

Open Shortest Path First IGP
Internet-Draft
Intended status: Standards Track
Expires: November 5, 2017

P. Psenak, Ed.
S. Previdi, Ed.
C. Filsfils
Cisco Systems, Inc.
H. Gredler
RtBrick Inc.
R. Shakir
Google, Inc.
W. Henderickx
Nokia
J. Tantsura
Individual
May 4, 2017

OSPF Extensions for Segment Routing
draft-ietf-ospf-segment-routing-extensions-13

Abstract

Segment Routing (SR) allows 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 OSPF extensions 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 November 5, 2017.

Copyright Notice

Copyright (c) 2017 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	4
3.	Segment Routing Capabilities	4
3.1.	SR-Algorithm TLV	4
3.2.	SID/Label Range TLV	6
3.3.	SR Local Block Sub-TLV	8
3.4.	SRMS Preference Sub-TLV	9
4.	OSPF Extended Prefix Range TLV	10
5.	Prefix SID Sub-TLV	12
6.	SID/Label Binding Sub-TLV	16
6.1.	ERO Metric Sub-TLV	17
6.2.	ERO Sub-TLVs	18
6.2.1.	IPv4 ERO Sub-TLV	18
6.2.2.	Unnumbered Interface ID ERO Sub-TLV	19
6.2.3.	IPv4 Backup ERO Sub-TLV	21
6.2.4.	Unnumbered Interface ID Backup ERO Sub-TLV	22
7.	Adjacency Segment Identifier (Adj-SID)	23
7.1.	Adj-SID Sub-TLV	23
7.2.	LAN Adj-SID Sub-TLV	25
8.	Elements of Procedure	26
8.1.	Intra-area Segment routing in OSPFv2	26
8.2.	Inter-area Segment routing in OSPFv2	27
8.3.	Segment Routing for External Prefixes	28
8.4.	Advertisement of Adj-SID	29
8.4.1.	Advertisement of Adj-SID on Point-to-Point Links	29
8.4.2.	Adjacency SID on Broadcast or NBMA Interfaces	29
9.	IANA Considerations	29
9.1.	OSPF OSPF Router Information (RI) TLVs Registry	29

9.2.	OSPF Extended Prefix LSA TLV Registry	30
9.3.	OSPF Extended Prefix LSA Sub-TLV Registry	30
9.4.	OSPF Extended Link LSA Sub-TLV Registry	30
10.	Implementation Status	30
11.	Security Considerations	32
12.	Contributors	32
13.	Acknowledgements	32
14.	References	32
14.1.	Normative References	32
14.2.	Informative References	33
	Authors' Addresses	34

[1.](#) Introduction

Segment Routing (SR) allows 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 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 signalling (other than IGP extensions). The IPv6 data plane is out of the scope of this specification - it is not applicable to OSPFv2 which only supports the IPv4 address-family. For example, when used in MPLS networks, SR paths do not require any LDP or RSVP-TE signalling. However, SR can interoperate in the presence of LSPs established with RSVP or LDP.

There are additional segment types, e.g., Binding SID defined in [\[I-D.ietf-spring-segment-routing\]](#).

This draft describes the OSPF 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.

also be advertised. If the SR-Algorithm TLV is not advertised by the node, such node is considered as not being segment routing capable.

