

Open Shortest Path First IGP
Internet-Draft
Intended status: Standards Track
Expires: December 28, 2015

P. Psenak, Ed.
S. Previdi, Ed.
C. Filsfils
Cisco Systems, Inc.
H. Gredler
Juniper Networks, Inc.
R. Shakir
British Telecom
W. Henderickx
Alcatel-Lucent
J. Tantsura
Ericsson
June 26, 2015

OSPFv3 Extensions for Segment Routing
draft-ietf-ospf-ospfv3-segment-routing-extensions-03

Abstract

Segment Routing (SR) allows for a flexible definition of end-to-end paths within IGP topologies by encoding paths as sequences of topological sub-paths, called "segments". These segments are advertised by the link-state routing protocols (IS-IS and OSPF).

This draft describes the OSPFv3 extensions that are required for Segment Routing.

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 December 28, 2015.

Copyright Notice

Copyright (c) 2015 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	3
2.	Segment Routing Identifiers	3
2.1.	SID/Label Sub-TLV	3
3.	Segment Routing Capabilities	4
3.1.	SR-Algorithm TLV	4
3.2.	SID/Label Range TLV	5
3.3.	SR-Forwarding Capabilities	7
4.	OSPFv3 Extended Prefix Range TLV	8
5.	Prefix SID Sub-TLV	9
6.	SID/Label Binding Sub-TLV	13
6.1.	ERO Metric Sub-TLV	15
6.2.	ERO Sub-TLVs	16
6.2.1.	IPv4 ERO Sub-TLV	16
6.2.2.	IPv6 ERO Sub-TLV	17
6.2.3.	Unnumbered Interface ID ERO Sub-TLV	18
6.2.4.	IPv4 Backup ERO Sub-TLV	19
6.2.5.	IPv6 Backup ERO Sub-TLV	20
6.2.6.	Unnumbered Interface ID Backup ERO Sub-TLV	21
7.	Adjacency Segment Identifier (Adj-SID)	22
7.1.	Adj-SID Sub-TLV	22
7.2.	LAN Adj-SID Sub-TLV	24
8.	Elements of Procedure	26
8.1.	Intra-area Segment routing in OSPFv3	26
8.2.	Inter-area Segment routing in OSPFv3	27
8.3.	SID for External Prefixes	28
8.4.	Advertisement of Adj-SID	28
8.4.1.	Advertisement of Adj-SID on Point-to-Point Links	28
8.4.2.	Adjacency SID on Broadcast or NBMA Interfaces	28
9.	IANA Considerations	29

9.1.	OSPF Router Information (RI) TLVs Registry	29
9.2.	OSPFv3 Extend-LSA TLV Registry	29
9.3.	OSPFv3 Extend-LSA Sub-TLV registry	29
10.	Security Considerations	30
11.	Acknowledgements	30
12.	References	30
12.1.	Normative References	30
12.2.	Informative References	30
	Authors' Addresses	31

[1.](#) Introduction

Segment Routing (SR) allows for a flexible definition of end-to-end paths within IGP topologies by encoding paths as sequences of topological sub-paths, called "segments". These segments are advertised by the link-state routing protocols (IS-IS and OSPF). Prefix segments represent an ecmp-aware shortest-path to a prefix (or a node), as per the state of the IGP topology. Adjacency segments represent a hop over a specific adjacency between two nodes in the IGP. A prefix segment is typically a multi-hop path while an adjacency segment, in most of the cases, is a one-hop path. SR's control-plane can be applied to both IPv6 and MPLS data-planes, and does not require any additional signaling (other than the regular IGP). For example, when used in MPLS networks, SR paths do not require any LDP or RSVP-TE signaling. Still, SR can interoperate in the presence of LSPs established with RSVP or LDP.

This draft describes the OSPFv3 extensions required for segment routing.

Segment Routing architecture is described in [\[I-D.ietf-spring-segment-routing\]](#).

Segment Routing use cases are described in [\[I-D.filsfils-spring-segment-routing-use-cases\]](#).

[2.](#) Segment Routing Identifiers

Segment Routing defines various types of Segment Identifiers (SIDs): Prefix-SID, Adjacency-SID, LAN Adjacency SID and Binding SID.

[2.1.](#) SID/Label Sub-TLV

The SID/Label Sub-TLV appears in multiple TLVs or Sub-TLVs defined later in this document. It is used to advertise the SID or label associated with a prefix or adjacency. The SID/Label TLV has following format:

An OSPFv3 router may use various algorithms when calculating reachability to other nodes in area or to prefixes attached to these nodes. Examples of these algorithms are metric based Shortest Path First (SPF), various sorts of Constrained SPF, etc. The SR-Algorithm TLV allows a router to advertise the algorithms that the router is currently using to other routers in an area. The SR-Algorithm TLV has following structure:


```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Type               |               Length               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Algorithm 1 | Algorithm... | Algorithm n |                           |
+-+-----+-----+-----+-----+-----+-----+-----+-----+
|                                                     |
+                                                     +

```

where:

Type: TBD, suggested value 8

Length: variable

Algorithm: Single octet identifying the algorithm. The following value has been defined:

0: Shortest Path First (SPF) algorithm based on link metric. This is the standard shortest path algorithm as computed by the OSPF protocol. Consistent with the deployed practice for link-state protocols, Algorithm 0 permits any node to overwrite the SPF path with a different path based on its local policy.

1: Strict Shortest Path First (SPF) algorithm based on link metric. The algorithm is identical to Algorithm 0 but Algorithm 1 requires that all nodes along the path will honor the SPF routing decision. Local policy at the node claiming the support of Algorithm 1 MUST NOT alter the forwarding decision computed by Algorithm 1.

