCs) [RFC4202], are required to hold and advertise resource information on link states and topology, just like other nodes (hosting a single ISC). They may also have to consider some portions of internal node resources use to terminate hierarchical LSPs, since in circuit-switching technologies (such as TDM, LSC, and FSC) LSPs require the use of resources allocated in a discrete manner (as predetermined by the switching type). For example, a node with PSC+LSC hierarchical switching capability can switch a lambda LSP, but cannot terminate the Lambda LSP if there is no available (i.e., not already in use) adjustment capability between the LSC and the PSC switching components. Another example occurs when L2SC (Ethernet) switching can be adapted in the Link Access Procedure-SDH (LAPS) X.86 and

Generic Framing Procedure (GFP) for instance, before reaching the TDM switching matrix. Similar circumstances can occur, for example, 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 adjustment capability to terminate LSPs (not the interface capability since the latter can be inferred from the bandwidth available for each switching capability) provides the 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 adjustment capability, e.g., to terminate LSPs at the PSC or LSC level.

Hence, we introduce the capability of discriminating the (internal) adjustment capability from the (interface) switching capability by defining an Interface Adjustment Capability Descriptor (IACD).

A more detailed problem statement can be found in [RFC5339].

3.2. Interface Adjustment Capability Descriptor (IACD)

The Interface Adjustment Capability Descriptor (IACD) provides the information for the forwarding/switching capability.

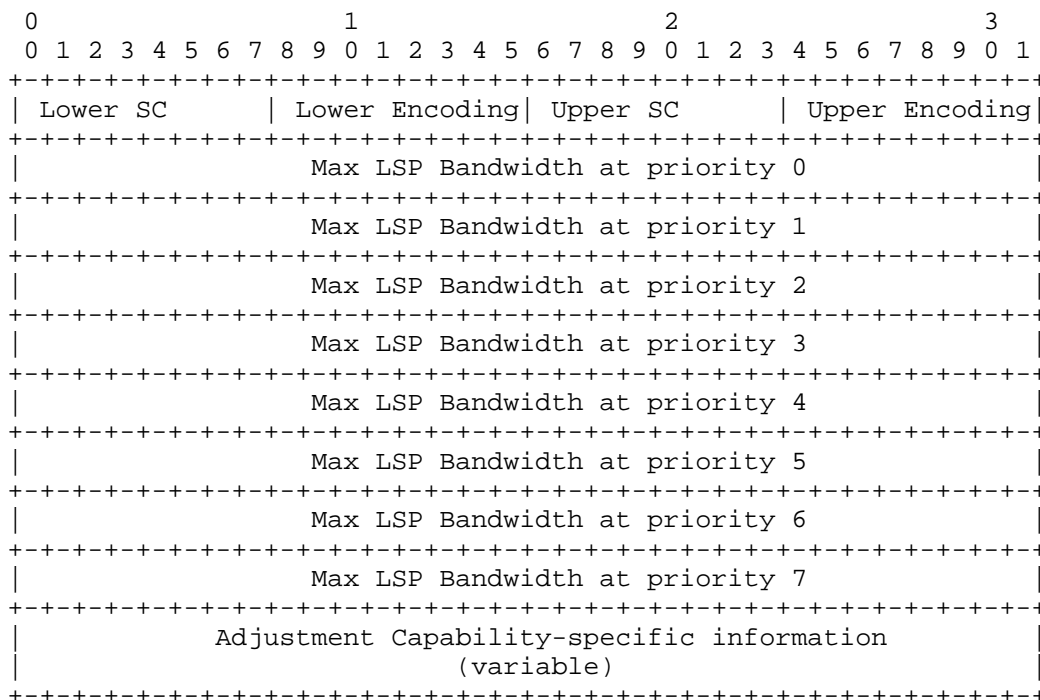
Note that the addition of the IACD as a TE link attribute does not modify the format of the Interface Switching Capability Descriptor (ISCD) defined in [RFC4202], and does not change how the ISCD sub-TLV is carried in the routing protocols or how it is processed when it is received [RFC4201], [RFC4203].