An SR Router may use various algorithms when calculating reachability to OSPF routers or prefixes in an OSPF area. Examples of these algorithms are metric based Shortest Path First (SPF), various flavors of Constrained SPF, etc. The SR-Algorithm TLV allows a router to advertise the algorithms currently used by the router to other routers in an OSPF area. The SR-Algorithm 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     |                                     |
|               Type                 |               Length                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Algorithm 1 | Algorithm... | Algorithm n |                               |
+-+-----+-----+-----+-----+-----+-----+-----+-----+
|                                                     |
+
```

where:

Type: TBD, suggested value 8

Length: Variable

Algorithm: Single octet identifying the algorithm. The following values are defined by this document:

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. If the SR-Algorithm TLV is advertised, Algorithm 0 MUST be included.

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 support for Algorithm 1 MUST NOT alter the SPF paths computed by Algorithm 1.

When multiple SR-Algorithm TLVs are received from a given router, the receiver SHOULD use the first occurrence of the TLV in the Router Information LSA. If the SR-Algorithm TLV appears in multiple Router

- 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 inside the Router Information Opaque LSA. The originating router MUST ensure the order is the 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 a SID index.

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) 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 SID/Label Range TLV advertisement, area scope flooding is required.

3.3. SR Local Block Sub-TLV

The SR Local Block (SRLB) Sub-TLV contains the range of labels the node has reserved for local SIDs. Local SIDs are used, e.g., for Adjacency-SIDs, and may also be allocated by components other than the OSPF protocol. As an example, an application or a controller may instruct the router to allocate a specific local SID. Therefore, in order for such applications or controllers to know what local SIDs are available on the router, it is required that the router advertises its SRLB. The SRLB Sub-TLV is used for that purpose.

The SR Local Block (SRLB) Sub-TLV is a top-level TLV of the Router Information Opaque LSA (defined in [\[RFC7770\]](#)).

The SR Local Block Sub-TLV MAY appear multiple times in the Router Information Opaque LSA 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                 |               Length                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     |               Reserved                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     |               Sub-TLVs (variable)         |
+-+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

where:

Type: TBD, suggested value 12

Length: Variable

Range Size: 3-octet SID/label range size (i.e., the number of SIDs or labels in the range including the first SID/label). It MUST be greater than 0.

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.

The originating router MUST NOT advertise overlapping ranges.

Each time a SID from the SRLB is allocated, it SHOULD also be reported to all components (e.g., controller or applications) in order for these components to have an up-to-date view of the current

SRLB allocation. This is required to avoid collisions between allocation instructions.

Within the context of OSPF, the reporting of local SIDs is done through OSPF Sub-TLVs such as the Adjacency-SID ([Section 7](#)). However, the reporting of allocated local SIDs may also be done through other means and protocols which are outside the scope of this document.

A router advertising the SRLB TLV may also have other label ranges, outside of the SRLB, used for its local allocation purposes which are NOT advertised in the SRLB. For example, it is possible that an Adjacency-SID is allocated using a local label that is not part of the SRLB.

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

3.4. SRMS Preference Sub-TLV

The Segment Routing Mapping Server (SRMS) Preference sub-TLV is used to advertise a preference associated with the node that acts as an SR Mapping Server. SRMS preference is defined in [\[I-D.ietf-spring-conflict-resolution\]](#).

The SRMS Preference Sub-TLV is a top-level TLV of the Router Information Opaque LSA (defined in [\[RFC7770\]](#)).

The SRMS Preference Sub-TLV MAY only be advertised once in the Router Information Opaque LSA 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Preference                   | Reserved                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

Type: TBD, suggested value 13

Length: 4 octets

Preference: 1 octet. SRMS preference value from 0 to 255.

When multiple SRMS Preference sub-TLVs are received from a given router, the receiver SHOULD use the first occurrence of the sub-TLV in the Router Information LSA. If the SRMS Preference sub-TLV appears in multiple Router Information LSAs that have different flooding scopes, the SRMS sub-TLV in the Router Information LSA with the narrowest flooding scope SHOULD be used. If the SRMS Preference sub-TLV appears in multiple Router Information LSAs that have the same flooding scope, the SRMS Preference sub-TLV in the Router Information LSA with the numerically smallest Instance ID SHOULD be used and subsequent instances of the SRMS Preference sub-TLV SHOULD be ignored.

The RI LSA can be advertised at any of the defined flooding scopes (link, area, or autonomous system (AS)). For the purpose of the SRMS Preference Sub-TLV advertisement, AS scope flooding is required. If the SRMS advertisements from the SRMS server are only used inside the area to which the SRMS server is attached, area scope flooding may be used.

4. OSPF Extended Prefix Range TLV