The RI LSA can be advertised at any of the defined flooding scopes (link, area, or autonomous system (AS)). For the purpose of the SR-Algorithm TLV propagation, area scope flooding is required.

3.2. SID/Label Range TLV

The SID/Label Range TLV is a TLV of the OSPFv3 Router Information LSA (defined in [[RFC4970](#)]).

The SID/Label Sub-TLV MAY appear multiple times and has following format:


```

0                                     1                                     2                                     3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Type                                     |                                     Length                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Range Size                                     |                                     Reserved                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Sub-TLVs (variable)                                     |
+-
|
+

```

where:

Type: TBD, suggested value 9

Length: variable

Range Size: 3 octets of SID/label range

Initially, the only supported Sub-TLV is the SID/Label TLV as defined in [Section 2.1](#). The SID/Label advertised in the SID/Label TLV represents the first SID/Label in the advertised range.

Multiple occurrence of the SID/Label Range TLV MAY be advertised, in order to advertise multiple ranges. In such case:

- o The originating router MUST encode each range into a different SID/Label Range TLV.
- o The originating router decides the order in which the set of SID/Label Range TLVs are advertised in the OSPFv3 Router Information LSA. The originating router MUST ensure the order is same after a graceful restart (using checkpointing, non-volatile storage or any other mechanism) in order to assure the SID/label range and SID index correspondence is preserved across graceful restarts.
- o The receiving router must adhere to the order in which the ranges are advertised when calculating a SID/label from the SID index.
- o A router not supporting multiple occurrences of the SID/Label Range TLV MUST use first advertised SID/Label Range TLV.

The following example illustrates the advertisement of multiple ranges:

The originating router advertises the following ranges:

Range 1: [100, 199]
Range 2: [1000, 1099]
Range 3: [500, 599]

The receiving routers concatenate the ranges and build the Segment Routing Global Block (SRGB) is as follows:

SRGB = [100, 199]
[1000, 1099]
[500, 599]

The indexes span multiple ranges:

index=0 means label 100
...
index 99 means label 199
index 100 means label 1000
index 199 means label 1099
...
index 200 means label 500
...

The RI LSA can be advertised at any of the defined flooding scopes (link, area, or autonomous system (AS)). For the purpose of the SR-Capability TLV propagation, area scope flooding is required.

3.3. SR-Forwarding Capabilities

OSPFv3 router supporting Segment Routing needs to advertise its SR data-plane capabilities. Data-plane capabilities are advertised in OSPF Router Informational Capabilities TLV, which is defined in [section 2.3 of RFC 4970](#) [[RFC4970](#)].

Two new bits are allocated in the OSPF Router Informational Capability Bits as follows:

Bit-6 - MPLS IPv6 flag. If set, then the router is capable of processing SR MPLS encapsulated IPv6 packets on all interfaces.

Bit-7 - If set, then the router is capable of processing the IPv6 Segment Routing Header on all interfaces as defined in [[I-D.previdi-6man-segment-routing-header](#)].

For the purpose of the SR-Forwarding Capabilities propagation, area scope flooding is required.

4. OSPFv3 Extended Prefix Range TLV

In some cases it is useful to advertise attributes for a range of prefixes. Segment Routing Mapping Server, which is described in [[I-D.filsfils-spring-segment-routing-ldp-interop](#)], is an example where we need a single advertisement to advertise SIDs for multiple prefixes from a contiguous address range. The OSPFv3 Extended Prefix Range TLV is defined for this purpose.

The OSPFv3 Extended Prefix Range TLV is a new top level TLV of the following LSAs defined in [[I-D.ietf-ospf-ospfv3-lsa-extend](#)]:

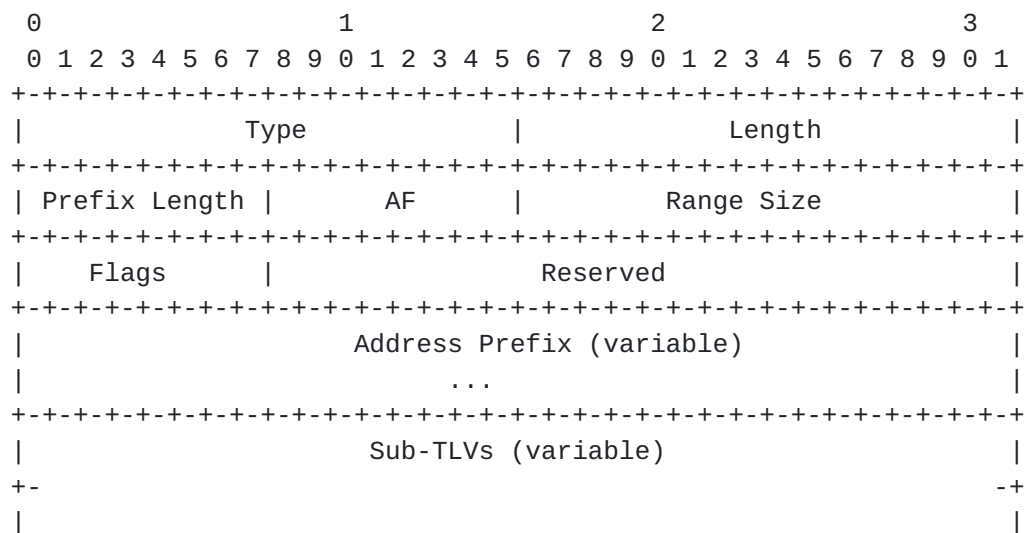
E-Intra-Area-Prefix-LSA

E-Inter-Area-Prefix-LSA

E-AS-External-LSA

E-Type-7-LSA

Multiple OSPFv3 Extended Prefix Range TLVs MAY be advertised in these extended LSAs. The OSPFv3 Extended Prefix Range TLV has the following format:



where:

Type: TBD, suggested value 9.