The receiving LSR uses its Link State Database to determine the IACD(s) of the far end of the link. Different Interface Adjustment Capabilities at two ends of a TE link are allowed.

3.2.1. OSPF

In OSPF, the IACD sub-TLV is defined as an optional sub-TLV of the TE Link TLV (Type 2, see [RFC3630]), with Type 25 and variable length.

The IACD sub-TLV format is defined as follows:



Lower Switching Capability (SC) field (byte 1) - 8 bits

Indicates the lower switching capability associated with the Lower Encoding field (byte 2). The value of the Lower Switching Capability field MUST be set to the value of Switching Capability of the ISCD sub-TLV advertised for this TE link. If multiple ISCD sub-TLVs are advertised for that TE link, the Lower Switching Capability (SC) value MUST be set to the value of SC to which the adjustment capacity is associated.

Lower Encoding (byte 2) - 8 bits

Contains one of the LSP Encoding Type values specified in Section 3.1.1 of [RFC3471] and updates.

Upper Switching Capability (SC) field (byte 3) - 8 bits

Indicates the upper switching capability. The Upper Switching Capability field MUST be set to one of the values defined in [RFC4202].

Upper Encoding (byte 4) - 8 bits

Set to the encoding of the available adjustment capacity 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. The adjustment capacity is the set of resources associated to the upper switching capability.

Max LSP Bandwidth

The Maximum LSP Bandwidth is encoded as a list of eight 4-octet fields in the IEEE floating point format [IEEE], with priority 0 first and priority 7 last. The units are bytes per second. Processing MUST follow the rules specified in [RFC4202].

The Adjustment Capability-specific information - variable

This field is defined so as to leave the possibility for future addition of technology-specific information associated to the adjustment capability.

Other fields MUST be processed as specified in [RFC4202] and [RFC4203].

The bandwidth values provide an indication of the resources still available to perform insertion/extraction for a given adjustment at a given priority (resource pool concept: set of shareable available resources that can be assigned dynamically).

Multiple IACD sub-TLVs MAY be present within a given TE Link TLV.

The presence of the IACD sub-TLV as part of the TE Link TLV does not modify the format/messaging and the processing associated to the ISCD sub-TLV defined in [RFC4203].

3.2.2. IS-IS

In IS-IS, the IACD sub-TLV is an optional sub-TLV of the Extended IS Reachability TLV (see [RFC5305]) with Type 27.

The IACD sub-TLV format is identical to the OSPF sub-TLV format defined in Section 3.2.1. The fields of the IACD sub-TLV have the same processing and interpretation rules as defined in Section 3.2.1.

Multiple IACD sub-TLVs MAY be present within a given extended IS reachability TLV.

The presence of the IACD sub-TLV as part of the extended IS reachability TLV does not modify format/messaging and processing associated to the ISCD sub-TLV defined in [RFC5307].

4. Multi-Region Signaling

Section 6.2 of [RFC4206] specifies that when a region boundary node receives a Path message, the node determines whether or not it is at the edge of an LSP region with respect to the Explicit Route Object (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 Explicit Route Object (ERO), using the IGP database. The node then extracts from the ERO the sub-sequence of hops from itself to the other end of the region.

The node then compares the sub-sequence of hops with all existing Forwarding Agency LSPs (FA-LSPs) originated by the node:

- o If a match is found, that FA-LSP has enough unreserved bandwidth for the LSP being signaled, and the Generalized PID (G-PID) of the FA-LSP is compatible with the G-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 previous hop (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 endpoint of the FA-LSP.
- o 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.

Note: compatible G-PID implies that traffic can be processed by both ends of the FA-LSP without dropping traffic after its establishment.

Applying the procedure of [RFC4206] in an MRN environment MAY lead to the setup of single-hop FA-LSPs between each pair of nodes. 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 over which link. Particularly, in case a TE link has multiple SCs advertised as part of its ISCD sub-TLVs, an ERO does not provide a mechanism to select a particular SC.