In some cases it is useful to advertise attributes for a range of prefixes. The 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 OSPF Extended Prefix Range TLV, which is a top level TLV of the Extended Prefix LSA described in [[RFC7684](#)] is defined for this purpose.

Multiple OSPF Extended Prefix Range TLVs MAY be advertised in each OSPF Extended Prefix Opaque LSA, but all prefix ranges included in a single OSPF Extended Prefix Opaque LSA MUST have the same flooding scope. The OSPF Extended Prefix Range TLV has the following format:

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

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

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

Address Prefix: The prefix, encoded as 32-bit value, padded with zero bits as necessary. The 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 OSPF Extended Prefix TLV described in [RFC7684] and the OSPF Extended Prefix Range TLV described in Section 4. 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                |          Length              |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Flags               |          Reserved             |
|                               |          MT-ID               |
|                               |          Algorithm            |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               SID/Index/Label (variable)     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

Type: TBD, suggested value 2.

Length: Variable

Flags: Single 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 packets to the node that advertised the Prefix-SID.

M-Flag: Mapping Server Flag. If set, the SID was advertised by a Segment Routing Mapping Server 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 the Explicit-NULL label (0 for IPv4) before forwarding the packet.

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.

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.

MT-ID: Multi-Topology ID (as defined in [\[RFC4915\]](#)).

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

A router receiving a Prefix-SID from a remote node and with an algorithm value that such remote node has not advertised in the SR-Algorithm sub-TLV ([Section 3.1](#)) MUST ignore the Prefix-SID sub-TLV.

SID/Index/Label: According to the V and L flags, it contains either:

- A 32-bit index defining the offset in the SID/Label space advertised by this router.

- A 24-bit label where the 20 rightmost bits are used for encoding the label value.

If multiple Prefix-SIDs are advertised for the same prefix, the receiving router MUST use the first encoded SID and MAY use 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 the E and P flags advertised by the next-hop router if that 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, the NP-Flag SHOULD NOT be set.

The NP-Flag (No-PHP) MUST be set for Prefix-SIDs allocated to redistributed prefixes, unless the redistributed prefix is directly attached to the 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 not set 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 Area Border Router (prefix propagation from one area to another) or at an AS Boundary 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 an Explicit-NULL label. 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 the M-Flag is set, the NP-flag and the E-flag MUST be ignored at reception.

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 SHOULD be done in following cases:

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

The Prefix is inter-area type and downstream neighbor is an ABR, which is advertising prefix reachability and is also generating the Extended Prefix TLV with the A-flag set for this prefix as described in [section 2.1 of \[RFC7684\]](#).

The Prefix is external type and downstream neighbor is an ASBR, which is advertising prefix reachability and is also generating the Extended Prefix TLV with the A-flag set for this prefix as described in [section 2.1 of \[RFC7684\]](#).

When a Prefix-SID is advertised in an Extended Prefix Range TLV, then the value advertised in the 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: 192.0.2.1/32, Prefix-SID: Index 1
Router-B: 192.0.2.2/32, Prefix-SID: Index 2
Router-C: 192.0.2.3/32, Prefix-SID: Index 3
Router-D: 192.0.2.4/32, Prefix-SID: Index 4

then the Prefix field in the Extended Prefix Range TLV would be set to 192.0.2.1, Prefix Length would be set to 32, 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:

10.1.1/24, Prefix-SID: Index 51
10.1.2/24, Prefix-SID: Index 52
10.1.3/24, Prefix-SID: Index 53
10.1.4/24, Prefix-SID: Index 54
10.1.5/24, Prefix-SID: Index 55
10.1.6/24, Prefix-SID: Index 56
10.1.7/24, Prefix-SID: Index 57

then the Prefix field in the Extended Prefix Range TLV would be set to 10.1.1.0, Prefix Length would be set to 24, Range Size would be 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 a SID/Label mapping for a path to the a prefix.

The SID/Label Binding Sub-TLV MAY be originated by any router in an OSPF 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 in a simple path description language. Analogous 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. Additionally, SID/Label Bindings from external protocols can be easily re-advertised.

The SID/Label Binding Sub-TLV may be used for advertising SID/Label Bindings and their associated Primary and Backup paths. In a single TLV, a primary ERO Path, backup ERO Path, or both can be 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.

The SID/Label Binding Sub-TLV is a Sub-TLV of the OSPF Extended Prefix TLV described in [[RFC7684](#)] and the OSPF Extended Prefix Range TLV described in [Section 4](#). Multiple SID/Label Binding TLVs can be present in their parent TLV. 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
|      Type                    |      Length                  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Flags                   |      Reserved                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |      Sub-TLVs (variable)     |
+-+-----+-----+-----+-----+-----+-----+-----+---+
|                               |                               |