Length: variable

Prefix length: length of the prefix

AF: 0 - IPv6 unicast

Range size: represents the number of prefixes that are covered by the advertisement. The Range Size MUST NOT exceed the number of prefixes that could be satisfied by the prefix length without including addresses from other than the IPv6 unicast address class.

Flags: 1 octet field. The following flags are defined:

```

  0  1  2  3  4  5  6  7
+--+--+--+--+--+--+--+
|IA|  |  |  |  |  |  |
+--+--+--+--+--+--+--+

```

where:

IA-Flag: Inter-Area flag. If set, advertisement is of inter-area type. ABR that is advertising the OSPF Extended Prefix Range TLV between areas MUST set this bit.

This bit is used to prevent redundant flooding of Prefix Range TLVs between areas as follows:

An ABR always prefers intra-area Prefix Range advertisement over inter-area one.

An ABR does not consider inter-area Prefix Range advertisements coming from non backbone area.

An ABR propagates inter-area Prefix Range advertisement from backbone area to connected non backbone areas only if such advertisement is considered to be the best one.

Address Prefix: the prefix, encoded as an even multiple of 32-bit words, padded with zeroed bits as necessary. This encoding consumes $((\text{PrefixLength} + 31) / 32)$ 32-bit words. The Address Prefix represents the first prefix in the prefix range.

5. Prefix SID Sub-TLV

The Prefix SID Sub-TLV is a Sub-TLV of the following OSPFv3 TLVs as defined in [[I-D.ietf-ospf-ospfv3-lsa-extend](#)] and in [Section 4](#):

Intra-Area Prefix TLV

Inter-Area Prefix TLV

External Prefix TLV

OSPFv3 Extended Prefix Range TLV

It MAY appear more than once in the parent TLV and has the following format:

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Type                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Flags   | Algorithm |           Reserved           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               SID/Index/Label (variable)         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

Type: TBD, suggested value 4.

Length: variable

Flags: 1 octet field. The following flags are defined:

```

      0  1  2  3  4  5  6  7
+---+---+---+---+---+---+---+
| |NP|M |E |V |L | | |
+---+---+---+---+---+---+---+

```

where:

NP-Flag: No-PHP flag. If set, then the penultimate hop MUST NOT pop the Prefix-SID before delivering the packet to the node that advertised the Prefix-SID.

M-Flag: Mapping Server Flag. If set, the SID is advertised from the Segment Routing Mapping Server functionality as described in [[I-D.filsfils-spring-segment-routing-ldp-interop](#)].

E-Flag: Explicit-Null Flag. If set, any upstream neighbor of the Prefix-SID originator MUST replace the Prefix-SID with a Prefix-SID having an Explicit-NULL value (0 for IPv4) before forwarding the packet.

The V-Flag: Value/Index Flag. If set, then the Prefix-SID carries an absolute value. If not set, then the Prefix-SID carries an index.

The L-Flag: Local/Global Flag. If set, then the value/index carried by the Prefix-SID has local significance. If not set, then the value/index carried by this Sub-TLV has global significance.

Other bits: Reserved. These MUST be zero when sent and are ignored when received.

Algorithm: one octet identifying the algorithm the Prefix-SID is associated with as defined in [Section 3.1](#).

SID/Index/Label: label or index value depending on the V-bit setting.

Examples:

A 32 bit global index defining the offset in the SID/Label space advertised by this router - in this case the V and L flags MUST NOT be set.

A 24 bit local label where the 20 rightmost bits are used for encoding the label value - in this case the V and L flags MUST be set.

If multiple Prefix-SIDs are advertised for the same prefix, the receiving router MUST use the first encoded SID and MAY use the subsequent SIDs.

When propagating Prefix-SIDs between areas, if multiple prefix-SIDs are advertised for a prefix, an implementation SHOULD preserve the original order when advertising prefix-SIDs to other areas. This allows implementations that only support a single Prefix-SID to have a consistent view across areas.

When calculating the outgoing label for the prefix, the router MUST take into account E and P flags advertised by the next-hop router, if next-hop router advertised the SID for the prefix. This MUST be done regardless of whether the next-hop router contributes to the best path to the prefix.

The NP-Flag (No-PHP) MUST be set for Prefix-SIDs allocated to inter-area prefixes that are originated by the ABR based on intra-area or inter-area reachability between areas. When the inter-area prefix is

generated based on a prefix which is directly attached to the ABR, NP-Flag SHOULD NOT be set

The NP-Flag (No-PHP) MUST be set on the Prefix-SIDs allocated to redistributed prefixes, unless the redistributed prefix is directly attached to ASBR, in which case the NP-Flag SHOULD NOT be set.

If the NP-Flag is not set then any upstream neighbor of the Prefix-SID originator MUST pop the Prefix-SID. This is equivalent to the penultimate hop popping mechanism used in the MPLS dataplane. In such case, MPLS EXP bits of the Prefix-SID are not preserved for the final destination (the Prefix-SID being removed). If the NP-Flag is clear then the received E-flag is ignored.