In order to limit the modifications to existing RSVP-TE procedures ([RFC3473] and referenced), this document defines a new subobject of the eXclude Route Object (XRO), see [RFC4874], called the Switching Capability subobject. This subobject enables (when desired) the explicit identification of at least one switching capability to be excluded from the resource selection process described above.

Including this subobject 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).

The approach followed here is to concentrate exclusions in XRO and inclusions in ERO. Indeed, the ERO specifies the topological characteristics of the path to be signaled. Usage of Explicit Exclusion Route Subobjects (EXRSs) would also lead in the exclusion over certain portions of the LSP during the FA-LSP setup. Thus, it is more suited to extend generality of the elements excluded by the XRO but also prevent complex consistency checks as well as transpositions between EXRS and XRO at FA-LSP head-ends.


4.1. XRO Subobjects

The contents of an EXCLUDE_ROUTE object defined in [RFC4874] are a series of variable-length data items called subobjects.

This document defines the Switching Capability (SC) subobject of the XRO (Type 35), its encoding, and processing. It also complements the subobjects defined in [RFC4874] with a Label subobject (Type 3).

4.1.1. SC Subobject

XRO subobject Type 35: Switching Capability



L (1 bit)

0 indicates that the attribute specified MUST be excluded.

1 indicates that the attribute specified SHOULD be avoided.

Type (7 bits)

The Type of the XRO SC subobject is 35.

Length (8 bits)

The total length of the subobject in bytes (including the Type and Length fields). The Length of the XRO SC subobject is 4.

Attribute (8 bits)

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.


The Switching Capability subobject MUST follow the set of one or more numbered or unnumbered interface subobjects to which this subobject refers.

In the case of a loose-hop ERO subobject, the XRO subobject MUST precede the loose-hop subobject identifying the tail-end node/interface of the traversed region(s).

4.1.2. Label Subobject

The encoding of the XRO Label subobject is identical to the Label ERO subobject defined in [RFC3473] with the exception of the L bit. The XRO Label subobject is defined as follows:

XRO Subobject Type 3: Label Subobject



L (1 bit)

0 indicates that the attribute specified MUST be excluded.

1 indicates that the attribute specified SHOULD be avoided.

Type (7 bits)

The Type of the XRO Label subobject is 3.

Length (8 bits)

The total length of the subobject in bytes (including the Type and Length fields). The Length is always divisible by 4.

U (1 bit)

See [RFC3471].

C-Type (8 bits)

The C-Type of the included Label Object. Copied from the Label Object (see [RFC3471]).

Label

See [RFC3471].

XRO Label subobjects MUST follow the numbered or unnumbered interface subobjects to which they refer, and, when present, MUST also follow the Switching Capability subobject.

When XRO Label subobjects are following the Switching Capability subobject, the corresponding label values MUST be compatible with the SC capability to be explicitly excluded.

5. Virtual TE Link

A virtual TE link is defined as a TE link between two upper-layer nodes that is not associated with a fully provisioned FA-LSP in a lower layer [RFC5212]. A virtual TE link is advertised as any TE link, following the rules in [RFC4206] defined for fully provisioned TE links. A virtual TE link represents thus the potentiality to set up an FA-LSP in the lower layer to support the TE link that has been advertised. In particular, the flooding scope of a virtual TE link is within an IGP area, as is the case for any 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. The procedures described in this section complement those defined in [RFC4206] and [HIER-BIS].

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 (see [RFC4974]). This technique consists of exchanging identification and TE attributes information directly between TE link endpoints through the establishment of a call between terminating LSRs. These TE link endpoints correspond to the LSP head-end and tail-end points of the LSPs that will be established. The endpoints MUST belong to the same (LSP) region.

Once the call is established, the resulting association populates the local Traffic Engineering DataBase (TEDB) and the resulting virtual TE link is advertised as any other TE link. The latter can then be used to attract traffic. When an upper-layer/region LSP tries to make use of this virtual TE link, one or more FA LSPs MUST be established using the procedures defined in [RFC4206] to make the virtual TE link "real" and allow it to carry traffic by nesting the upper-layer/region LSP.