```

where:

Type: TBD, suggested value 3

Length: Variable

Flags: Single octet field containing the following flags:

```

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

```

where:

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

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

MT-ID: Multi-Topology ID (as defined in [\[RFC4915\]](#)).

Weight: 8 bits of weight used for load-balancing purposes. The use of the weight is defined in [\[I-D.ietf-spring-segment-routing\]](#).

The SID/Label Binding Sub-TLV supports the 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 SHOULD only appear once. If the SID/Label Sub-TLV is not included in the SID/Label Binding Sub-TLV, the SID/Label Binding Sub-TLV MUST be ignored. If the SID/Label Sub-TLV appears in the SID/Label Binding Sub-TLV more than once, instances other than the first SHOULD be ignored and the condition SHOULD be logged for possible action by the network operator.

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

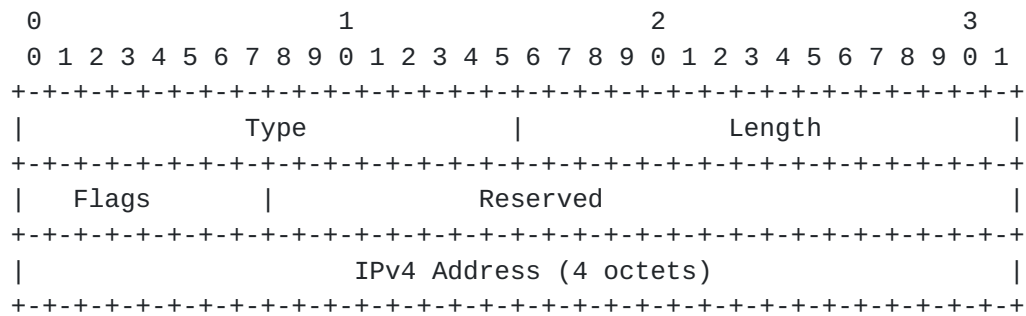
ERO Sub-TLVs as defined in [Section 6.2](#).

[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. ERO Metric Sub-TLV is an option sub-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

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




IPv4 ERO Sub-TLV format

where:

Type: TBD, suggested value 4

Length: 8 octets

Flags: Single octet field containing the following flags:



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'. The terms 'loose' and 'strict' are defined for RSVP subobjects in [\[RFC3209\]](#).

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

IPv4 Address - The address of the explicit route hop.

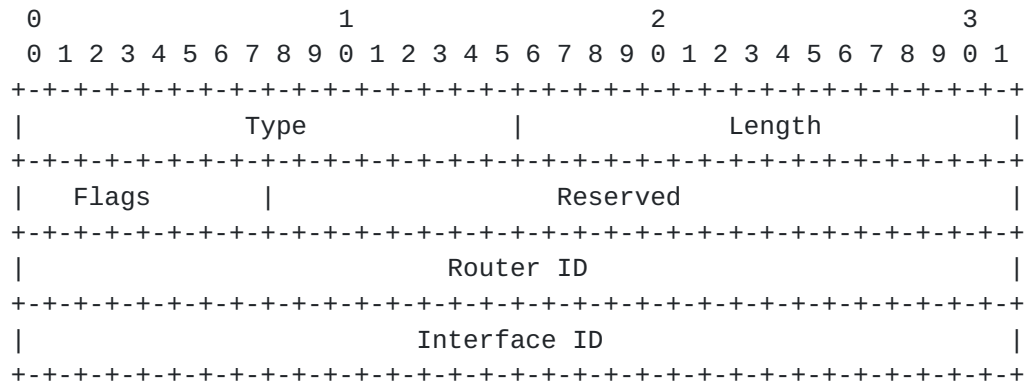
6.2.2.2. Unnumbered Interface ID ERO Sub-TLV