If the NP-Flag is set then:

If the E-flag is not set then any upstream neighbor of the Prefix-SID originator MUST keep the Prefix-SID on top of the stack. This is useful when the originator of the Prefix-SID must stitch the incoming packet into a continuing MPLS LSP to the final destination. This could occur at an inter-area border router (prefix propagation from one area to another) or at an inter-domain border router (prefix propagation from one domain to another).

If the E-flag is set then any upstream neighbor of the Prefix-SID originator MUST replace the Prefix-SID with a Prefix-SID having an Explicit-NULL value. This is useful, e.g., when the originator of the Prefix-SID is the final destination for the related prefix and the originator wishes to receive the packet with the original EXP bits.

When M-Flag is set, NP-Flag MUST be set and E-bit MUST NOT be set.

As the Mapping Server does not specify the originator of a prefix advertisement it is not possible to determine PHP behavior solely based on the Mapping Server advertisement. However, PHP behavior may safely be done in following cases:

Prefix is of intra-area type and the downstream neighbor is the originator of the prefix.

Prefix is of inter-area type and downstream neighbor is an ABR, which is advertising the prefix reachability and is setting LA-bit in the Prefix Options as described in section 3.1 of [\[I-D.ietf-ospf-ospfv3-lsa-extend\]](#).

Prefix is of external type and downstream neighbor is an ASBR, which is advertising the prefix reachability and is setting LA-bit in the Prefix Options as described in section 3.1 of [\[I-D.ietf-ospf-ospfv3-lsa-extend\]](#).

When a Prefix-SID is advertised in an Extended Prefix Range TLV, then the value advertised in Prefix SID Sub-TLV is interpreted as a starting SID value.

Example 1: if the following router addresses (loopback addresses) need to be mapped into the corresponding Prefix SID indexes:

```
Router-A: 192::1/128, Prefix-SID: Index 1
Router-B: 192::2/128, Prefix-SID: Index 2
Router-C: 192::3/128, Prefix-SID: Index 3
Router-D: 192::4/128, Prefix-SID: Index 4
```

then the Address Prefix field in the OSPFv3 Extended Prefix Range TLV is set to 192::1, Prefix Length would be set to 128, Range Size would be set to 4 and the Index value in the Prefix-SID Sub-TLV would be set to 1.

Example 2: If the following prefixes need to be mapped into the corresponding Prefix-SID indexes:

```
10:1:1::0/120, Prefix-SID: Index 51
10:1:1::100/120, Prefix-SID: Index 52
10:1:1::200/120, Prefix-SID: Index 53
10:1:1::300/120, Prefix-SID: Index 54
10:1:1::400/120, Prefix-SID: Index 55
10:1:1::500/120, Prefix-SID: Index 56
10:1:1::600/120, Prefix-SID: Index 57
```

then the Address Prefix field in the OSPFv3 Extended Prefix Range TLV is set to 10:1:1::0, Prefix Length would be set to 120, Range Size would be set to 7 and the Index value in the Prefix-SID Sub-TLV would be set to 51.

6. SID/Label Binding Sub-TLV

The SID/Label Binding Sub-TLV is used to advertise SID/Label mapping for a path to the prefix.

The SID/Label Binding TLV MAY be originated by any router in an OSPFv3 domain. The router may advertise a SID/Label binding to a FEC along with at least a single 'nexthop style' anchor. The protocol supports more than one 'nexthop style' anchor to be attached to a SID/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 anchoring SID/label bindings to both link and node IP addresses, any Label Switched Path (LSP) can be described. Furthermore, SID/Label Bindings from external protocols can also be re-advertised.

The SID/Label Binding TLV may be used for advertising SID/Label Bindings and their associated Primary and Backup paths. In one single TLV, either a primary ERO Path, backup ERO Path, or both are advertised. If a router wants to advertise multiple parallel paths, then it can generate several TLVs for the same Prefix/FEC. Each occurrence of a Binding TLV for a given FEC Prefix will add a new path.

SID/Label Binding Sub-TLV is a Sub-TLV of the following OSPFv3 TLVs, as defined in [[I-D.ietf-ospf-ospfv3-lsa-extend](#)] and in [Section 4](#):

Intra-Area Prefix TLV

Inter-Area Prefix TLV

External Prefix TLV

OSPFv3 Extended Prefix Range TLV

Multiple SID/Label Binding Sub-TLVs can be present in these TLVs. The SID/Label Binding Sub-TLV has following format:

0																1																2																3															
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9																								
Type																Length																																															
Flags								Weight								Reserved																																															
Sub-TLVs (variable)																																																															

where:

Type: TBD, suggested value 7

Length: variable

Flags: 1 octet field of following flags:


```

 0 1 2 3 4 5 6 7
+---+---+---+---+
|M|               |
+---+---+---+---+

```

where:

M-bit - When the bit is set the binding represents the mirroring context as defined in [\[I-D.minto-rsvp-lsp-egress-fast-protection\]](#).

Weight: weight used for load-balancing purposes. The use of the weight is defined in section 3.5.1 of [\[I-D.ietf-spring-segment-routing\]](#).

SID/Label Binding Sub-TLV currently supports following Sub-TLVs:

SID/Label Sub-TLV as described in [Section 2.1](#). This Sub-TLV MUST appear in the SID/Label Binding Sub-TLV and it MUST only appear once.

ERO Metric Sub-TLV as defined in [Section 6.1](#).

ERO Sub-TLVs as defined in [Section 6.2](#).

[6.1](#). ERO Metric Sub-TLV

The ERO Metric Sub-TLV is a Sub-TLV of the SID/Label Binding TLV.

