

Internet Engineering Task Force (IETF)
Request for Comments: 6107
Updates: 3477, 4206
Category: Standards Track
ISSN: 2070-1721

K. Shiomoto, Ed.
NTT
A. Farrel, Ed.
Old Dog Consulting
February 2011

Procedures for Dynamically Signaled Hierarchical Label Switched Paths

Abstract

Label Switched Paths (LSPs) set up in Multiprotocol Label Switching (MPLS) or Generalized MPLS (GMPLS) networks can be used to form links to carry traffic in those networks or in other (client) networks.

Protocol mechanisms already exist to facilitate the establishment of such LSPs and to bundle traffic engineering (TE) links to reduce the load on routing protocols. This document defines extensions to those mechanisms to support identifying the use to which such LSPs are to be put and to enable the TE link endpoints to be assigned addresses or unnumbered identifiers during the signaling process.

The mechanisms defined in this document deprecate the technique for the signaling of LSPs that are to be used as numbered TE links described in RFC 4206.

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/rfc6107>.

Copyright Notice

Copyright (c) 2011 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.

Table of Contents

1. Introduction and Problem Statement	3
1.1. Background	4
1.1.1. Hierarchical LSPs	4
1.1.2. LSP Stitching Segments	5
1.1.3. Private Links	5
1.1.4. Routing Adjacencies	5
1.1.5. Forwarding Adjacencies	5
1.1.6. Client/Server Networks	6
1.1.7. Link Bundles	6
1.2. Desired Function	7
1.3. Existing Mechanisms	7
1.3.1. LSP Setup	7
1.3.2. Routing Adjacency Establishment and Link State Advertisement	7
1.3.3. TE Link Advertisement	7
1.3.4. Configuration and Management Techniques	8
1.3.5. Signaled Unnumbered FAs	8
1.3.6. Establishing Numbered FAs through Signaling and Routing	9
1.4. Overview of Required Extensions	10
1.4.1. Efficient Signaling of Numbered FAs	10
1.4.2. LSPs for Use as Private Links	10
1.4.3. Signaling an LSP for Use in Another Network	10
1.4.4. Signaling an LSP for Use in a Link Bundle	11
1.4.5. Support for IPv4 and IPv6	11
1.4.6. Backward Compatibility	11
2. Overview of Solution	11
2.1. Common Approach for Numbered and Unnumbered Links	11
2.2. LSP Usage Indication	12
2.3. IGP Instance Identification	12
2.4. Link Bundle Identification	12

2.5. Backward Compatibility	12
3. Mechanisms and Protocol Extensions	13
3.1. LSP_TUNNEL_INTERFACE_ID Object	13
3.1.1. Existing Definition and Usage	13
3.1.2. Unnumbered Links with Action Identification	13
3.1.3. IPv4 Numbered Links with Action Identification	16
3.1.4. IPv6 Numbered Links with Action Identification	17
3.2. Target IGP Identification TLV	18
3.3. Component Link Identification TLV	19
3.3.1. Unnumbered Component Link Identification	20
3.3.2. IPv4 Numbered Component Link Identification	20
3.3.3. IPv6 Numbered Component Link Identification	21
3.4. Link State Advertisement	21
3.5. Message Formats	22
3.6. Error Cases and Non-Acceptance	22
3.7. Backward Compatibility	24
4. Security Considerations	25
5. IANA Considerations	25
5.1. New Class Types	25
5.2. Hierarchy Actions	26
5.3. New Error Codes and Error Values	26
6. Acknowledgements	27
7. References	27
7.1. Normative References	27
7.2. Informative References	28

1. Introduction and Problem Statement

Traffic Engineering (TE) links in a Multiprotocol Label Switching (MPLS) or a Generalized MPLS (GMPLS) network may be constructed from Label Switched Paths (LSPs) [RFC4206]. Such LSPs are known as hierarchical LSPs (H-LSPs).

The LSPs established in one network may be used as TE links in another network, and this is particularly useful when a server layer network (for example, an optical network) provides LSPs for use as TE links in a client network (for example, a packet network). This enables the construction of a multilayer network (MLN) [RFC5212].

When the number of TE links (created from LSPs or otherwise) between a pair of nodes grows large, it is inefficient to advertise them individually. This may cause scaling issues in configuration and in the routing protocols used to carry the advertisements. The solution (described in [RFC4201]) is to collect the TE links together and to advertise them as a single TE link called a link bundle.

These various mechanisms have proved to be very powerful in building dynamically provisioned networks, but, as set out later in this document, several issues have been identified during deployment with how LSPs are established and made available for use as H-LSPs or as components of a link bundle, and with how these links are advertised appropriately in an interior gateway protocol (IGP). These issues relate to how the LSP's endpoints coordinate two things: the use to which the LSP is put and the identifiers of the endpoints.

This document provides solutions to these issues by defining mechanisms so that the ends of signaled LSPs can exchange information about:

- Whether the LSP is an ordinary LSP or an H-LSP.
- In which IGP instances the LSP should be advertised as a link.
- How the client networks should make use of the new links.
- Whether the link should form part of a bundle (and if so, which bundle).
- How the link endpoints should be identified when advertised.

This document deprecates one of the mechanisms defined in [RFC4206] for the signaling of LSPs that are to be used as numbered TE links (see Sections 1.3.6 and 1.4.6 for more details), and provides extensions to the other mechanisms defined in [RFC4206] so that the use to which the new LSP is to be put may be indicated during signaling. It also extends the mechanisms defined in [RFC3477] for signaling unnumbered TE links.

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].

1.1. Background

1.1.1. Hierarchical LSPs

[RFC3031] describes how MPLS labels may be stacked so that LSPs may be nested with one LSP running through another. This concept of H-LSPs is formalized in [RFC4206] with a set of protocol mechanisms for the establishment of an H-LSP that can carry one or more other LSPs.

