

Open Shortest Path First IGP
Internet-Draft
Intended status: Standards Track
Expires: November 22, 2013

H. Gredler, Ed.
Juniper Networks, Inc.
S. Amante
Level 3 Communications, Inc.
T. Scholl
Amazon
L. Jalil
Verizon
May 21, 2013

Advertising MPLS labels in OSPF
draft-gredler-ospf-label-advertisement-03

Abstract

Historically MPLS label distribution was driven by protocols like LDP, RSVP and LBGP. All of those protocols are session oriented. In order to obtain label binding for a given destination FEC from a given router one needs first to establish an LDP/RSVP/LBGP session with that router.

Advertising MPLS labels in IGPs
[I-D.gredler-rtgwg-igp-label-advertisement] describes several use cases where utilizing the flooding machinery of link-state protocols for MPLS label distribution allows to obtain the binding without requiring to establish an LDP/RSVP/LBGP session with that router.

This document describes the protocol extension to distribute MPLS label bindings by the OSPFv2 and OSPFv3 protocol.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months

and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on November 22, 2013.

Copyright Notice

Copyright (c) 2013 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	5
2.	Motivation, Rationale and Applicability	5
2.1.	Issue: Bi-directionality semantics	6
2.2.	Issue: IP path semantics	6
2.3.	Issue: Lack of 'path' notion	6
2.4.	Motivation	7
3.	OSPF MPLS LSA Format	7
3.1.	Common LSA Type	7
3.2.	OSPFv2 LSA ID	7
3.3.	OSPFv2 LSA Format Overview	8
3.4.	OSPFv3 LSA ID	8
3.5.	OSPFv3 LSA Format Overview	9
3.6.	TLV Header	9
4.	LSA payload details	10
4.1.	ERO TLVs	10
4.1.1.	IPv4 Prefix ERO TLV	10
4.1.2.	IPv6 Prefix ERO TLV	11
4.1.3.	Unnumbered Interface ID ERO TLV	12
4.1.4.	IPv4 Prefix Bypass ERO TLV	13
4.1.5.	IPv6 Prefix Bypass ERO TLV	13
4.1.6.	Unnumbered Interface ID Bypass ERO TLV	14
4.2.	Flags TLV	15
4.3.	All Router Block TLV	15
4.4.	All Router ID IPv4 Map TLV	17
4.5.	All Router ID IPv6 Map TLV	18
5.	Advertising Label Examples	18
5.1.	Sample Topology	18
5.1.1.	Transport IP addresses and router-IDs	19
5.1.2.	Link IP addresses	19
5.2.	One-hop LSP to an adjacent Router	20
5.3.	One-hop LSP to an adjacent Router using a specific link	20
5.4.	One-hop LSP to an adjacent external Router	20
5.5.	Advertisement of an RSVP LSP	21
5.6.	Advertisement of an LDP LSP	21
5.7.	Interarea advertisement of diverse paths	21
5.8.	Advertisement of SPT labels using 'All Router Block' TLV	22
5.9.	Expansion of an 'All Router Block' TLV	23
6.	Inter Area Protocol Procedures	24
6.1.	Applicability	24
6.2.	Data plane operations	24
6.3.	Control plane operations	24
6.3.1.	MPLS Label operations	24
6.3.2.	MPLS Label Block operations	25
7.	Acknowledgements	25
8.	IANA Considerations	25

9. Security Considerations 26
10. References 26
 10.1. Normative References 26
 10.2. Informative References 27
Authors' Addresses 27

1. Introduction

MPLS label allocations are predominantly distributed by using the LDP [RFC5036], RSVP [RFC5151] or labeled BGP [RFC3107] protocol. All of those protocols have in common that they are session oriented, which means that in order to obtain label binding for a given destination FEC from a given router one needs first to establish a direct control plane (LDP/RSVP/LBGP) session with that router.

There are a couple of practical use cases [I-D.gredler-rtgwg-igp-label-advertisement] where the consumer of a MPLS label binding may not be adjacent to the router that performs the binding. Bringing up an explicit session using the existing label distribution protocols between the non-adjacent router that binds the label and the router that acts as a consumer of this binding is the existing remedy for this dilemma.

This document describes an OSPFv2 and OSPFv3 protocol extension which allows routers to advertise MPLS label bindings within and beyond an IGP domain, and controlling inter-area distribution.

2. Motivation, Rationale and Applicability

Distributing MPLS labels in an IGP (IS-IS) has been described in Segment Routing [I-D.previdi-filsfils-isis-segment-routing]. The authors propose to re-use existing traffic-engineering extensions for carrying the label information. While retrofitting existing protocol machinery for new purposes is generally a good thing, Segment Routing [I-D.previdi-filsfils-isis-segment-routing] falls short of addressing some use-cases defined in [I-D.gredler-rtgwg-igp-label-advertisement].

The dominant issue around re-using traffic-engineering extensions is that both have existing protocol semantics, which might not be applicable to advertising MPLS label switched paths in a generic fashion. These are specifically:

- o Bi-directionality semantics
- o IP path semantics
- o Lack of 'path' notion

2.1. Issue: Bi-directionality semantics

'Bi-directionality semantics', affects the complexity around advertisement of unidirectional LSPs. Label advertisement of per-link labels or 'Adj-SIDs' [I-D.previdi-filsfils-isis-segment-routing] is done by attaching label information to adjacency advertisement TLVs. Usually implementations need to have an adjacency in 'Up' state prior to advertising this adjacency as TE-Link in its Link State advertisement. In order to advertise a per-link LSP an implementation first needs to have an adjacency, which only transitions to 'Up' state after passing the 3-way check. This implies bi-directionality. If an implementation wants to advertise per-link MPLS LSPs to e.g. outside the IGP domain then it would need to fake-up an adjacency. Changing existing IGP Adjacency code to support such cases defeats the purpose of re-using existing functionality as there is not much common functionality to be shared.