The ERO Metric Sub-TLV advertises the cost of an ERO path. It is used to compare the cost of a given source/destination path. A router SHOULD advertise the ERO Metric Sub-TLV in an advertised ERO TLV. The cost of the ERO Metric Sub-TLV SHOULD be set to the cumulative IGP or TE path cost of the advertised ERO. Since manipulation of the Metric field may attract or repel traffic to and from the advertised segment, it MAY be manually overridden.

```

 0                               1                               2                               3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Type                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Metric (4 octets)                   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

ERO Metric Sub-TLV format

where:

Type: TBD, suggested value 8

Length: Always 4

Metric: A 4 octet metric representing the aggregate IGP or TE path cost.

6.2. ERO Sub-TLVs

All 'ERO' information represents an ordered set which describes the segments of a path. The first ERO Sub-TLV describes the first segment of a path. Similarly, the last ERO Sub-TLV describes the segment closest to the egress point. If a router extends or stitches a path, it **MUST** prepend the new segment's path information to the ERO list. This applies equally to advertised backup EROs.

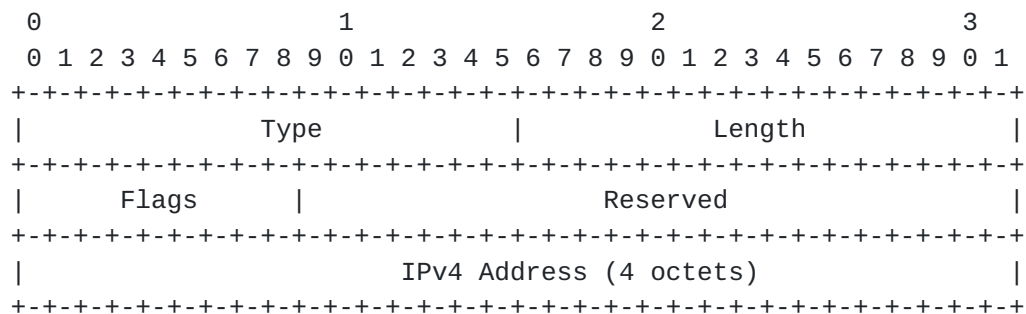
All ERO Sub-TLVs must immediately follow the (SID)/Label Sub-TLV.

All Backup ERO Sub-TLVs must immediately follow the last ERO Sub-TLV.

6.2.1. IPv4 ERO Sub-TLV

IPv4 ERO Sub-TLV is a Sub-TLV of the SID/Label Binding Sub-TLV.

The IPv4 ERO Sub-TLV describes a path segment using IPv4 Address style of encoding. Its semantics have been borrowed from [RFC3209].



IPv4 ERO Sub-TLV format

where:

Type: TBD, suggested value 9

Length: 8 bytes

Flags: 1 octet field of following flags:


```

 0 1 2 3 4 5 6 7
+---+---+---+---+
| L |           |
+---+---+---+---+

```

where:

L-bit - If the L-bit is set, then the segment path is designated as 'loose'. Otherwise, the segment path is designated as 'strict'.

IPv6 Address - the address of the explicit route hop.

6.2.3. Unnumbered Interface ID ERO Sub-TLV

The Unnumbered Interface ID ERO Sub-TLV is a Sub-TLV of the SID/Label Binding Sub-TLV.

The appearance and semantics of the 'Unnumbered Interface ID' have been borrowed from [\[RFC3477\]](#).

The Unnumbered Interface-ID ERO Sub-TLV 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.

```

      0              1              2              3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
|      Type                    |      Length                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Flags                   |      Reserved               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |      Router ID              |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |      Interface ID           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

Unnumbered Interface ID ERO Sub-TLV format

Type: TBD, suggested value 11

Length: 12 bytes

Flags: 1 octet field of following flags:

```

0 1 2 3 4 5 6 7
+---+---+---+---+
|L|               |
+---+---+---+---+

```

where:

L-bit - If the L-bit is set, then the segment path is designated as 'loose'. Otherwise, the segment path is designated as 'strict'.

Router-ID: Router-ID of the next-hop.

Interface ID: is the identifier assigned to the link by the router specified by the Router-ID.

6.2.4. IPv4 Backup ERO Sub-TLV

IPv4 Prefix Backup ERO Sub-TLV is a Sub-TLV of the SID/Label Binding Sub-TLV.

The IPv4 Backup ERO Sub-TLV describes a path segment using IPv4 Address style of encoding. Its semantics have been borrowed from [\[RFC3209\]](#).

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Type                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Flags                         |                               Reserved                         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IPv4 Address (4 octets)                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

IPv4 Backup ERO Sub-TLV format

where:

Type: TBD, suggested value 12

Length: 8 bytes

Flags: 1 octet field of following flags:


```

 0 1 2 3 4 5 6 7
+---+---+---+---+
|L|               |
+---+---+---+---+

```

where:

L-bit - If the L-bit is set, then the segment path is designated as 'loose'. Otherwise, the segment path is designated as 'strict'.

IPv4 Address - the address of the explicit route hop.

6.2.5. IPv6 Backup ERO Sub-TLV

The IPv6 ERO Sub-TLV is a Sub-TLV of the SID/Label Binding Sub-TLV.

The IPv6 Backup ERO Sub-TLV describes a Backup path segment using IPv6 Address style of encoding. Its appearance and semantics have been borrowed from [\[RFC3209\]](#).

The 'L' bit in the Flags 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.'