[RFC4206] goes on to explain that an H-LSP may carry other LSPs only according to their switching types. This is a function of the way labels are carried. In a packet switch capable (PSC) network, the H-LSP can carry other PSC LSPs using the MPLS label stack. In non-packet networks where the label is implicit, label stacks are not possible, and H-LSPs rely on the ability to nest switching

technologies. Thus, for example, a lambda switch capable (LSC) LSP can carry a time-division multiplexing (TDM) LSP, but cannot carry another LSC LSP.

Signaling mechanisms defined in [RFC4206] allow an H-LSP to be treated as a single hop in the path of another LSP (i.e., one hop of the LSP carried by the H-LSP). This mechanism is known as "non-adjacent signaling".

1.1.2. LSP Stitching Segments

LSP stitching is defined in [RFC5150]. It enables LSPs of the same switching type to be included (stitched) as hops in an end-to-end LSP. The stitching LSP (S-LSP) is used in the control plane in the same way as an H-LSP, but in the data plane the LSPs are stitched so that there is no label stacking or nesting of technologies. Thus, an S-LSP must be of the same switching technology as the end-to-end LSP that it facilitates.

1.1.3. Private Links

An H-LSP or S-LSP can be used as a private link. Such links are known by their endpoints, but are not more widely known and are not advertised by routing protocols. They can be used to carry traffic between the endpoints, but are not usually used to carry traffic that is going beyond the egress of the LSP.

1.1.4. Routing Adjacencies

A routing adjacency is formed between two IGP speakers that are logically adjacent. In this sense, 'logically adjacent' means that the routers have a protocol tunnel between them through which they can exchange routing protocol messages. The tunnel is also usually available for carrying IP data although a distinction should be made in GMPLS networks between data-plane traffic and control-plane traffic.

Routing adjacencies for forwarding data traffic are only relevant in PSC networks (i.e., IP/MPLS networks).

1.1.5. Forwarding Adjacencies

A Forwarding Adjacency (FA) is defined in [RFC4206] as a data link created from an LSP and advertised in the same instance of the control plane that advertises the TE links from which the LSP is constructed. The LSP itself is called an FA-LSP.

Thus, an H-LSP or S-LSP may form an FA such that it is advertised as a TE link in the same instance of the routing protocol as was used to advertise the TE links that the LSP traverses.

As observed in [RFC4206], the nodes at the ends of an FA would not usually have a routing adjacency.

1.1.6. Client/Server Networks

An LSP may be created in one network and used as a link (sometimes called a virtual link) in another network [RFC5212]. In this case, the networks are often referred to as server and client networks, respectively.

The server network LSP is used as an H-LSP or an S-LSP as described above, but it does not form an FA because the client and server networks run separate instances of the control-plane routing protocols.

The virtual link may be used in the client network as a private link or may be advertised in the client network IGP. Additionally, the link may be used in the client network to form a routing adjacency and/or as a TE link.

1.1.7. Link Bundles

[RFC4201] recognizes that a pair of adjacent routers may have a large number of TE links that run between them. This can be a considerable burden to the operator who may need to configure them and to the IGP that must distribute information about each of them. A TE link bundle is defined by [RFC4201] as a TE link that is advertised as an aggregate of multiple TE links that could have been advertised in their own right. All TE links that are collected into a TE link bundle have the same TE properties.

Thus, a link bundle is a shorthand that improves the scaling properties of the network.

Since a TE link may, itself, be an LSP (either an FA or a virtual link), a link bundle may be constructed from FA-LSPs or virtual links.

1.2. Desired Function

It should be possible to signal an LSP and automatically coordinate its use and advertisement in any of the ways described in Section 1.3 with minimum involvement from an operator. The mechanisms used should be applicable to numbered and unnumbered links and should not require implementation complexities.

1.3. Existing Mechanisms

This section briefly introduces existing protocol mechanisms used to satisfy the desired function described in Section 1.2.

1.3.1. LSP Setup

Both unidirectional LSPs and bidirectional LSPs are signaled from the ingress label switching router (LSR) to the egress LSR. That is, the ingress LSR is the initiator, and the egress learns about the LSP through the signaling protocol [RFC3209] [RFC3473].

1.3.2. Routing Adjacency Establishment and Link State Advertisement

Although hosts can discover routers (for example, through ICMP [RFC1256]), routing adjacencies are usually configured at both ends of the adjacency. This requires that each router know the identity of the router at the other end of the link connecting the routers, and know that the IGP is to be run on this link.

Once a routing adjacency has been established, a pair of routers may use the IGP to share information about the links available for carrying IP traffic in the network.

Suitable routing protocols are OSPF version 2 [RFC2328], OSPF version 3 [RFC5340], and IS-IS [RFC1195].

1.3.3. TE Link Advertisement

Extensions have been made to IGPs to advertise TE link properties ([RFC3630], [RFC5329], [RFC5305], [RFC5308], and [ISIS-IPV6-TE]) and also to advertise link properties in GMPLS networks ([RFC4202], [RFC4203], and [RFC5307]).

TE link advertisement can be performed by the same instance of the IGP as is used for normal link state advertisement, or can use a separate instance. Furthermore, the ends of a TE link advertised in an IGP do not need to form a routing adjacency. This is particularly the case with FAs as described in Section 1.1.5.

1.3.4. Configuration and Management Techniques

LSPs are usually created as the result of a management action. This is true even when a control plane is used: the management action is a request to the control plane to signal the LSP.

If the LSP is to be used as an H-LSP or S-LSP, management commands can be used to install the LSP as a link. The link must be defined, interface identifiers allocated, and the endpoints configured to know about (and advertise, if necessary) the new link.

If the LSP is to be used as part of a link bundle, the operator must decide which bundle it forms part of and ensure that information is configured at the ingress and egress, along with the necessary interface identifiers.

These mechanisms are perfectly functional and quite common in MLNs, but require configuration coordination and additional management. They are open to user error and misconfiguration that might result in the LSP not being used (a waste of resources) or the LSP being made available in the wrong way with some impact on traffic engineering.