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 includes an unnumbered interface. Unnumbered interfaces are referenced using the interface index. Interface indices are assigned

local to the router and therefore are 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.




where:

Unnumbered Interface ID ERO Sub-TLV format

Type: TBD, suggested value 5

Length: 12 octets

Flags: Single octet field containing the following flags:



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'. The terms 'loose' and 'strict' are defined for RSVP subobjects in [[RFC3209](#)]

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

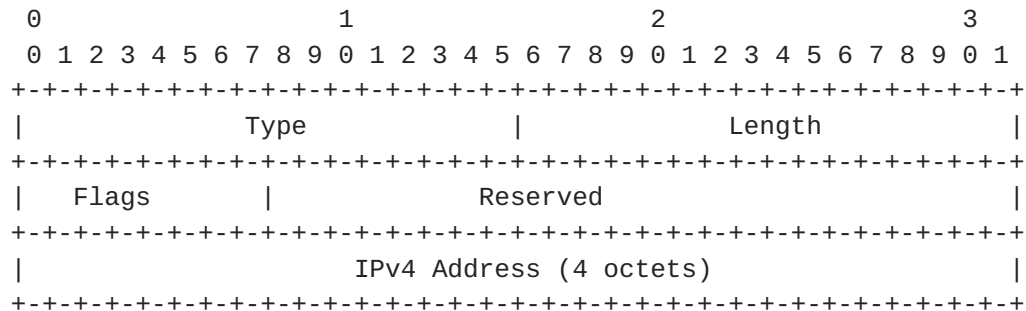
Router ID: Router ID of the next-hop.

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

6.2.3. 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 [\[RFC3209\]](#).



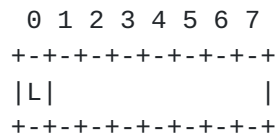
IPv4 Backup ERO Sub-TLV format

where:

Type: TBD, suggested value 6

Length: 8 octets

Flags: Single octet field containing the following flags:



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'. The terms 'loose' and 'strict' are defined for RSVP subobjects in [\[RFC3209\]](#)

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

IPv4 Address - The address of the explicit route hop.

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

Router ID: Router ID of the next-hop.

Interface ID: 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

Adj-SID is an optional Sub-TLV of the Extended Link TLV defined in [RFC7684]. It MAY appear multiple times in the Extended 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:

```

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             |
|          MT-ID                |          Weight               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
|          SID/Label/Index (variable)                           |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

where:

Type: TBD, suggested value 2.

Length: Variable.

Flags: Single octet field containing the following flags:

```

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

```

where:

B-Flag: Backup Flag. If set, the Adj-SID refers to an adjacency that is eligible for protection (e.g., using IPFRR or MPLS-FRR) as described in section 3.5 of [\[I-D.ietf-spring-segment-routing\]](#).

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

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

The G-Flag: Group Flag. When set, the G-Flag indicates that the Adj-SID refers to a group of adjacencies (and therefore MAY be assigned to other adjacencies as well).

P-Flag. Persistent flag. When set, the P-Flag indicates that the Adj-SID is persistently allocated, i.e., the Adj-SID value remains consistent across router restart and/or interface flap.

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

MT-ID: Multi-Topology ID (as defined in [\[RFC4915\]](#)).

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

SID/Index/Label: According to the V and L flags, it contains either:

A 32-bit index defining the offset in the SID/Label space advertised by this router.

A 24-bit label where the 20 rightmost bits are used for encoding the label value.

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 [\[I-D.ietf-spring-segment-routing\]](#).

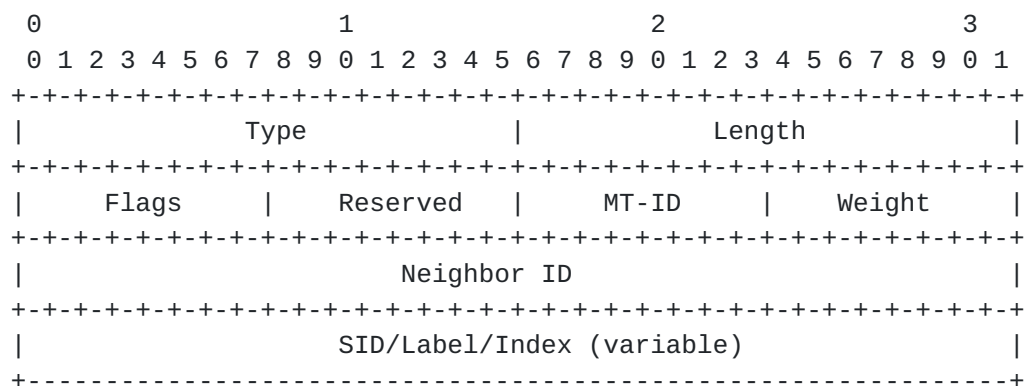
An SR capable router MAY allocate more than one Adj-SID to an adjacency

An SR capable router MAY allocate the same Adj-SID to different adjacencies

When the P-flag is not set, the Adj-SID MAY be persistent. When the P-flag is set, the Adj-SID MUST be persistent.

7.2. LAN Adj-SID Sub-TLV

LAN Adj-SID is an optional Sub-TLV of the Extended Link TLV defined in [RFC7684]. It MAY appear multiple times in the Extended-Link TLV. It is used to advertise a SID/Label for an adjacency to a non-DR router on a broadcast, NBMA, or hybrid [RFC6845] network.




where:

Type: TBD, suggested value 3.

Length: Variable.

Flags: Single octet field containing the following flags:



where:

B-Flag: Backup-flag. If set, the LAN-Adj-SID refers to an adjacency that is eligible for protection (e.g.: using IPFRR or MPLS-FRR) as described in section 3.5 of [I-D.ietf-spring-segment-routing].

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.

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

P-Flag. Persistent flag. When set, the P-Flag indicates that the Adj-SID is persistently allocated, i.e., the Adj-SID value remains consistent across router restart and/or interface flap.

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

MT-ID: Multi-Topology ID (as defined in [\[RFC4915\]](#)).

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

Neighbor ID: The Router ID of the neighbor for which the LAN-Adj-SID is advertised.

SID/Index/Label: According to the V and L flags, it contains either:

A 32-bit index defining the offset in the SID/Label space advertised by this router.

A 24-bit label where the 20 rightmost bits are used for encoding the label value.

When the P-flag is not set, the Adj-SID MAY be persistent. When the P-flag is set, the Adj-SID MUST be persistent.

8. Elements of Procedure

8.1. Intra-area Segment routing in OSPFv2

An OSPFv2 router that supports segment routing MAY advertise Prefix-SIDs for any prefix to which it is advertising reachability (e.g., a 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 OSPFv2 routing domain.

A Prefix-SID can also be advertised by the SR Mapping Servers (as described in [[I-D.filsfils-spring-segment-routing-ldp-interop](#)]). A Mapping Server advertises Prefix-SIDs for remote prefixes that exist in the OSPFv2 routing domain. Multiple Mapping Servers can advertise Prefix-SIDs for the same prefix, in which case the same Prefix-SID MUST be advertised by all of them. The flooding scope of the OSPF Extended Prefix Opaque LSA that is generated by the SR Mapping Server could be either area scoped or AS scoped and is determined based on the configuration of the SR Mapping Server.

An SR Mapping Server MUST use the 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 an SR Mapping Server.

Area-scoped OSPF Extended Prefix Range TLV are propagated between areas. Similar to propagation of prefixes between areas, an 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 are described in [Section 4](#).

When propagating an OSPF Extended Prefix Range TLV between areas, ABRs 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](#).

8.2. Inter-area Segment routing in OSPFv2

In order to support SR in a multi-area environment, OSPFv2 must propagate Prefix-SID information between areas. The following procedure is used to propagate Prefix SIDs between areas.