2.2. Issue: IP path semantics

LSPs pointing to a Node are advertised as 'Node-SIDs' [I-D.previdi-filsfils-isis-segment-routing] using IP Prefix containers. That means that in order to advertise a MPLS LSP, one is inheriting the semantics of advertising an IP path. Consider router A has got existing LSPs to its entire one-hop neighborhood and is re-advertising those LSPs using IP reachability semantics. Now we have two exact matching IP advertisements. One from the owning router (router B) which advertises its stable transport loopback address and another one from router A re-advertising a LSP path to router B. Existing routing software may get confused now as the 'stable transport' address shows up from multiple places in the network and more worse the IP forwarding path for control-plane protocols may get mingled with the MPLS data plane.

2.3. Issue: Lack of 'path' notion

Both existing traffic-engineering extension containers have limited semantics describing MPLS label-switched paths in the sense of a 'path'. Both encoding formats allow to express a pointer to some specific router, but not to describe a MPLS label switched path containing all of its path segments. [I-D.previdi-filsfils-isis-segment-routing] allows to define 'Forwarding Adjacencies' as per [RFC4206]. The way to describe a path of a given forwarding adjacency is to enlist a list of "Segment IDs". That implies that nodes which do not yet participate in 'Segment routing' or are outside of a 'Segment routing' domain can not be expressed using those path semantics.

A protocol for advertising MPLS label switched paths, should be

generic enough to express paths sourced by existing MPLS LSPs, such that ingress routers can flexibly combine them according to application needs.

2.4. Motivation

IGP advertisement of MPLS label switched paths requires a new set of protocol semantics (undirectional paradigm, path paradigm), which hardly can be expressed using the existing OSPF and OSPF-TE protocol semantics. This document describes protocol extensions which allows generic advertisement of MPLS label bindings in the OSPFv2 and OSPFv3 protocol.

The Protocol extensions described in this document are equally applicable to IPv4 and IPv6 carried over MPLS. Furthermore the proposed use of distributing MPLS Labels using IGP protocols adheres to the architectural principles laid out in [RFC3031].

3. OSPF MPLS LSA Format

3.1. Common LSA Type

One new LSA is defined, the MPLS Label LSA. This LSA advertises MPLS labels along with their path information. The LSA contains more specific information encoded in TLVs. Those TLV extensions are shared between the OSPFv2 and OPSFv3 protocols.

3.2. OSPFv2 LSA ID

The LSA ID of an Opaque LSA is defined as having eight bits of type data and 24 bits of type-specific data. The MPLS Label LSA uses type 149. The remaining 24 bits are 4 zero bits followed by the MPLS Label or MPLS Label Base value as follows:

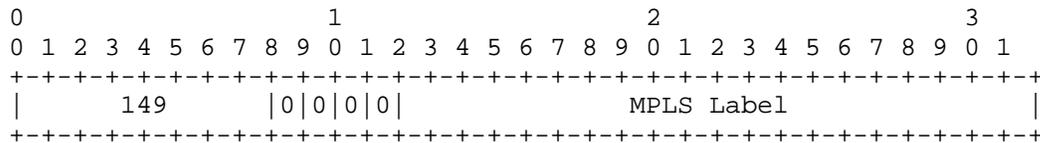


Figure 1: OSPFv2 MPLS Label LSA-ID format

The 'MPLS Label' field holds the 20 Bit MPLS label for MPLS Label base. Therefore a maximum of 2^20 MPLS Label LSAs may be sourced by a single system.

3.5. OSPFv3 LSA Format Overview

The MPLS Label LSA for OSPFv3 starts with the standard OSPFv3 LSA header:

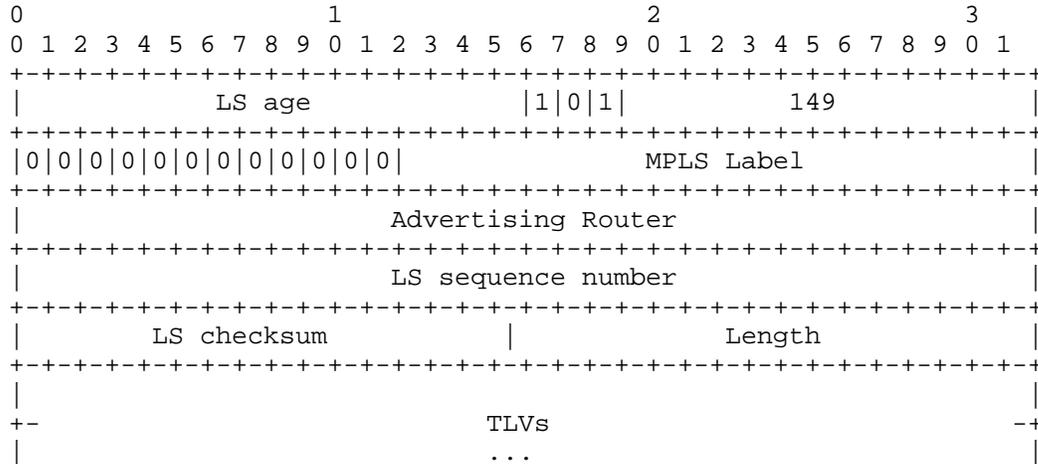


Figure 4: OSPFv3 MPLS Label LSA format

The OSPFv3 'U' Bit will always be set such that routers which do not understand the new MPLS Label LSA will store and forward it further.

In analogy to the OSPFv2 opaque LSA 10 the flooding scope will be set to 'Area scoping'.

3.6. TLV Header

The LSA payload consists of one or more nested Type/Length/Value (TLV) triplets for extensibility. The format of each TLV is:

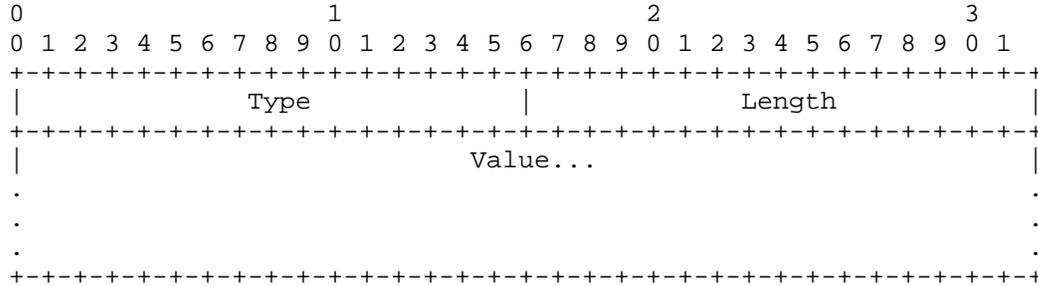


Figure 5: TLV format

The Length field defines the length of the value portion in octets (thus a TLV with no value portion would have a length of zero). The TLV is padded to four-octet alignment; padding is not included in the length field (so a three octet value would have a length of three, but the total size of the TLV would be eight octets). Nested TLVs are also 32-bit aligned. Unrecognized types are ignored.