1.3.5. Signaled Unnumbered FAs

[RFC3477] describes how to establish an LSP and to make it available automatically as a TE link in the same instance of the routing protocol as advertises the TE links it traverses, using IPv4-based unnumbered identifiers to identify the new TE link. That is, [RFC3477] describes how to create an unnumbered FA.

The mechanism, as defined in Section 3 of [RFC3477], is briefly as follows:

- The ingress of the LSP signals the LSP as normal using a Path message, and includes an LSP_TUNNEL_INTERFACE_ID object. The LSP_TUNNEL_INTERFACE_ID object identifies:
 - The sender's LSR Router ID
 - The sender's interface ID for the TE link being created
- The egress of the LSP responds as normal to accept the LSP and set it up, and includes an LSP_TUNNEL_INTERFACE_ID object. The LSP_TUNNEL_INTERFACE_ID object identifies:
 - The egress's LSR Router ID
 - The egress's interface ID for the TE link being created.

- Following the exchange of the Path and Resv messages, both the ingress and the egress know that the LSP is to be advertised as a TE link in the same instance of the routing protocol as was used to advertise the TE links that it traverses, and also know the unnumbered interface identifiers to use in the advertisement.

[RFC3477] does not include mechanisms to support IPv6-based unnumbered identifiers, nor IPv4 or IPv6 numbered identifiers.

1.3.6. Establishing Numbered FAs through Signaling and Routing

[RFC4206] describes procedures to establish an LSP and to make it available automatically as a TE link that is identified using IPv4 addresses in the same instance of the routing protocol as advertises the TE links it traverses (that is, as a numbered FA).

The mechanism, as defined in [RFC4206], is briefly as follows:

- The ingress of the LSP signals the LSP as normal using a Path message, and sets the IPv4 tunnel sender address to the IP address it will use to identify its interface for the TE link being created. This is one address from a /31 pair.
- The egress of the LSP responds as normal to accept the LSP and set it up. It infers the address that it must assign to identify its interface for the TE link being created as the partner address of the /31 pair.
- The ingress decides whether the LSP is to be advertised as a TE link (i.e., as an FA). If so, it advertises the link in the IGP in the usual way.
- If the link is unidirectional, nothing more needs to be done. If the link is bidirectional, the egress must also advertise the link, but it does not know that advertisement is required as there is no indication in the signaling messages.
- When the ingress's advertisement of the link is received by the egress, it must check to see whether it is the egress of the LSP that formed the link. Typically, this means the egress:
 - Checks to see if the link advertisement is new.
 - Checks to see if the Link-ID address in the received advertisement matches its own TE Router ID.
 - Checks the advertising router ID from the advertisement against the ingress address of each LSP for which it is the egress.
 - Deduces the LSP that has been advertised as a TE link and issues the corresponding advertisement for the reverse direction.

It is possible that some reduction in processing can be achieved by mapping based on the /31 pairing, but nevertheless, there is a fair amount of processing required, and this does not scale well in large networks.

Note that this document deprecates this procedure as explained in Section 1.4.6.

No explanation is provided in [RFC4206] of how to create numbered IPv6 FAs.

1.4. Overview of Required Extensions

This section provides a brief outline of the required protocol extensions.

1.4.1. Efficient Signaling of Numbered FAs

The mechanism described in Section 1.3.6 is inefficient. The egress must wait until it receives an advertisement from the ingress before it knows that the LSP is to be installed as a TE link and advertised as an FA. Further, it must parse all received advertisements to determine if any is the trigger for it to advertise an FA.

An efficient signaling mechanism is required so that the egress is informed by the ingress during LSP establishment.

1.4.2. LSPs for Use as Private Links

There is currently no mechanism by which an ingress can indicate that an LSP is set up for use as a private link. Any attempt to make it a link would currently cause it to be advertised as an FA.

A signaling mechanism is needed to identify the use to which an LSP is to be put.

1.4.3. Signaling an LSP for Use in Another Network

The existing signaling/routing mechanisms are designed for use in the creation of FAs. That is, the TE link created is advertised in the single IGP instance.

The numbered TE link mechanism (Section 1.3.6) could, in theory, be used in an MLN with multiple IGP instances if the addressing model is kept consistent across the networks, and if the egress searches all advertisements in all IGP instances. However, this is complex and does not work for unnumbered interfaces.

A signaling mechanism is required to indicate in which IGP instance the TE link should be advertised.

1.4.4. Signaling an LSP for Use in a Link Bundle

A signaling mechanism is required to indicate that an LSP is intended to form a component link of a link bundle, and to identify the bundle and the IGP instance in which the bundle is advertised.

1.4.5. Support for IPv4 and IPv6

The protocol mechanisms must support IPv4 and IPv6 numbered and unnumbered TE links.

1.4.6. Backward Compatibility

The existing protocol mechanisms for unnumbered FAs as defined in [RFC4206] and [RFC3477] must continue to be supported, and new mechanisms must be devised such that their introduction will not break existing implementations or deployments.

Note that an informal survey in the CCAMP working group established that there are no significant deployments of numbered FAs established using the technique described in [RFC4206] and set out in Section 1.3.6. Therefore, this document deprecates this procedure.

2. Overview of Solution

This section provides an overview of the mechanisms and protocol extensions defined in this document. Details are presented in Section 3.

2.1. Common Approach for Numbered and Unnumbered Links

The LSP_TUNNEL_INTERFACE_ID object [RFC3477] is extended for use for all H-LSPs and S-LSPs whether they form numbered or unnumbered IPv4 or IPv6 links. Different Class Types of the object identify the address type for which it applies.

2.2. LSP Usage Indication

The LSP_TUNNEL_INTERFACE_ID object is given flags in a new Actions field to say how the LSP is to be used. The following categories are supported:

- The LSP is used as an advertised TE link.
- The LSP is used to form a routing adjacency.
- The LSP is used to form an advertised TE link and a routing adjacency.
- The LSP forms a private link and is not advertised.
- The LSP is used as part of a link bundle.
- The LSP is used as a hierarchical LSP or a stitching segment.