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Type               |               Length               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Flags               |               Reserved               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               |               |               |
+-+               +-+               +-+
|               |               |               |
+-+               +-+               +-+
|               |               |               |
+-+               +-+               +-+
|               |               |               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

IPv6 Backup ERO Sub-TLV format

where:

Type: TBD, suggested value 13

Length: 8 bytes

Flags: 1 octet field of following flags:

```

 0 1 2 3 4 5 6 7
+---+---+---+---+
|L|               |
+---+---+---+---+

```

where:

L-bit - If the L-bit is set, then the segment path is designated as 'loose'. Otherwise, the segment path is designated as 'strict'.

IPv6 Address - the address of the explicit route hop.

6.2.6. Unnumbered Interface ID Backup ERO Sub-TLV

The Unnumbered Interface ID Backup Sub-TLV is a Sub-TLV of the SID/Label Binding Sub-TLV.

The appearance and semantics of the 'Unnumbered Interface ID' have been borrowed from [[RFC3477](#)].

The Unnumbered Interface-ID Backup ERO Sub-TLV 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 are 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 with specification of the unique Router-ID.

```

      0              1              2              3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type      |                               Length      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Flags      |                               Reserved      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Router ID                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Interface ID                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Unnumbered Interface ID Backup ERO Sub-TLV format

where:

Type: TBD, suggested value 14

Length: 12 bytes

Flags: 1 octet field of following flags:

```

0 1 2 3 4 5 6 7
+---+---+---+---+
|L|           |
+---+---+---+---+

```

where:

L-bit - If the L-bit is set, then the segment path is designated as 'loose'. Otherwise, the segment path is designated as 'strict'.

Router-ID: Router-ID of the next-hop.

Interface ID: is the identifier assigned to the link by the router specified by the Router-ID.

7. Adjacency Segment Identifier (Adj-SID)

An Adjacency Segment Identifier (Adj-SID) represents a router adjacency in Segment Routing.

7.1. Adj-SID Sub-TLV

The extended OSPFv3 LSAs, as defined in [[I-D.ietf-ospf-ospfv3-lsa-extend](#)], are used to advertise prefix SID in OSPFv3

The Adj-SID Sub-TLV is an optional Sub-TLV of the Router-Link TLV as defined in [[I-D.ietf-ospf-ospfv3-lsa-extend](#)]. It MAY appear multiple times in Router-Link TLV. Examples where more than one Adj-SID may be used per neighbor are described in section 4 of [[I-D.filsfils-spring-segment-routing-use-cases](#)]. The Adj-SID Sub-TLV has the following format:

Weight: weight used for load-balancing purposes. The use of the weight is defined in section 3.5.1 of [\[I-D.ietf-spring-segment-routing\]](#).

SID/Index/Label: label or index value depending on the V-bit setting.

Examples:

A 32 bit global index defining the offset in the SID/Label space advertised by this router - in this case the V and L flags MUST NOT be set.


A 24 bit local label where the 20 rightmost bits are used for encoding the label value - in this case the V and L flags MUST be set.

16 octet IPv6 address - in this case the V-flag MUST be set. The L-flag MUST be set for link-local IPv6 address and MUST NOT be set for IPv6 global unicast address.

An SR capable router MAY allocate an Adj-SID for each of its adjacencies and set the B-Flag when the adjacency is eligible for protection by an FRR mechanism (IP or MPLS) as described in [section 3.5](#) of [\[I-D.ietf-spring-segment-routing\]](#).

7.2. LAN Adj-SID Sub-TLV

The LAN Adj-SID is an optional Sub-TLV of the Router-Link TLV. It MAY appear multiple times in the Router-Link TLV. It is used to advertise a SID/Label for an adjacency to a non-DR neighbor on a broadcast or NBMA network.



where:

Type: TBD, suggested value 6.

Length: variable.

Flags. 1 octet field of following flags:

```

 0 1 2 3 4 5 6 7
+---+---+---+---+
|B|V|L|S|   |
+---+---+---+---+

```

where:

B-Flag: Backup-flag: set if the LAN-Adj-SID refer to an adjacency that is eligible for protection (e.g.: using IPFRR or MPLS-FRR) as described in section 3.1 of [\[I-D.filsfils-spring-segment-routing-use-cases\]](#).

The V-Flag: Value/Index Flag. If set, then the LAN Adj-SID carries an absolute value. If not set, then the LAN Adj-SID carries an index.

The L-Flag: Local/Global Flag. If set, then the value/index carried by the LAN Adj-SID has local significance. If not set, then the value/index carried by this subTLV has global significance.

The S-Flag. Set Flag. When set, the S-Flag indicates that the LAN Adj-SID refers to a set of adjacencies (and therefore MAY be assigned to other adjacencies as well).

Other bits: Reserved. These MUST be zero when sent and are ignored when received.

Weight: weight used for load-balancing purposes. The use of the weight is defined in section 3.5.1 of [\[I-D.ietf-spring-segment-routing\]](#).

SID/Index/Label: label or index value depending on the V-bit setting.

Examples:

A 32 bit global index defining the offset in the SID/Label space advertised by this router - in this case the V and L flags MUST NOT be set.

A 24 bit local label where the 20 rightmost bits are used for encoding the label value - in this case the V and L flags MUST be set.

16 octet IPv6 address - in this case the V-flag MUST be set. The L-flag MUST be set for link-local IPv6 address and MUST NOT be set for IPv6 global unicast address.

8. Elements of Procedure

8.1. Intra-area Segment routing in OSPFv3

An OSPFv3 router that supports segment routing MAY advertise Prefix-SIDs for any prefix that it is advertising reachability for (e.g., loopback IP address) as described in [Section 5](#).

If multiple routers advertise a Prefix-SID for the same prefix, then the Prefix-SID MUST be the same. This is required in order to allow traffic load-balancing when multiple equal cost paths to the destination exist in the network.

The Prefix-SID can also be advertised by the SR Mapping Servers (as described in [[I-D.filsfils-spring-segment-routing-ldp-interop](#)]). The Mapping Server advertises Prefix-SID for remote prefixes that exist in the network. Multiple Mapping Servers can advertise Prefix-SID for the same prefix, in which case the same Prefix-SID MUST be advertised by all of them. The SR Mapping Server could use either area scope or autonomous system flooding scope when advertising Prefix SID for prefixes, based on the configuration of the SR Mapping Server. Depending on the flooding scope used, the SR Mapping Server chooses the LSA that will be used. If the area flooding scope is needed, E-Intra-Area-Prefix-LSA ([[I-D.ietf-ospf-ospfv3-lsa-extend](#)]) is used. If autonomous system flooding scope is needed, E-AS-External-LSA ([[I-D.ietf-ospf-ospfv3-lsa-extend](#)]) is used.

When a Prefix-SID is advertised by the Mapping Server, which is indicated by the M-flag in the Prefix-SID Sub-TLV ([Section 5](#)), the route type as implied by the LSA type is ignored and the Prefix-SID is bound to the corresponding prefix independent of the route type.

Advertisement of the Prefix-SID by the Mapping Server using Inter-Area Prefix TLV, External-Prefix TLV or Intra-Area-Prefix TLV ([[I-D.ietf-ospf-ospfv3-lsa-extend](#)]) does not itself contribute to the prefix reachability. The NU-bit MUST be set in the PrefixOptions field of the LSA which is used by the Mapping Server to advertise SID or SID range, which prevents the advertisement to contribute to prefix reachability.

SR Mapping Server MUST use OSPF Extended Prefix Range TLV when advertising SIDs for prefixes. Prefixes of different route-types can be combined in a single OSPF Extended Prefix Range TLV advertised by the SR Mapping Server.

Area scoped OSPF Extended Prefix Range TLV are propagated between areas. Similar to propagation of prefixes between areas, ABR only propagates the OSPF Extended Prefix Range TLV that it considers to be the best from the set it received. The rules used to pick the best OSPF Extended Prefix Range TLV is described in [Section 4](#).

When propagating OSPF Extended Prefix Range TLV between areas, ABR MUST set the IA-Flag, that is used to prevent redundant flooding of the OSPF Extended Prefix Range TLV between areas as described in [Section 4](#).

If the Prefix-SID that is advertised in Prefix SID Sub-TLV is also covered by the OSPF Extended Prefix Range TLV, the Prefix-SID advertised in Prefix SID Sub-TLV MUST be preferred.

8.2. Inter-area Segment routing in OSPFv3

In order to support SR in a multi-area environment, OSPFv3 must propagate Prefix-SID information between areas. The following procedure is used in order to propagate Prefix SIDs between areas.

When an OSPFv3 ABR advertises a Inter-Area-Prefix-LSA from an intra-area prefix to all its connected areas, it will also include Prefix-SID Sub-TLV, as described in [Section 5](#). The Prefix-SID value will be set as follows:

The ABR will look at its best path to the prefix in the source area and find out the advertising router associated with the best path to that prefix.

The ABR will then determine if such router advertised a Prefix-SID for the prefix and use it when advertising the Prefix-SID to other connected areas.

If no Prefix-SID was advertised for the prefix in the source area by the router that contributes to the best path to the prefix, the originating ABR will use the Prefix-SID advertised by any other router (e.g.: a Prefix-SID coming from an SR Mapping Server as defined in [[I-D.filsfils-spring-segment-routing-ldp-interop](#)]) when propagating Prefix-SID for the prefix to other areas.

When an OSPFv3 ABR advertises Inter-Area-Prefix-LSA LSAs from an inter-area route to all its connected areas it will also include Prefix-SID Sub-TLV, as described in [Section 5](#). The Prefix-SID value will be set as follows:

The ABR will look at its best path to the prefix in the source area and find out the advertising router associated with the best path to that prefix.

The ABR will then look if such router advertised a Prefix-SID for the prefix and use it when advertising the Prefix-SID to other connected areas.

If no Prefix-SID was advertised for the prefix in the source area by the ABR that contributes to the best path to the prefix, the originating ABR will use the Prefix-SID advertised by any other router when propagating Prefix-SID for the prefix to other areas.