This memo defines TLV Types 1 through 8. See the IANA Considerations section for allocation of new Types.

4. LSA payload details

The MPLS Label LSA may be originated by any Traffic Engineering [RFC3630] capable router in an OSPF domain. The router may advertise a single label binding or a block of label bindings. For single label binding advertisement a router needs to provide at least a single 'nexthop style' anchor. The protocol supports more than one 'nexthop style' anchor to be attached to a Label binding, which results into a simple path description language. In analogy to RSVP the terminology for this is called an 'Explicit Route Object' (ERO). Since ERO style path notation allows to anchor label bindings to both link and node IP addresses any label switched path, can be described. Furthermore also Label Bindings from other protocols can get easily re-advertised.

An LSA contains one or more TLVs, describing properties of the advertised MPLS label.

The following TLV extensions may be shared in both OSPV2 and OSPFv3. Passing an IP address of the other address family (IPv4 in OPSFv3 or IPv6 in OSPFv2) is possible as the information carried are related describing the hops along a path. The receiver of this information is a protocol agnostic path computation module.

4.1. ERO TLVs

All 'ERO' information represents an ordered set which describes the segments of a label-switched path. The last ERO TLV describes the segment closest to the egress point of the LSP. Contrary the first ERO TLV describes the first segment of a label switched path. If a router extends or stitches a label switched path it MUST prepend the new segments path information to the ERO list.

4.1.1. IPv4 Prefix ERO TLV

The IPv4 ERO TLV (Type 1) describes a path segment using IPv4 Prefix style of encoding. Its appearance and semantics have been borrowed

from Section 4.3.3.2 [RFC3209].

the 'IPv4 Address' is treated as a prefix based on the prefix length value below. Bits beyond the prefix are ignored on receipt and SHOULD be set to zero on transmission.

The 'Prefix Length' field contains the length of the prefix in bits.

The 'L' bit in the TLV is a one-bit attribute. If the L bit is set, then the value of the attribute is 'loose.' Otherwise, the value of the attribute is 'strict.'

The 'Reserved' bits are for future use. They should be zero on transmission and ignored on receipt.

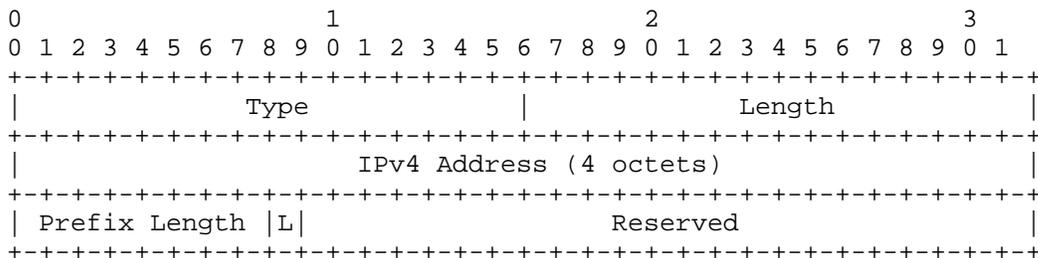


Figure 6: IPv4 Prefix ERO TLV format

4.1.2. IPv6 Prefix ERO TLV

The IPv6 ERO TLV (Type 2) describes a path segment using IPv6 Prefix style of encoding. Its appearance and semantics have been borrowed from Section 4.3.3.2 [RFC3209].

the 'IPv6 Address' is treated as a prefix based on the prefix length value below. Bits beyond the prefix are ignored on receipt and SHOULD be set to zero on transmission.

The 'Prefix Length' field contains the length of the prefix in bits.

The 'L' bit in the TLV is a one-bit attribute. If the L bit is set, then the value of the attribute is 'loose.' Otherwise, the value of the attribute is 'strict.'

The 'Reserved' bits are for future use. They should be zero on transmission and ignored on receipt.

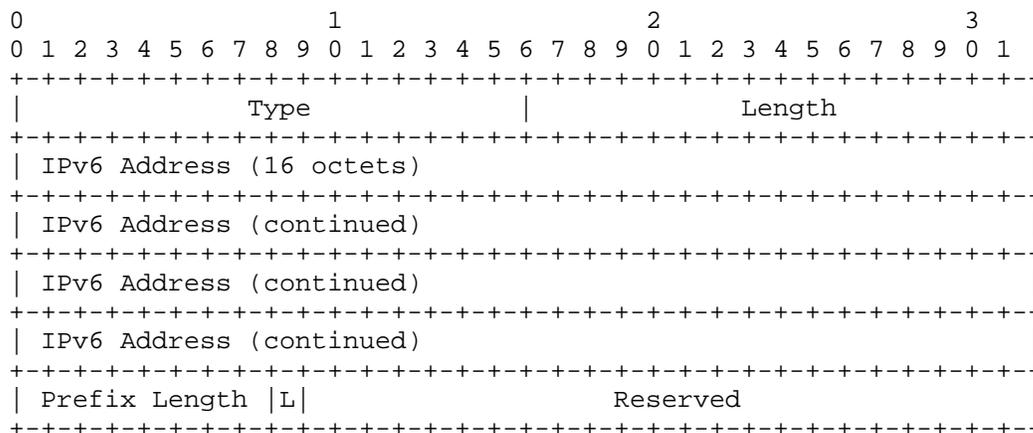


Figure 7: IPv6 Prefix ERO TLV format

4.1.3. Unnumbered Interface ID ERO TLV

The appearance and semantics of the 'Unnumbered Interface ID' have been borrowed from Section 4 [RFC3477].

The Unnumbered Interface-ID ERO TLV (Type 9) describes a path segment that spans over an unnumbered interface. Unnumbered interfaces are referenced using the interface index. Interface indices are assigned local to the router and therefore not unique within a domain. All elements in an ERO path need to be unique within a domain and hence need to be disambiguated using a domain unique Router-ID.