2.3. IGP Instance Identification

An optional TLV is added to the LSP_TUNNEL_INTERFACE_ID object to identify the IGP instance into which the link formed by the LSP is to be advertised. If the TLV is absent and the link is to be advertised as indicated by the Actions field, the link is advertised in the same instance of the IGP as was used to advertise the TE links it traverses.

2.4. Link Bundle Identification

When an LSP is to be used as a component link of a link bundle, the LSP_TUNNEL_INTERFACE_ID object refers to the bundle indicating how the bundle is addressed and used, and a new TLV indicates the component link identifier for the link that the LSP creates.

2.5. Backward Compatibility

Backward compatibility has three aspects.

- Existing mechanisms for unnumbered FAs described in [RFC3477] and [RFC4206] must continue to work, so that ingress nodes do not have to be updated when egresses are updated.
- Existing mechanisms for establishing numbered FAs described in [RFC4206] are safely deprecated by this document, as they are not significantly deployed.
- New mechanisms must be gracefully rejected by existing egress implementations so that egress nodes do not have to be updated when ingresses are updated.

3. Mechanisms and Protocol Extensions

This section defines protocol mechanisms and extensions to achieve the function described in the previous section.

3.1. LSP_TUNNEL_INTERFACE_ID Object

The principal signaling protocol element used to achieve all of the required functions is the LSP_TUNNEL_INTERFACE_ID object defined in [RFC3477]. The existing definition and usage continues to be supported as described in the next section. Subsequent sections describe new variants of the object (denoted by new Class Types), and additional information carried in the object by means of extensions.

3.1.1. Existing Definition and Usage

This document does not deprecate the mechanisms defined in [RFC3477] and [RFC4206]. Those procedures must continue to operate as described in Section 3.7.

That means that the LSP_TUNNEL_INTERFACE_ID object with Class Type 1 remains unchanged, and can be used to establish an LSP that will be advertised as an unnumbered TE link in the same instance of the IGP as was used to advertise the TE links that the LSP traverses (that is, as an FA). The procedure is unchanged and operates as summarized in Section 1.3.5.

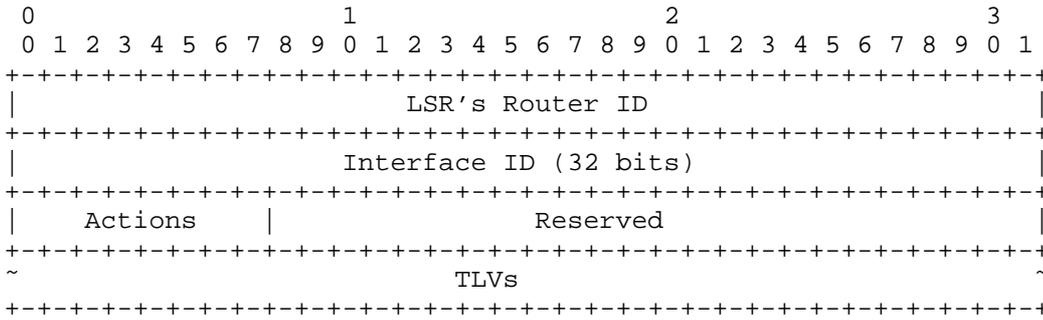
[RFC3477] does not make any suggestions about where in Path or Resv messages the LSP_TUNNEL_INTERFACE_ID object should be placed. See Section 3.5 for recommended placement of this object in new implementations.

3.1.2. Unnumbered Links with Action Identification

A new C-Type variant of the LSP_TUNNEL_INTERFACE_ID object is defined to carry an unnumbered interface identifier and to indicate into which instance of the IGP the consequent TE link should be advertised. This does not deprecate the use of C-Type 1.

The format of the object is as shown below.

C-NUM = 193, C-Type = 4



LSR's Router ID

Unchanged from the definition in [RFC3477].

Interface ID

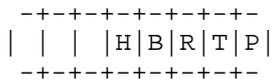
Unchanged from the definition in [RFC3477].

Actions

This field specifies how the LSP that is being set up is to be treated.

The field has meaning only on a Path message. On a Resv message, the field SHOULD be set to reflect the value received on the corresponding Path message, and it MUST be ignored on receipt.

The field is composed of bit flags as follows:



P-flag (Private)

- 0 means that the LSP is to be advertised as a link according to the settings of the other flags.
- 1 means the LSP is to form a private link and is not to be advertised in the IGP, but is to be used according to the settings of the other flags.

T-flag (TE link)

- 0 means that the LSP is to be used as a TE link.
- 1 means that the LSP is not to be used as a TE link. It may still be used as an IP link providing a routing adjacency as defined by the R-flag.

R-flag (Routing adjacency)

0 means that the link is not to be used as a routing adjacency.

1 means that the link is to be used to form a routing adjacency.

B-flag (Bundle)

0 means that the LSP will not form part of a link bundle.

1 means that the LSP will form part of a bundle. See Section 3.3 for more details.

H-flag (Hierarchy/stitching)

The use of an LSP as an H-LSP or as an S-LSP is usually implicit from the network technologies of the networks and the LSP, but this is not always the case (for example, in PSC networks).

0 means that the LSP is to be used as a hierarchical LSP.

1 means that the LSP is to be used as a stitching segment.

Other bits are reserved for future use. They MUST be set to zero on transmission and SHOULD be ignored on receipt.

Note that all defined bit flags have meaning at the same time. An LSP that is to form an FA would carry the Actions field set to 0x00. That is:

- P=0 (advertised)
- T=0 (TE link)
- R=0 (not a routing adjacency)
- B=0 (not a bundle)
- H=0 (hierarchical LSP)

Reserved

The Reserved bits MUST be set to zero on transmission and SHOULD be ignored on receipt.