8.3. SID for External Prefixes

AS-External-LSAs are flooded domain wide. When an ASBR, which supports SR, generates E-AS-External-LSA, it should also include Prefix-SID Sub-TLV, as described in [Section 5](#). The Prefix-SID value will be set to the SID that has been reserved for that prefix.

When an NSSA ASBR translates an E-NSSA-LSA into an E-AS-External-LSA, it should also advertise the Prefix-SID for the prefix. The NSSA ABR determines its best path to the prefix advertised in the translated E-NSSA-LSA and finds the advertising router associated with that path. If the advertising router has advertised a Prefix-SID for the prefix, then the NSSA ABR uses it when advertising the Prefix-SID in the E-AS-External-LSA. Otherwise the Prefix-SID advertised by any other router will be used.

8.4. Advertisement of Adj-SID

The Adjacency Segment Routing Identifier (Adj-SID) is advertised using the Adj-SID Sub-TLV as described in [Section 7](#).

8.4.1. Advertisement of Adj-SID on Point-to-Point Links

An Adj-SID MAY be advertised for any adjacency on p2p link that is in a state 2-Way or higher. If the adjacency on a p2p link transitions from the FULL state, then the Adj-SID for that adjacency MAY be removed from the area. If the adjacency transitions to a state lower than 2-Way, then the Adj-SID advertisement MUST be removed from the area.

8.4.2. Adjacency SID on Broadcast or NBMA Interfaces

Broadcast or NBMA networks in OSPFv3 are represented by a star topology where the Designated Router (DR) is the central point to which all other routers on the broadcast or NBMA network connect. As

a result, routers on the broadcast or NBMA network advertise only their adjacency to the DR. Routers that do not act as DR do not form or advertise adjacencies with each other. They do, however, maintain a 2-Way adjacency state with each other and are directly reachable.

When Segment Routing is used, each router on the broadcast or NBMA network MAY advertise the Adj-SID for its adjacency to the DR using Adj-SID Sub-TLV as described in [Section 7.1](#).

SR capable routers MAY also advertise an Adj-SID for other neighbors (e.g. BDR, DR-OTHER) on the broadcast or NBMA network using the LAN ADJ-SID Sub-TLV as described in [Section 7.2](#).

9. IANA Considerations

This specification updates several existing OSPF registries.

9.1. OSPF Router Information (RI) TLVs Registry

- o 8 (IANA Preallocated) - SR-Algorithm TLV
- o 9 (IANA Preallocated) - SID/Label Range TLV

9.2. OSPFv3 Extend-LSA TLV Registry

Following values are allocated:

- o suggested value 9 - OSPF Extended Prefix Range TLV

9.3. OSPFv3 Extend-LSA Sub-TLV registry

- o suggested value 3 - SID/Label Sub-TLV
- o suggested value 4 - Prefix SID Sub-TLV
- o suggested value 5 - Adj-SID Sub-TLV
- o suggested value 6 - LAN Adj-SID Sub-TLV
- o suggested value 7 - SID/Label Binding Sub-TLV
- o suggested value 8 - ERO Metric Sub-TLV
- o suggested value 9 - IPv4 ERO Sub-TLV
- o suggested value 10 - IPv6 ERO Sub-TLV
- o suggested value 11 - Unnumbered Interface ID ERO Sub-TLV

- o suggested value 12 - IPv4 Backup ERO Sub-TLV
- o suggested value 13 - IPv6 Backup ERO Sub-TLV
- o suggested value 14 - Unnumbered Interface ID Backup ERO Sub-TLV

10. Security Considerations

Implementations must assure that malformed permutations of the newly defined sub-TLVs do not result in errors which cause hard OSPFv3 failures.

11. Acknowledgements

Thanks to Acee Lindem for the detail review of the draft, corrections, as well as discussion about details of the encoding.

We would like to thank Anton Smirnov for his contribution.

Many thanks to Yakov Rekhter, John Drake and Shraddha Hedge for their contribution on earlier incarnations of the "Binding / MPLS Label TLV" in [[I-D.gredler-ospf-label-advertisement](#)].

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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.
- [RFC4970] Lindem, A., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", [RFC 4970](#), July 2007.

12.2. Informative References

[I-D.filsfils-spring-segment-routing-ldp-interop]

Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., and E. Crabbe, "Segment Routing interoperability with LDP", [draft-filsfils-spring-segment-routing-ldp-interop-02](#) (work in progress), September 2014.

[I-D.filsfils-spring-segment-routing-use-cases]

Filsfils, C., Francois, P., Previdi, S., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., Kini, S., and E. Crabbe, "Segment Routing Use Cases", [draft-filsfils-spring-segment-routing-use-cases-01](#) (work in progress), October 2014.

[I-D.gredler-ospf-label-advertisement]

Gredler, H., Amante, S., Scholl, T., and L. Jalil, "Advertising MPLS labels in OSPF", [draft-gredler-ospf-label-advertisement-03](#) (work in progress), May 2013.

[I-D.ietf-ospf-ospfv3-lsa-extend]

Lindem, A., Mirtorabi, S., Roy, A., and F. Baker, "OSPFv3 LSA Extendibility", [draft-ietf-ospf-ospfv3-lsa-extend-06](#) (work in progress), February 2015.

[I-D.ietf-spring-segment-routing]

Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Shakir, R., Tantsura, J., and E. Crabbe, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-01](#) (work in progress), February 2015.

[I-D.minto-rsvp-lsp-egress-fast-protection]

Jeganathan, J., Gredler, H., and Y. Shen, "RSVP-TE LSP egress fast-protection", [draft-minto-rsvp-lsp-egress-fast-protection-03](#) (work in progress), November 2013.

[I-D.previdi-6man-segment-routing-header]

Previdi, S., Filsfils, C., Field, B., and I. Leung, "IPv6 Segment Routing Header (SRH)", [draft-previdi-6man-segment-routing-header-06](#) (work in progress), May 2015.

Authors' Addresses

Peter Psenak (editor)
Cisco Systems, Inc.
Apollo Business Center
Mlynske nivy 43
Bratislava 821 09
Slovakia

Email: ppsenak@cisco.com

Stefano Previdi (editor)
Cisco Systems, Inc.
Via Del Serafico, 200
Rome 00142
Italy

Email: sprevidi@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Brussels
Belgium

Email: cfilsfil@cisco.com

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

Email: hannes@juniper.net

Rob Shakir
British Telecom
London
UK

Email: rob.shakir@bt.com

Wim Henderickx
Alcatel-Lucent
Copernicuslaan 50
Antwerp 2018
BE

Email: wim.henderickx@alcatel-lucent.com

Jeff Tantsura
Ericsson
300 Holger Way
San Jose, CA 95134
US

Email: Jeff.Tantsura@ericsson.com