The 'Router-ID' field contains the router ID of the router which has assigned the 'Interface ID' field. Its purpose is to disambiguate the 'Interface ID' field from other routers in the domain.

The 'Interface ID' is the identifier assigned to the link by the router specified by the router ID.

The 'L' bit in the TLV is a one-bit attribute. If the L bit is set, then the value of the attribute is 'loose.' Otherwise, the value of the attribute is 'strict.'

The 'Reserved' bits are for future use. They should be zero on transmission and ignored on receipt.

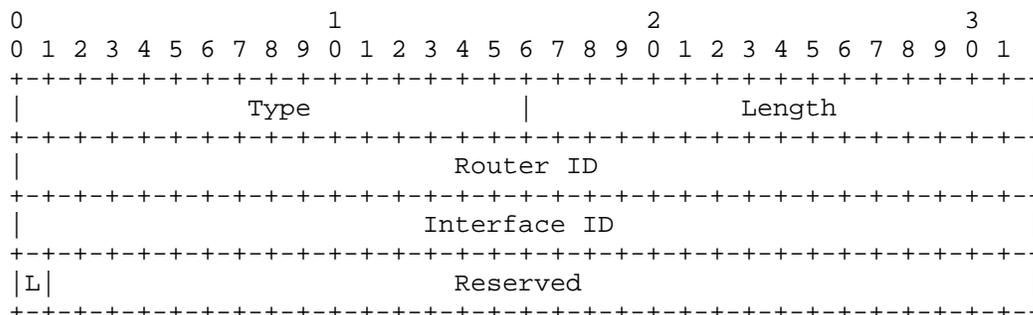


Figure 8: Unnumbered Interface ID ERO TLV format

4.1.4. IPv4 Prefix Bypass ERO TLV

The IPv4 Bypass ERO TLV (Type 3) describes a Bypass LSP path segment using IPv4 Prefix style of encoding. Its appearance and semantics have been borrowed from Section 4.3.3.2 [RFC3209].

The 'Prefix Length' field contains the length of the prefix in bits.

The 'L' bit in the TLV is a one-bit attribute. If the L bit is set, then the value of the attribute is 'loose.' Otherwise, the value of the attribute is 'strict.'

The 'Reserved' bits are for future use. They should be zero on transmission and ignored on receipt.

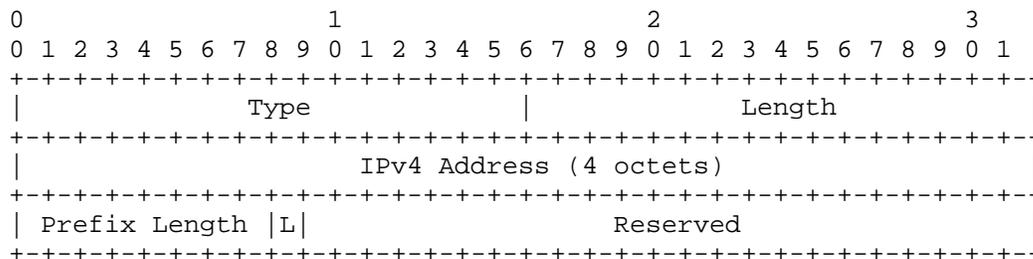


Figure 9: IPv4 Prefix Bypass ERO TLV format

4.1.5. IPv6 Prefix Bypass ERO TLV

The IPv6 ERO TLV (Type 4) describes a Bypass LSP path segment using IPv6 Prefix style of encoding. Its appearance and semantics have been borrowed from Section 4.3.3.3 [RFC3209].

The 'Prefix Length' field contains the length of the prefix in bits.

The 'L' bit in the TLV is a one-bit attribute. If the L bit is set, then the value of the attribute is 'loose.' Otherwise, the value of the attribute is 'strict.'

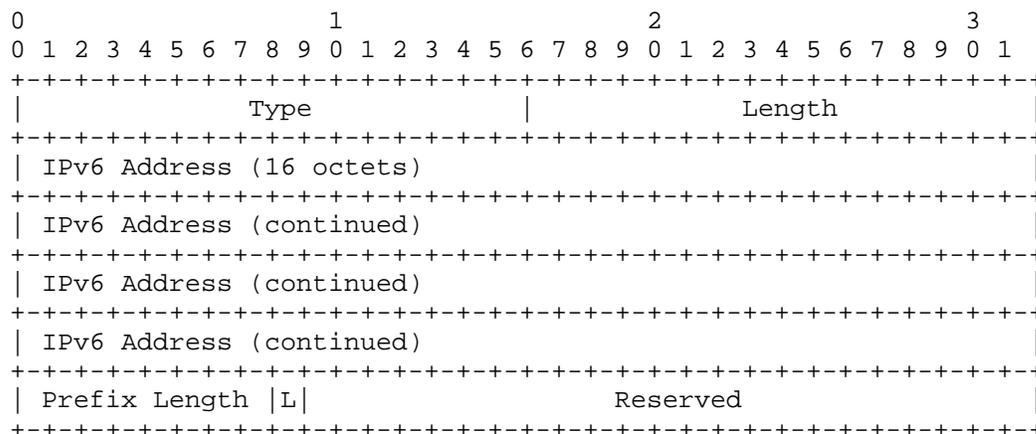


Figure 10: IPv6 Prefix Bypass ERO TLV format

4.1.6. Unnumbered Interface ID Bypass ERO TLV

The appearance and semantics of the 'Unnumbered Interface ID' have been borrowed from Section 4 [RFC3477].

The Unnumbered Interface-ID Bypass ERO TLV (Type 10) describes a Bypass path segment that spans over an unnumbered interface. Unnumbered interfaces are referenced using the interface index. Interface indices are assigned local to the router and therefore not unique within a domain. All elements in an ERO path need to be unique within a domain and hence need to be disambiguated using a domain unique Router-ID.

The 'Router-ID' field contains the router ID of the router which has assigned the 'Interface ID' field. Its purpose is to disambiguate the 'Interface ID' field from other routers in the domain.

The 'Interface ID' is the identifier assigned to the link by the router specified by the router ID.