TLVs

Zero, one, or more TLVs may be present. Each TLV is encoded as follows:

Type (16 bits)

The identifier of the TLV. Two type values are defined in this document:

- 1 IGP Instance Identifier TLV
- 2 Unnumbered Component Link Identifier TLV
- 3 IPv4 Numbered Component Link Identifier TLV
- 4 IPv6 Numbered Component Link Identifier TLV

Length (16 bits)

Indicates the total length of the TLV in octets, i.e., 4 + the length of the value field in octets. A value field whose length is not a multiple of four MUST be zero-padded so that the TLV is four-octet aligned.

Value

The data for the TLV padded as described above.

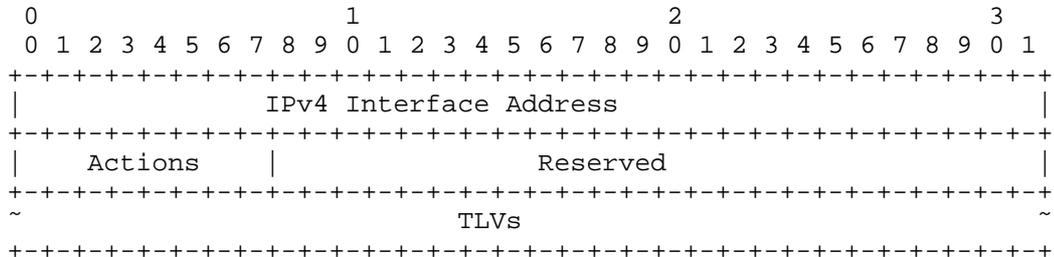
If this object is carried in a Path message, it is known as the "Forward Interface ID" for the LSP that is being set up. On a Resv message, the object is known as the "Reverse Interface ID" for the LSP.

3.1.3. IPv4 Numbered Links with Action Identification

A new C-Type variant of the LSP_TUNNEL_INTERFACE_ID object is defined to carry an IPv4 numbered interface address.

The format of the object is as shown below.

C-NUM = 193, C-Type = 2



IPv4 Interface Address

The address assigned to the interface that the sender applies to this LSP.

Actions

See Section 3.1.2.

Reserved

See Section 3.1.2.

TLVs

See Section 3.1.2.

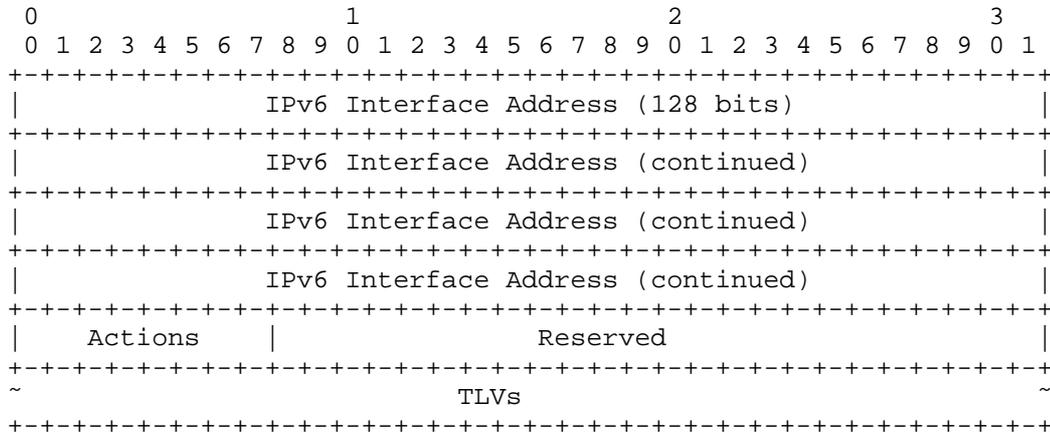
If this object is carried in a Path message, it is known as the "Forward Interface ID" for the LSP that is being set up. On a Resv message, the object is known as the "Reverse Interface ID" for the LSP.

3.1.4. IPv6 Numbered Links with Action Identification

A new C-Type variant of the LSP_TUNNEL_INTERFACE_ID object is defined to carry an IPv6 numbered interface address.

The format of the object is as shown below.

C-NUM = 193, C-Type = 3



IPv6 Interface Address

The address assigned to the interface the sender applies to this LSP.

Actions

See Section 3.1.2.

Reserved

See Section 3.1.2.

TLVs

See Section 3.1.2.

If this object is carried in a Path message, it is known as the "Forward Interface ID" for the LSP that is being set up. On a Resv message, the object is known as the "Reverse Interface ID" for the LSP.

3.2. Target IGP Identification TLV

If the LSP being set up is to be advertised as a link, the egress needs to know which instance of the IGP it should use to make the advertisement. The default in [RFC4206] and [RFC3477] is that the LSP is advertised as an FA, that is, in the same instance of the IGP as was used to advertise the TE links that the LSP traverses.

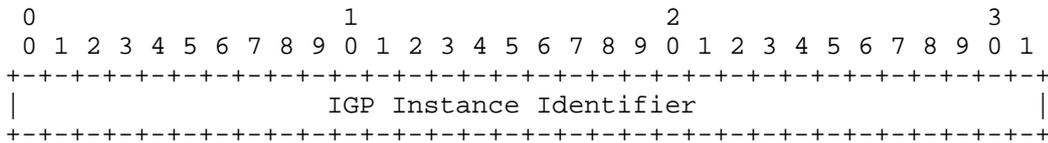
In order to facilitate an indication from the ingress to the egress of which IGP instance to use, the IGP Identification TLV is defined for inclusion in the new variants of the LSP_TUNNEL_INTERFACE_ID object defined in this document.

The TLV has meaning only in a Path message. It SHOULD NOT be included in the LSP_TUNNEL_INTERFACE_ID object in a Resv message and MUST be ignored if found.

If the P-flag in the Actions field of the LSP_TUNNEL_INTERFACE_ID object in a Path message is clear (i.e., zero), this TLV indicates the IGP instance to use for the advertisement. If the TLV is absent, the same instance of the IGP should be used as is used to advertise the TE links that the LSP traverses. Thus, for an FA, the TLV can be omitted; alternatively, the IGP Instance TLV may be present and identify the IGP instance or carry the reserved value 0xffffffff.