In order to distinguish usage of such call from the call and associated procedures defined in [RFC4974], 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 modeled on the LSP_ATTRIBUTES object defined in [RFC5420]. The CALL_ATTRIBUTES object MAY also be used to report call operational state on a Notify message.

The CALL_ATTRIBUTES object class is 202 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.

CALL_ATTRIBUTES class = 202, C-Type = 1



The Call Attributes TLVs are encoded as described in Section 5.1.3.

5.1.2. Processing


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.

5.1.3. Call Attributes TLVs

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

There are no ordering rules for Call Attributes TLVs, and no interpretation SHOULD be placed on the order in which these TLVs are received.

Each Call Attributes TLV carried by the CALL_ATTRIBUTES object is encoded as follows:



Type

The identifier of the TLV.

Length

Indicates the total length of the TLV in octets. That is, the combined length of the Type, Length, and Value fields, i.e., four plus the length of the Value field in octets.

The entire TLV MUST be padded with between zero and three trailing zeros to make it four-octet aligned. The Length field does not count any padding.

Value

The data field for the TLV padded as described above.

Assignment of Call Attributes TLV types MUST follow the rules specified in Section 8 (IANA Considerations).

5.1.4. Call Attributes Flags TLV

The Call Attributes TLV of Type 1 defines the Call Attributes Flags TLV. The Call Attributes Flags TLV MAY be present in a CALL_ATTRIBUTES object.

The Call Attributes 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 MUST therefore always be a multiple of 4 bytes, regardless of the number of bits carried and no padding is required.

Unassigned bits are considered reserved and MUST be set to zero on transmission by the originator of the object. Bits not contained in the Call Attributes Flags TLV MUST be assumed to be set to zero. If the Call Attributes Flags 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. In other terms, assigned bits that are not present either because the Call Attributes Flags TLV is deliberately foreshortened or because the TLV is not included MUST be treated as though they are present and are set to zero.

5.1.5. Call Inheritance Flag

This document introduces a specific Call Inheritance Flag at position bit 0 (most significant bit) in the Call Attributes Flags TLV. This flag indicates that the association initiated between the endpoints 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 SHALL by default be 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) if the association had been established with this flag set to 1. In the latter case, the corresponding FA-LSP SHOULD also be torn down to prevent unused resources.

The Notify message used for establishing the association is defined as per [RFC4974]. Additionally, the Notify message MUST carry an LSP_TUNNEL_INTERFACE_ID Object, that allows identifying unnumbered FA-LSPs ([RFC3477], [RFC4206], [HIER-BIS]) and numbered FA-LSPs ([RFC4206], [HIER-BIS]).

5.2. Soft Forwarding Adjacency (Soft FA)

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 an FA is established, the corresponding TE link can be advertised following the procedures described in [RFC4206].

There are two techniques to set up Soft FAs:

- o The first one consists in setting up the FA LSP by precluding resource commitment during its establishment. These are known as pre-planned LSPs.
- o The second technique consists in making use of path-provisioned LSPs only. In this case, there is no associated resource demand during the LSP establishment. This can be considered as the RSVP-TE equivalent of the Null service type specified in [RFC2997].

5.2.1. Pre-Planned LSP Flag

The LSP ATTRIBUTES object and Attributes Flags TLV are defined in [RFC5420]. 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 bit 6 in accordance with IANA assignment. This flag, part of the Attributes Flags TLV, follows general processing of [RFC5420] for LSP_REQUIRED_ATTRIBUTE 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 along the LSP being established.

The Pre-Planned LSP flag can take one of the following values:

- o When set to 0, this means that the LSP MUST 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 MUST 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 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).