The 'L' bit in the TLV is a one-bit attribute. If the L bit is set, then the value of the attribute is 'loose.' Otherwise, the value of the attribute is 'strict.'

The 'Reserved' bits are for future use. They should be zero on transmission and ignored on receipt.

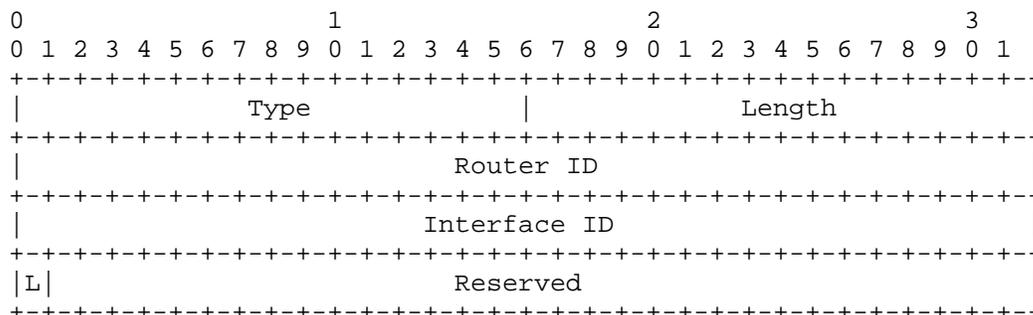


Figure 11: Unnumbered Interface ID Bypass ERO TLV format

4.2. Flags TLV

The Flags TLV (Type 5) describes Flags for further MPLS LSA treatment.

Up/Down 'U' Bit: A router may flood MPLS label information across area boundaries. In order to prevent flooding loops, a router will Set the Up/Down (U-Bit) when propagating MPLS labels from Area 0 to a non-zero Area.

The 'Reserved' bits are for future use. They should be zero on transmission and ignored on receipt.

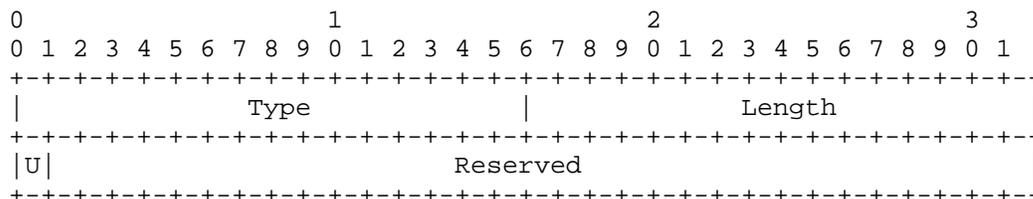


Figure 12: Flags TLV format

4.3. All Router Block TLV

The 'All Router Block' TLV (Type 6) denominates the label block size of an MPLS Label advertisement and its semantics to connect to all routers in a given OSPF domain using a local assigned [RFC3031] label range. Note that the actual mapping of a router within the label range is done using the TLVs described in Section 4.4 and Section 4.5. Since generation of an IPv4 or IPv6 Map TLV is a local policy decision, it might be the case that connectivity is provided not to 'All' but rather a subset of 'All' routers. Keeping policy decisions aside, for simplicity reasons, assume that All Routers in a

domain do generate either the 'All Router ID IPv4 Map' or 'All Router ID IPv6 Map' TLVs and therefore all routers desire construction of a Label switched path from every source router in the network. The basic concept of using label blocks to provide connectivity to a set of routers has been borrowed from [RFC4761] which allows to advertise labels from multiple end-points using a single control-plane message. The difference to [RFC4761] is that rather than advertising where a particular packet came from (=source semantics), destination semantics (where a particular packet will be going to) is advertised.

Along with each label block a router advertises one for more 'IDs'. The 'ID' must be unique within a given domain. The 'ID' serves as ordinal to determine the actual label value inside the set of all advertised label ranges of a given router. A receiving router uses the ordinal to determine the actual label value in order to construct forwarding state to a particular destination router. The 'ID' is separately advertised using the TLVs described in Section 4.4 and Section 4.5.

The ability to advertise more than one label block eases operational procedures for increasing the number of supported routers within a domain. For example consider a given domain has got support for <M> routers and runs out of ID space. It simply advertises one more label block to cover additional ordinals outside the range of the first label block. An example of label-block expansion is described in more detail in Section 5.9

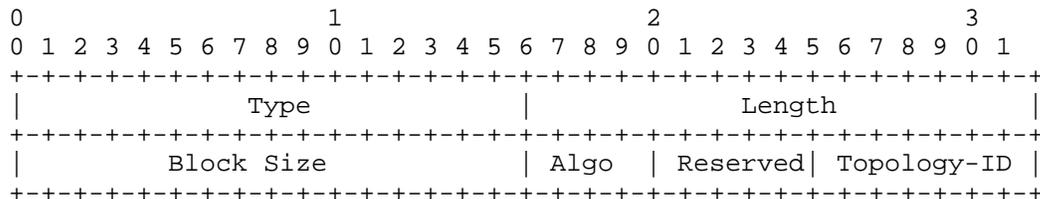


Figure 13: All Router Block TLV format

The 'Block Size' value contains the size of the label advertisement. The 'value determines the amount of reachable router endpoints within a given Label block. It MUST contain a value greater or equal than two. Note that the label base is inferred from the LSA-ID in the LSA header. For example if a router wants to advertise a label range of 5000-5099 then it would need to generate a LSA-ID of 5000 (= 0.0.19.136) and a Block Size of 100.

The 'Algo' value denominates the path computation algorithm in order to calculate the forwarding topology. The basic SPF algorithm has an assigned 'Algo' code point of zero. The purpose of the 'Algo' field

is to extend the notion of Label Block Signaling to arbitrary algorithms like for example 'MRT' ([I-D.ietf-rtgwg-mrt-frr-architecture]. Advertised Label Blocks with an unknown, unsupported or non-configured algorithm MUST be silently ignored.

The 'Reserved' bits are for future use. They should be zero on transmission and ignored on receipt.

The 'Topology-ID' field contains the Multi Topology ID ([RFC4915]) for which the advertised Label Block does apply. The basic IPv4 unicast Topology has an assigned 'Topology-ID' code point of zero. Advertised Label Blocks with an unknown, unsupported or non-configured Topology-ID MUST be silently ignored.