If the P-flag in the Actions field in the LSP_TUNNEL_INTERFACE_ID object in a Resv message is set (i.e., one) indicating that the LSP is not to be advertised as a link, this TLV SHOULD NOT be present and MUST be ignored if encountered.

The TLV is formatted as described in Section 3.1.2. The Type field has the value 1, and the Value field has the following content:



IGP Instance Identifier

A 32-bit identifier assigned to each of the IGP instances within a network, such that ingress and egress LSRs have the same understanding of these numbers. This is a management configuration exercise outside the scope of this document.

Note that the IGP Instance Identifier might be mapped from an instance identifier used in the IGP itself (such as section 2.4 of [RFC5340] for OSPFv3, or [OSPFv2-MI] for OSPFv2) as a matter of network policy. See [OSPF-TI] for further discussion of this topic in OSPF, and [ISIS-GENAP] for IS-IS.

The value 0xffffffff is reserved to mean that the LSP is to be advertised in the same instance of the IGP as was used to advertise the TE links that the LSP traverses.

3.3. Component Link Identification TLV

If the LSP that is being set up is to be used as a component link in a link bundle [RFC4201], it is necessary to indicate both the identity of the component link and the identity of the link bundle. Furthermore, it is necessary to indicate how the link bundle (that may be automatically created by the establishment of this LSP) is to be used and advertised.

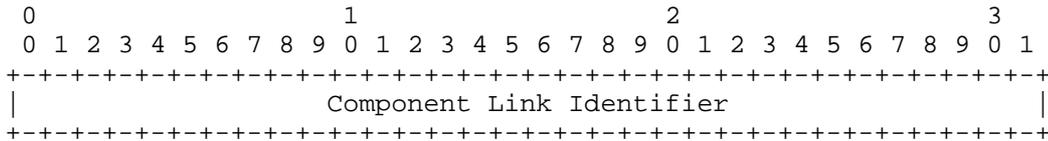
If the B-flag in the Actions field of the LSP_TUNNEL_INTERFACE_ID object is set, the other fields of the object apply to the link bundle itself. That is, the interface identifiers (numbered or unnumbered) and the other flags in the Actions field apply to the link bundle and not to the component link that the LSP will form.

Furthermore, the IGP Instance Identifier TLV (if present) also applies to the link bundle and not to the component link.

In order to exchange the identity of the component link, the Component Link Identifier TLVs are introduced as set out in the next sections. If the B-flag is set in the Actions field of the LSP_TUNNEL_INTERFACE_ID object in the Path message, exactly one of these TLVs MUST be present in the LSP_TUNNEL_INTERFACE_ID object in both the Path and Resv objects.

3.3.1. Unnumbered Component Link Identification

If the component link is to be unnumbered, the Unnumbered Component Link Identifier TLV is used. The TLV is formatted as described in Section 3.1.2. The Type field has the value 2, and the Value field has the following content:

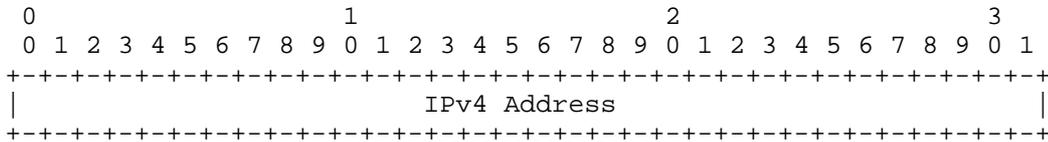


Component Link Identifier

Unnumbered identifier that is assigned to this component link within the bundle [RFC4201].

3.3.2. IPv4 Numbered Component Link Identification

If the component link is identified with an IPv4 address, the IPv4 Numbered Component Link Identifier TLV is used. The TLV is formatted as described in Section 3.1.2. The Type field has the value 3, and the Value field has the following content:

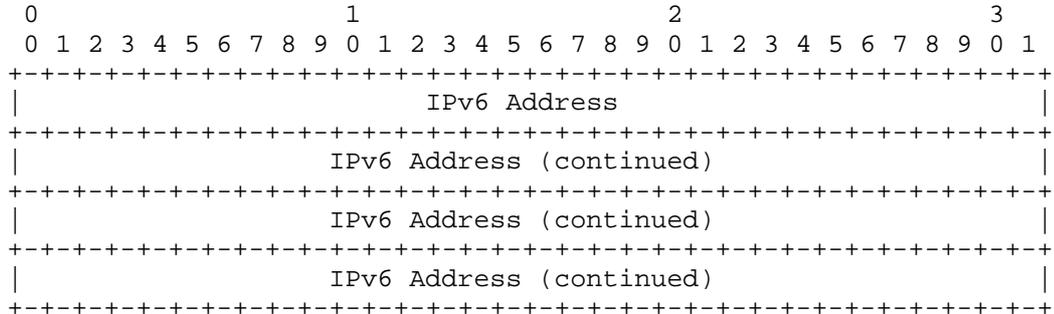


IPv4 Address

The IPv4 address that is assigned to this component link within the bundle.

3.3.3. IPv6 Numbered Component Link Identification

If the component link is identified with an IPv6 address, the IPv6 Numbered Component Link Identifier TLV is used. The TLV is formatted as described in Section 3.1.2. The Type field has the value 4, and the Value field has the following content:



IPv6 Address

The IPv6 address that is assigned to this component link within the bundle.

3.4. Link State Advertisement

The ingress and egress of an LSP that is set up using the LSP_TUNNEL_INTERFACE_ID object MUST advertise the LSP as agreed in the parameters of the object.

Where a TE link is created from the LSP, the TE link SHOULD inherit the TE properties of the LSP as described in [RFC5212], but this process is subject to local and network-wide policy.

It is possible that an LSP will be used to offer capacity and connectivity to multiple other networks. In this case, multiple instances of the LSP_TUNNEL_INTERFACE_ID object are permitted in the same Path and Resv messages. Each instance MUST have a different IGP Instance Identifier.