Mechanisms for provisioning (pre-planned or not) LSP with 0 bandwidth is straightforward for PSC LSP: in the SENDER_TSPEC/FLOWSPEC object, the Peak Data Rate field of IntServ objects (see [RFC2210]) MUST be set to 0. For L2SC LSP: the Committed Information Rate (CIR), Excess Information Rate (EIR), Committed Burst Size (CBS), and Excess Burst Size (EBS) values MUST be set to 0 in the Type 2 sub-TLV of the Ethernet Bandwidth Profile TLV. In both cases, upon LSP resource commitment, actual traffic parameter values are used to perform corresponding resource reservation.

However, mechanisms for provisioning (pre-planned or not) a TDM or LSC LSP with 0 bandwidth is currently not possible because the exchanged label value is tightly coupled with resource allocation during LSP signaling (e.g., see [RFC4606] for a SONET/SDH LSP). For TDM and LSC LSP, a NULL Label value is used to prevent resource allocation at the data plane level. In these cases, upon LSP resource commitment, actual label value exchange is performed to commit allocation of timeslots/ wavelenghts.

6. Backward Compatibility

New objects and procedures defined in this document are running within a given TE domain, defined as group of LSRs that enforces a common TE policy. Thus, the extensions defined in this document are expected to run in the context of a consistent TE policy. Specification of a consistent TE policy is outside the scope of this document.

In such TE domains, we distinguish between edge LSRs and intermediate LSRs. Edge LSRs MUST be able to process Call Attributes as defined in Section 5.1 if this is the method selected for 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

This document does not introduce any new security considerations from the ones already detailed in [RFC5920] that describes the MPLS and GMPLS security threats, the related defensive techniques, and the mechanisms for detection and reporting. Indeed, the applicability of the proposed GMPLS extensions is limited to single TE domain. Such a domain is under the authority of a single administrative entity. In this context, multiple switching layers comprised within such TE domain are under the control of a single GMPLS control plane instance.

Nevertheless, Call initiation, as depicted in Section 5.1, MUST strictly remain under control of the TE domain administrator. To prevent any abuse of Call setup, edge nodes MUST ensure isolation of their call controller (i.e., the latter is not reachable via external TE domains). To further prevent man-in-the-middle attacks, security associations MUST be established between edge nodes initiating and terminating calls. For this purpose, Internet Key Exchange (IKE) protocol [RFC5996] MUST be used for performing mutual authentication and establishing and maintaining these security associations.

8. IANA Considerations

8.1. RSVP

IANA has made the following assignments in the "Class Names, Class Numbers, and Class Types" section of the "RSVP PARAMETERS" registry available from <http://www.iana.org>.

This document introduces a new class named CALL_ATTRIBUTES, which has been created in the 11bbbbbb range with the following definition:

Class Number	Class Name	Reference
202	CALL_ATTRIBUTES	[RFC6001]

Class Type (C-Type):

1	Call Attributes	[RFC6001]
---	-----------------	-----------

IANA has established a "Call Attributes TLV" registry. The following types are defined:

TLV Value	Name	Reference
0	Reserved	[RFC6001]
1	Call Attributes Flags TLV	[RFC6001]

The values should be allocated based on the following allocation policy as defined in [RFC5226].

Range	Registration Procedures
0-32767	RFC Required
32768-65535	Reserved for Private Use

IANA has established a "Call Attributes Flags" registry. The following flags are defined:

Bit Number	32-bit Value	Name	Reference
0	0x80000000	Call Inheritance Flag	[RFC6001]

The values should be allocated based on the "RFC Required" policy as defined in [RFC5226].

This document introduces a new Flag in the Attributes Flags TLV defined in [RFC5420]:

Bit Number	Name	Reference
6	Pre-Planned LSP Flag	[RFC6001]

This document introduces two new subobjects for the EXCLUDE_ROUTE object [RFC4874], C-Type 1.

Subobject Type	Subobject Description
3	Label
35	Switching Capability (SC)

8.2. OSPF

IANA maintains the "Open Shortest Path First (OSPF) Traffic Engineering TLVs" registries including the "Types for sub-TLVs of TE link TLV (Value 2)" registry.

This document defines the following sub-TLV of TE link TLV (Value 2).