An LSA containing the 'All Router Block' TLV MUST only contain the Flags TLV (Section 4.2, the 'All Router IPv4 Map' TLV (Section 4.4) or the 'All Router IPv6 Map' TLV (Section 4.5).

4.4. All Router ID IPv4 Map TLV

The 'All Router ID IPv4 Map' TLV (Type 7) maps an 'ID' to a given stable transport IPv4 address. Its purpose is to associate a given transport IPv4 IP address to the ordinal inside a label range as described in Section 4.3.

A router MAY advertise more than one 'ID' to 'IPv4 address' mapping pair, in case it has more than one stable transport IPv4 address.

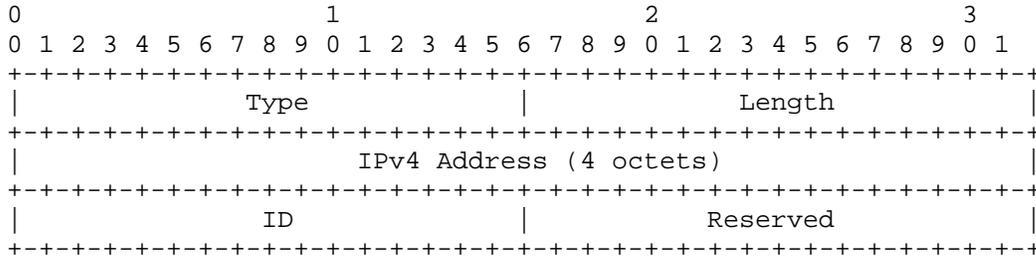


Figure 14: All Router ID IPv4 Map TLV format

The 'IPv4 address' contains stable IPv4 transport address of a given router.

The 'ID' contains the ordinal value of an advertising router inside the set of all advertised label blocks of a given router.

The 'Reserved' bits are for future use. They should be zero on

transmission and ignored on receipt.

4.5. All Router ID IPv6 Map TLV

The 'All Router ID IPv6 Map' TLV (Type 8) maps an 'ID' to a given stable transport IPv6 address. Its purpose is to associate a given transport IPv6 IP address to the ordinal inside a label range as described in Section 4.3.

A router MAY advertise more than one 'ID' to 'IPv6 address' mapping pair, in case it has more than one stable transport IPv6 address.

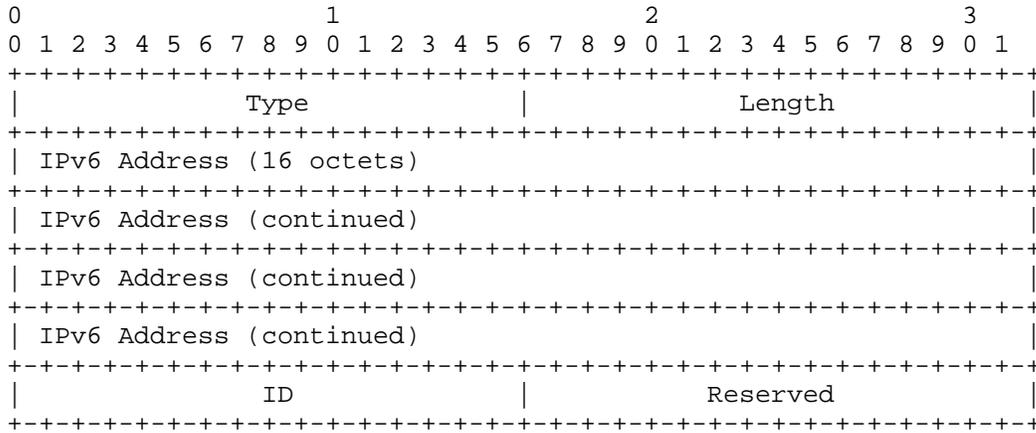


Figure 15: All Router ID IPv6 Map TLV format

The 'IPv6 address' contains the stable IPv6 transport address of a given router.

The 'ID' contains the ordinal value of an advertising router inside the set of all advertised label blocks of a given router.

The 'Reserved' bits are for future use. They should be zero on transmission and ignored on receipt.

5. Advertising Label Examples

5.1. Sample Topology

The following topology (Figure 16) and IP addresses shall be used throughout the Label advertisement examples.

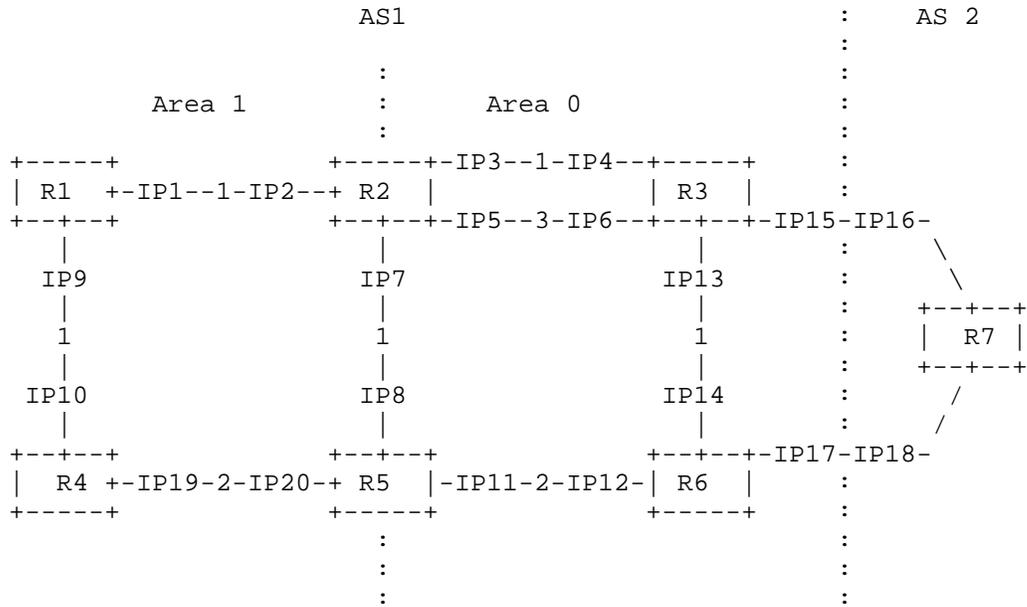


Figure 16: Sample Topology

5.1.1. Transport IP addresses and router-IDs

- o R1: 192.168.1.1
- o R2: 192.168.1.2
- o R3: 192.168.1.3
- o R4: 192.168.1.4
- o R5: 192.168.1.5
- o R6: 192.168.1.6
- o R7: 192.168.1.7