Note, however, that a Path or Resv message MUST NOT contain more than one instance of the LSP_TUNNEL_INTERFACE_ID object with C-Type 1, and if such an object is present, all other instances of the LSP_TUNNEL_INTERFACE_ID object MUST include an IGP Instance Identifier TLV with IGP Instance Identifier set to a value that identifies some other IGP instance (in particular, not the value 0xffffffff).

If the link created from an LSP is advertised in the same IGP instance as was used to advertise the TE links that the LSP traverses, the addresses for the new link and for the links from which it is built MUST come from the same address space.

If the link is advertised into another IGP instance, the addresses MAY be drawn from overlapping address spaces such that some addresses have different meanings in each IGP instance.

In the IGP, the TE Router ID of the ingress LSR is taken from the Tunnel Sender Address in the Sender Template object signaled in the Path message. It is assumed that the ingress LSR knows the TE Router ID of the egress LSR since it has chosen to establish an LSP to that LSR and plans to use the LSP as a TE link.

The link interface addresses or link interface identifiers for the forward and reverse direction links are taken from the LSP_TUNNEL_INTERFACE_ID object on the Path and Resv messages, respectively.

When an LSP is torn down through explicit action (a PathTear message or a PathErr message with the Path_State_Removed flag set), the ingress and egress LSRs SHOULD withdraw the advertisement of any link that the LSP created and that was previously advertised. The link state advertisement MAY be retained as a virtual link in another layer network according to network-wide policy [PCE-LAYER].

3.5. Message Formats

[RFC3477] does not state where in the Path message or Resv message the LSP_TUNNEL_INTERFACE_ID object should be placed.

It is RECOMMENDED that new implementations place the LSP_TUNNEL_INTERFACE_ID objects in the Path message immediately after the SENDER_TSPEC object, and in the Resv message immediately after the FILTER_SPEC object.

All implementations SHOULD be able to handle received messages with objects in any order, as described in [RFC3209].

3.6. Error Cases and Non-Acceptance

Error cases and non-acceptance of new object variants caused by back-level implementations are discussed in Section 3.7.

An egress LSR that receives an LSP_TUNNEL_INTERFACE_ID object may have cause to reject the parameters carried in the object for a number of reasons as set out below. In all cases, the egress SHOULD

respond with a PathErr message with the error code as indicated in the list below. In most cases, the error will arise during LSP setup, no Resv state will exist, and the PathErr will cause Path state to be removed. Where the error arises after the LSP has been successfully set up, the PathErr SHOULD be sent with the Path_State_Removed flag [RFC3473] clear so that the LSP remains operational.

The error cases identified are as follows and are reported using the new error code 'LSP Hierarchy Issue' (code 38) and the error values listed below.

Error code	Error value	Error-case
38	1	Link advertisement not supported
38	2	Link advertisement not allowed by policy
38	3	TE link creation not supported
38	4	TE link creation not allowed by policy
38	5	Routing adjacency creation not supported
38	6	Routing adjacency creation not allowed by policy
38	7	Bundle creation not supported
38	8	Bundle creation not allowed by policy
38	9	Hierarchical LSP not supported
38	10	LSP stitching not supported
38	11	Link address type or family not supported
38	12	IGP instance unknown
38	13	IGP instance advertisement not allowed by policy
38	14	Component link identifier not valid
38	15	Unsupported component link identifier address family

When an ingress LSR receives an LSP_TUNNEL_INTERFACE_ID object on a Resv message, it may need to reject it because of the setting of certain parameters in the object. Since these reasons all represent errors rather than mismatches of negotiable parameters, the ingress SHOULD respond with a PathTear to remove the LSP. The ingress MAY use a ResvErr with one of the following error codes, allowing the egress the option to correct the error in a new Resv message, or to tear down the LSP with a PathErr with the Path_State_Removed flag set. An ingress that uses the ResvErr MUST take precautions against a protocol loop where the egress responds with the same LSP_TUNNEL_INTERFACE_ID object with the same (or different) issues.

Error code	Error value	Error-case
38	11	Link address type or family not supported
38	14	Component link identifier not valid
38	15	Unsupported component link identifier address family
38	16	Component link identifier missing

3.7. Backward Compatibility

The LSP_TUNNEL_INTERFACE_ID object defined in [RFC3477] has a class number of 193. According to [RFC2205], this means that a node that does not understand the object SHOULD ignore the object but forward it, unexamined and unmodified. Thus, there are no issues with transit LSRs supporting the pre-existing or new Class Types of this object.

A back-level ingress node will behave as follows:

- It will not issue Path messages containing LSP_TUNNEL_INTERFACE_ID objects with the new Class Types defined in this document.
- It will reject Resv messages containing LSP_TUNNEL_INTERFACE_ID objects with the new Class Types defined in this document as described in [RFC2205]. In any case, such a situation would represent an error by the egress.
- It will continue to use the LSP_TUNNEL_INTERFACE_ID object with Class Type 1 as defined in [RFC3477]. This behavior is supported by back-level egresses and by egresses conforming to this document.
- According to an informal survey, there is no significant deployment of numbered FA establishment following the procedures defined in [RFC4206] and set out in Section 1.3.6 of this document. It is therefore safe to assume that back-level ingress LSRs will not use this mechanism.

A back-level egress node will behave as follows:

- It will continue to support the LSP_TUNNEL_INTERFACE_ID object with Class Type 1, as defined in [RFC3477], if issued by an ingress.
- It will reject a Path message that carries an LSP_TUNNEL_INTERFACE_ID object with any of the new Class Types defined in this document using the procedures of [RFC2205]. This will inform the ingress that the egress is a back-level LSR.

- It will not expect to use the procedures for numbered FA establishment defined in [RFC4206] and set out in Section 1.3.6 of this document.