Value	Sub-TLV
25	Interface Adjustment Capability Descriptor (IACD)

8.3. IS-IS

This document defines the following new sub-TLV type of top-level TLV 22 that has been reflected in the ISIS sub-TLV registry for TLV 22, 141, and 222:

Type	Description	Length
27	Interface Adjustment Capability Descriptor (IACD)	Var.

9. References

9.1. Normative References

- [IEEE] IEEE, "IEEE Standard for Binary Floating-Point Arithmetic", Standard 754-1985, 1985.
- [RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC2210] Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, September 1997.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2997] Bernet, Y., Smith, A., and B. Davie, "Specification of the Null Service Type", RFC 2997, November 2000.

- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, January 2003.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4201] Kompella, K., Rekhter, Y., and L. Berger, "Link Bundling in MPLS Traffic Engineering (TE)", RFC 4201, October 2005.
- [RFC4202] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005.
- [RFC4203] Kompella, K., Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, 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.
- [RFC4606] Mannie, E. and D. Papadimitriou, "Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control", RFC 4606, August 2006.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, April 2007.

- [RFC4974] Papadimitriou, D. and A. Farrel, "Generalized MPLS (GMPLS) RSVP-TE Signaling Extensions in Support of Calls", RFC 4974, August 2007.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 2008.
- [RFC5307] Kompella, K., Ed., and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 5307, October 2008.
- [RFC5420] Farrel, A., Ed., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.

9.2. Informative References

- [HIER-BIS] Shiomoto, K., Ed., and A. Farrel, "Procedures for Dynamically Signaled Hierarchical Label Switched Paths", Work in Progress, February 2010.
- [RFC5212] Shiomoto, 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.
- [RFC5339] Le Roux, JL., Ed., and D. Papadimitriou, Ed., "Evaluation of Existing GMPLS Protocols against Multi-Layer and Multi-Region Networks (MLN/MRN)", RFC 5339, September 2008.
- [RFC5710] Berger, L., Papadimitriou, D., and JP. Vasseur, "PathErr Message Triggered MPLS and GMPLS LSP Reroutes", RFC 5710, January 2010.
- [RFC5817] Ali, Z., Vasseur, JP., Zamfir, A., and J. Newton, "Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks", RFC 5817, April 2010.

[RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

Acknowledgments

The authors would like to thank Mr. Wataru Imajuku for the discussions on adjustment between regions.

Contributors

Eiji Oki
University of Electro-Communications
1-5-1 Chofugaoka
Chofu, Tokyo 182-8585, Japan
EMail: oki@ice.uec.ac.jp

Ichiro Inoue
NTT Network Service Systems Laboratories
3-9-11 Midori-cho
Musashino-shi, Tokyo 180-8585, Japan
Phone: +81 422 596076
EMail: ichiro.inoue@lab.ntt.co.jp

Emmanuel Dotaro
Alcatel-Lucent France
Route de Villejust
91620 Nozay, France
Phone: +33 1 69634723
EMail: emmanuel.dotaro@alcatel-lucent.fr

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

Authors' Addresses

Dimitri Papadimitriou
Alcatel-Lucent
Copernicuslaan 50
B-2018 Antwerpen, Belgium
Phone: +32 3 2408491
EMail: dimitri.papadimitriou@alcatel-lucent.com

Martin Vigoureux
Alcatel-Lucent
Route de Villejust
91620 Nozay, France
Phone: +33 1 30772669
EMail: martin.vigoureux@alcatel-lucent.fr

Kohei Shiomoto
NTT
3-9-11 Midori-cho
Musashino-shi, Tokyo 180-8585, Japan
Phone: +81 422 594402
EMail: shiomoto.kohei@lab.ntt.co.jp

Deborah Brungard
ATT
Rm. D1-3C22 - 200 S. Laurel Ave.
Middletown, NJ 07748, USA
Phone: +1 732 4201573
EMail: dbrungard@att.com

Jean-Louis Le Roux
France Telecom
Avenue Pierre Marzin
22300 Lannion, France
Phone: +33 2 96053020
EMail: jean-louis.leroux@rd.francetelecom.com