5.1.2. Link IP addresses

- o R1 to R2 link: 10.0.0.1, 10.0.0.2
- o R1 to R4 link: 10.0.0.9, 10.0.0.10
- o R2 to R3 link #1: 10.0.0.3, 10.0.0.4

- o R2 to R3 link #2: 10.0.0.5, 10.0.0.6
- o R2 to R5 link: 10.0.0.7, 10.0.0.8
- o R3 to R6 link: 10.0.0.13, 10.0.0.14
- o R3 to R7 link: 10.0.0.15, 10.0.0.16
- o R4 to R5 link: 10.0.0.19, 10.0.0.20
- o R5 to R6 link: 10.0.0.11, 10.0.0.12
- o R6 to R7 link: 10.0.0.17, 10.0.0.18

The IGP link metrics are displayed in the middle of the link. All of them are assumed to be bi-directional.

5.2. One-hop LSP to an adjacent Router

If R1 would advertise a label <N> bound to a one-hop LSP from R1 to R2 it would encode as follows:

LSA 149, LSA-ID <N>:

IPv4 Prefix ERO TLV: 192.168.1.2/32, Strict

5.3. One-hop LSP to an adjacent Router using a specific link

If R2 would advertise a label <N> bound to a one-hop LSP from R2 to R3, using the link #2 it would encode as follows

LSA 149, LSA-ID <N>:

IPv4 Prefix ERO TLV: 10.0.0.6/32, Strict

5.4. One-hop LSP to an adjacent external Router

If R3 would advertise a label <N> bound to a one-hop LSP from R3 to R7 (which is outside of the IGP domain), it would encode as follows:

LSA 149, LSA-ID <N>:

IPv4 Prefix ERO TLV: 192.168.1.7/32, Strict

As you can see the representation of an MPLS label crossing an external link is identical as an internal link Section 5.2.

5.5. Advertisement of an RSVP LSP

Consider a RSVP LSP name "R2-to-R6" traversing (R2 to R3 using link #1, R6):

If R2 would advertise a label <N> bound to the RSVP LSP named 'R2-to-R6', it would encode as follows

LSA 149, LSA-ID <N>:

IPv4 Prefix ERO TLV: 10.0.0.4/32, Strict

IPv4 Prefix ERO TLV: 192.168.1.6/32, Strict

5.6. Advertisement of an LDP LSP

Consider R2 that creates a LDP label binding for FEC 172.16.0.0./12 using label <N>.

If R2 would re-advertise this label binding in OSPF it would encode as follows

LSA 149, LSA-ID <N>:

IPv4 Prefix ERO TLV: 172.16.0.0/12, Loose

5.7. Interarea advertisement of diverse paths

Consider two R2->R6 paths: {R2, R3, R6} and {R2, R5, R6}

Consider two R5->R3 paths: {R5, R2, R3} and {R5, R6, R3}

R2 encodes its two paths to R6 as follows:

LSA 149, LSA-ID <N1>:

IPv4 Prefix ERO TLV: 192.168.1.3, Loose

IPv4 Prefix ERO TLV: 192.168.1.6, Loose

Flags TLV: Down

LSA 149, LSA-ID <N2>:

IPv4 Prefix ERO TLV: 192.168.1.5, Loose

IPv4 Prefix ERO TLV: 192.168.1.6, Loose

Flags TLV: Down

R5 encodes its two paths to R3 as follows:

LSA 149, LSA-ID <N1>:

IPv4 Prefix ERO TLV: 192.168.1.2, Loose

IPv4 Prefix ERO TLV: 192.168.1.3, Loose

Flags TLV: Down

LSA 149, LSA-ID <N2>:

IPv4 Prefix ERO TLV: 192.168.1.6, Loose

IPv4 Prefix ERO TLV: 192.168.1.3, Loose

Flags TLV: Down

A receiving non-backbone router does see now all 4 paths and may decide to load-balance across all or a subset of them.

5.8. Advertisement of SPT labels using 'All Router Block' TLV

All routers within a given area MUST advertise their Label Blocks along with an 'ID'.

If R2 would advertise a label block <N1> with a size of 10, declaring SPT label forwarding support to all routers within a given domain, it would encode as follows:

LSA 149, LSA-ID <N1>:

All Router Block TLV: Block Size 10, Algo 0, Topology-ID 0

All Router ID IPv4 Map TLV: ID 2, 192.168.1.2

If R3 would advertise a label block <N2> with a size of 10, declaring SPT label forwarding support to all routers within a given domain, it would encode as follows:

LSA 149, LSA-ID <N2>:

All Router Block TLV: Block Size 10, Algo 0, Topology-ID 0

All Router ID IPv4 Map TLV: ID 3, 192.168.1.3

If R5 would advertise a label block <N3> with a size of 10, declaring SPT label forwarding support to all routers within a given domain, it would encode as follows:

LSA 149, LSA-ID <N3>:

All Router Block TLV: Block Size 10, Algo 0, Topology-ID 0

All Router ID IPv4 Map TLV: ID 5, 192.168.1.5

If R6 would advertise a label block <N4> with a size of 10, declaring SPT label forwarding support to all routers within a given domain, it would encode as follows:

LSA 149, LSA-ID <N4>:

All Router Block TLV: Block Size 10, Algo 0, Topology-ID 0

All Router ID IPv4 Map TLV: ID 6, 192.168.1.6

Consider now R2 constructing a SPT label for R6. R2s SPT to R6 is {R2, IP4, R3, R6}. R2 first determines if its downstream router (R3) has advertised a label-block. Since R3 has advertised a label block 'N2' and it has received R6 'ID' of 6 it will be picking the 6th label value inside the advertised range of its downstream neighbor. Specifically R2 MUST be program a MPLS SWAP for its own label range Label(N1+6) to Label(N2+6), NH 10.0.0.4 into its MPLS transit RIB. Furthermore R2 MAY program a MPLS PUSH operation for IP 192.168.1.6 to Label (N2+6), NH 10.0.0.4 into its IPv4 tunnel RIB.