In summary, the new mechanisms defined in this document do not impact the method to exchange unnumbered FA information described in [RFC3477]. That mechanism can be safely used in combination with the new mechanisms described here and is functionally equivalent to using the new C-Type indicating an unnumbered link with target IGP instance identifier with the Target IGP Instance value set to 0xffffffff.

The mechanisms in this document obsolete the method to exchange the numbered FA information described in [RFC4206] as described in Section 1.4.6.

4. Security Considerations

[RFC3477] points out that one can argue that the use of the extra interface identifier that it provides could make an RSVP message harder to spoof. In that respect, the minor extensions to the protocol made in this document do not constitute an additional security risk, but could also be said to improve security.

It should be noted that the ability of an ingress LSR to request that an egress LSR advertise an LSP as a TE link MUST be subject to appropriate policy checks at the egress LSR. That is, the egress LSR MUST NOT automatically accept the word of the ingress unless it is configured with such a policy.

Further details of security for MPLS-TE and GMPLS can be found in [RFC5920].

5. IANA Considerations

5.1. New Class Types

IANA maintains a registry of RSVP parameters called "Resource Reservation Protocol (RSVP) Parameters" with a sub-registry called "Class Names, Class Numbers, and Class Types". There is an entry in this registry for the LSP_TUNNEL_INTERFACE_ID object defined in [RFC3477] with one Class Type defined.

IANA has allocated three new Class Types for this object as defined in Sections 3.1.2, 3.1.3, and 3.1.4 as follows:

C-Type	Meaning	Reference
2	IPv4 interface identifier with target	[RFC6107]
3	IPv6 interface identifier with target	[RFC6107]
4	Unnumbered interface with target	[RFC6107]

5.2. Hierarchy Actions

Section 3.1.2 defines an 8-bit flags field in the LSP_TUNNEL_INTERFACE_ID object. IANA has created a new sub-registry of the "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Parameters" registry called the "Hierarchy Actions" sub-registry as follows:

Registry Name: Hierarchy Actions
 Reference: [RFC6107]
 Registration Procedures: Standards Action

Registry:

Bit Number	Hex Value	Name	Reference
0-2		Unassigned	
3	0x10	Hierarchy/stitching (H)	[RFC6107]
4	0x08	Bundle (B)	[RFC6107]
5	0x04	Routing adjacency (R)	[RFC6107]
6	0x02	TE link (T)	[RFC6107]
7	0x01	Private (P)	[RFC6107]

5.3. New Error Codes and Error Values

IANA maintains a registry of RSVP error codes and error values as the "Error Codes and Globally-Defined Error Value Sub-Codes" sub-registry of the "Resource Reservation Protocol (RSVP) Parameters" registry.

IANA has defined a new error code with value 38 as below (see also Section 3.6).

Error Code	Meaning	Reference
38	LSP Hierarchy Issue	[RFC6107]

IANA has listed the following error values for this error code (see also Section 3.6).

This Error Code has the following globally-defined Error Value sub-codes:

1 = Link advertisement not supported	[RFC6107]
2 = Link advertisement not allowed by policy	[RFC6107]
3 = TE link creation not supported	[RFC6107]
4 = TE link creation not allowed by policy	[RFC6107]
5 = Routing adjacency creation not supported	[RFC6107]
6 = Routing adjacency creation not allowed by policy	[RFC6107]
7 = Bundle creation not supported	[RFC6107]
8 = Bundle creation not allowed by policy	[RFC6107]
9 = Hierarchical LSP not supported	[RFC6107]
10 = LSP stitching not supported	[RFC6107]
11 = Link address type or family not supported	[RFC6107]
12 = IGP instance unknown	[RFC6107]
13 = IGP instance advertisement not allowed by policy	[RFC6107]
14 = Component link identifier not valid	[RFC6107]
15 = Unsupported component link identifier address family	[RFC6107]
16 = Component link identifier missing	[RFC6107]

6. Acknowledgements

The authors would like to thank Lou Berger, Deborah Brungard, John Drake, Yakov Rekhter, Igor Bryskin, and Lucy Yong for their valuable discussions and comments.

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.
- [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.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, January 2003.
- [RFC4201] Kompella, K., Rekhter, Y., and L. Berger, "Link Bundling in MPLS Traffic Engineering (TE)", RFC 4201, October 2005.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.
- [RFC5150] Ayyangar, A., Kompella, K., Vasseur, JP., and A. Farrel, "Label Switched Path Stitching with Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)", RFC 5150, February 2008.

7.2. Informative References

- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.
- [RFC1256] Deering, S., Ed., "ICMP Router Discovery Messages", RFC 1256, September 1991.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.
- [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.

- [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.
- [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.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, October 2008.
- [RFC5329] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed., "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, September 2008.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
- [ISIS-GENAP] Ginsberg, L., Previdi, S., and M. Shand, "Advertising Generic Information in IS-IS", Work in Progress, November 2010.
- [ISIS-IPV6-TE] Harrison, J., Berger, J., and M. Bartlett, "IPv6 Traffic Engineering in IS-IS", Work in Progress, September 2010.
- [OSPF-TI] Lindem, A., Roy, A., and S. Mirtorabi, "OSPF Transport Instance Extensions", Work in Progress, October 2010.
- [OSPFv2-MI] Lindem, A., Roy, A., and S. Mirtorabi, "OSPF Multi-Instance Extensions", Work in Progress, October 2010.
- [PCE-LAYER] Takeda, T., Ed., and A. Farrel, "PCC-PCE Communication and PCE Discovery Requirements for Inter-Layer Traffic Engineering", Work in Progress, December 2010.

Authors' Addresses

Richard Rabbat
Google Inc.
1600 Amphitheatre Pkwy
Mountain View, CA 94043
United States of America
EMail: rabbat@alum.mit.edu

Arthi Ayyangar
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
United States of America
EMail: arthi@juniper.net

Zafar Ali
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, Ontario, K2K 3E8
Canada
EMail: zali@cisco.com

Editors' Addresses

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

Adrian Farrel
Old Dog Consulting
EMail: adrian@olddog.co.uk