Next walk down to R3, which is the next router on the SPT tree towards R6. R3s SPT to R6 is {R3, R6}. R3 determines if its downstream router (R6) has advertised a label-block. Since R6 has advertised a label block 'N4' and it has received R6 'ID' of 6 it will be picking the 6th label value inside the advertised range of its downstream neighbor. Since R3 is the penultimate router to R6 it MUST program a MPLS POP for its own label range Label(N2+6) NH 10.0.0.14 into its MPLS transit RIB. Furthermore R3 MAY program a MPLS NOP for IP 192.168.1.6, NH 10.0.0.14 into its IPv4 tunnel RIB.

5.9. Expansion of an 'All Router Block' TLV

All routers within a given area MUST advertise their Label Blocks along with an 'ID'. Now assume that the initial label block size assignment is too small to support all routers which generate an ordinal within an IGP domain. Consider the seven routers in Figure 16, and assume that R7 advertises a new ID '15' using an 'All Router ID Map' TLV. ID '15' is outside of the range of '10' as per

the previous example in Section 5.8. Now all the routers in an IGP domain need to advertise one more label block in order to map the ID '15' to an actual label value.

All routers would advertise in addition to their label block <N> with a size of 10, a second label block <N2> with a size sufficient enough, that the new ordinal can get covered. In this example the same block size 10 is used also for the second label block. For example router R2 would advertise the following label bindings.

LSA 149: LSA-ID <N1>:

All Router Block TLV: Block Size 10, Algo 0, Topology 0

All Router ID IPv4 Map TLV: ID 2, 192.168.1.2

LSA 149: LSA-ID <N2>:

All Router Block TLV: Block Size 10, Algo 0, Topology 0

Now the upstream router can map the new ID of R7 to an actual label value, as ID '15' corresponds to the 5th label inside the second Label block.

6. Inter Area Protocol Procedures

6.1. Applicability

Propagation of a MPLS LSPs and MPLS Block LSPs across an area boundary is a local policy decision.

6.2. Data plane operations

If local policy dictates that a given ABR router needs to re-advertise a MPLS LSPs from one area to another then it MUST allocate a new label and program its label forwarding table to connect the new label to the path in the respective other area. Depending on how to reach the re-advertised LSP, this is typically done using a MPLS 'SWAP' or 'SWAP/PUSH' data plane operation.

6.3. Control plane operations

6.3.1. MPLS Label operations

If local policy dictates that a given ABR router needs to re-advertise MPLS LSPs from one area to another then it must prepend its "Traffic-Engineering-ID" as a loose hop in the Prefix ERO TLV list.

Furthermore it MUST append the Flags TLV and set the 'Down' Bit.

6.3.2. MPLS Label Block operations

If local policy dictates that a given ABR router advertises its 'All Router Block' into another area, then it also MUST re-advertise all known 'ID' ordinals (again gated by policy) to the respective other area. Without knowledge of all 'ID's in the network no router is able to construct SPT label switched paths. Furthermore an ABR MUST append the Flags TLV and set the 'Down' Bit for all re-advertised 'CE' IDs.

7. Acknowledgements

Many thanks to Yakov Rekhter, John Drake and Shraddha Hedge for their useful comments.

8. IANA Considerations

This documents request allocation for one common OSPFv2 and OSPFv3 LSA Type and TLVs contained within.

LSA Type	TLV	TLV Type	#Occurence
149	IPv4 Prefix ERO	1	>=0
	IPv6 Prefix ERO	2	>=0
	Unnumbered Interface ID ERO	9	>=0
	IPv4 Prefix Bypass ERO	3	>=0
	IPv6 Prefix Bypass ERO	4	>=0
	Unnumbered Bypass Interface ID ERO	10	>=0
	Flags	5	0,1
	All Router Block	6	>=0
	All Router ID IPv4 Map	7	>=0
	All Router ID IPv6 Map	8	>=0

Table 1: IANA allocations

The MPLS Label LSA requires a new sub-registry, with a starting TLV value of 1, and managed by IETF consensus.

9. Security Considerations

This document does not introduce any change in terms of OSPF security. It simply proposes to flood MPLS label information via the IGP. All existing procedures to ensure message integrity do apply here.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", RFC 3107, May 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.
- [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.
- [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.
- [RFC4761] Kompella, K. and Y. Rekhter, "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, June 2007.
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", RFC 5036, October 2007.

- [RFC5151] Farrel, A., Ayyangar, A., and JP. Vasseur, "Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 5151, February 2008.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, July 2008.

10.2. Informative References

- [I-D.gredler-rtgwg-igp-label-advertisement]
Gredler, H., Amante, S., Scholl, T., and L. Jalil,
"Advertising MPLS labels in IGPs",
draft-gredler-rtgwg-igp-label-advertisement-05 (work in
progress), May 2013.
- [I-D.ietf-rtgwg-mrt-frr-architecture]
Atlas, A., Kebler, R., Envedi, G., Csaszar, A., Tantsura,
J., Konstantynowicz, M., White, R., and M. Shand, "An
Architecture for IP/LDP Fast-Reroute Using Maximally
Redundant Trees", draft-ietf-rtgwg-mrt-frr-architecture-02
(work in progress), February 2013.
- [I-D.previdi-filsfils-isis-segment-routing]
Previdi, S., Filsfils, C., Bashandy, A., Horneffer, M.,
Decraene, B., Litkowski, S., Milojevic, I., Shakir, R.,
Ytti, S., Henderickx, W., and J. Tantsura, "Segment
Routing with IS-IS Routing Protocol",
draft-previdi-filsfils-isis-segment-routing-02 (work in
progress), March 2013.

Authors' Addresses

Hannes Gredler (editor)
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: hannes@juniper.net

Shane Amante
Level 3 Communications, Inc.
1025 Eldorado Blvd
Broomfield, CO 80021
US

Email: shane@level3.net

Tom Scholl
Amazon
Seattle, WN
US

Email: tscholl@amazon.com

Luay Jalil
Verizon
1201 E Arapaho Rd.
Richardson, TX 75081
US

Email: luay.jalil@verizon.com