When an OSPF ABR advertises a Type-3 Summary LSA from an intra-area prefix to all its connected areas, it will also originate an Extended Prefix Opaque LSA, as described in [[RFC7684](#)]. The flooding scope of the Extended Prefix Opaque LSA type will be set to area-scope. The route-type in the OSPF Extended Prefix TLV is set to inter-area. The Prefix-SID Sub-TLV will be included in this LSA and 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 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 when propagating the Prefix-SID for the prefix to other areas.

When an OSPF ABR advertises Type-3 Summary LSAs from an inter-area route to all its connected areas, it will also originate an Extended Prefix Opaque LSA, as described in [RFC7684]. The flooding scope of the Extended Prefix Opaque LSA type will be set to area-scope. The route-type in OSPF Extended Prefix TLV is set to inter-area. The Prefix-SID Sub-TLV will be included in this LSA and the Prefix-SID will be set as follows:

The ABR will look at its best path to the prefix in the backbone area and find 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 backbone 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 the Prefix-SID for the prefix to other areas.

8.3. Segment Routing for External Prefixes

Type-5 LSAs are flooded domain wide. When an ASBR, which supports SR, generates Type-5 LSAs, it should also originate Extended Prefix Opaque LSAs, as described in [RFC7684]. The flooding scope of the Extended Prefix Opaque LSA type is set to AS-scope. The route-type in the OSPF Extended Prefix TLV is set to external. The Prefix-SID Sub-TLV is included in this LSA and the Prefix-SID value will be set to the SID that has been reserved for that prefix.

When an NSSA ABR translates Type-7 LSAs into Type-5 LSAs, it should also advertise the Prefix-SID for the prefix. The NSSA ABR determines its best path to the prefix advertised in the translated Type-7 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 for the Type-5 prefix. 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 a P2P link that is in neighbor 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 withdrawn from the area.

8.4.2. Adjacency SID on Broadcast or NBMA Interfaces

Broadcast, NBMA, or hybrid [[RFC6845](#)] networks in OSPF are represented by a star topology where the Designated Router (DR) is the central point to which all other routers on the broadcast, NBMA, or hybrid network connect. As a result, routers on the broadcast, NBMA, or hybrid 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 2-Way adjacency state with each other and are directly reachable.

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

SR capable routers MAY also advertise a LAN-Adj-SID for other neighbors (e.g., BDR, DR-OTHER) on the broadcast, NBMA, or hybrid 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
- o 12 - SR Local Block Sub-TLV
- o 13 - SRMS Preference Sub-TLV

9.2. OSPF Extended Prefix LSA TLV Registry

Following values are allocated:

- o 2 - OSPF Extended Prefix Range TLV

9.3. OSPF Extended Prefix LSA Sub-TLV Registry

Following values are allocated:

- o 1 - SID/Label Sub-TLV
- o 2 - Prefix SID Sub-TLV
- o 3 - SID/Label Binding Sub-TLV
- o 4 - IPv4 ERO Sub-TLV
- o 5 - Unnumbered Interface ID ERO Sub-TLV
- o 6 - IPv4 Backup ERO Sub-TLV
- o 7 - Unnumbered Interface ID Backup ERO Sub-TLV
- o 8 - ERO Metric Sub-TLV

9.4. OSPF Extended Link LSA Sub-TLV Registry

Following initial values are allocated:

- o 1 - SID/Label Sub-TLV
- o 2 - Adj-SID Sub-TLV
- o 3 - LAN Adj-SID/Label Sub-TLV

10. Implementation Status

An implementation survey with seven questions related to the implementer's support of OSPFv2 Segment Routing was sent to the OSPF WG list and several known implementers. This section contains responses from three implementers who completed the survey. No external means were used to verify the accuracy of the information submitted by the respondents. The respondents are considered experts on the products they reported on. Additionally, responses were omitted from implementers who indicated that they have not implemented the function yet.

Responses from Nokia (former Alcatel-Lucent):

Link to a web page describing the implementation:

https://infoproducts.alcatel-lucent.com/cgi-bin/dbaccessfilename.cgi/3HE10799AAAATQZZA01_V1_7450%20ESS%207750%20SR%20and%207950%20XRS%20Unicast%20Routing%20Protocols%20Guide%20R14.0.R1.pdf

The implementation's level of maturity: Production.

Coverage: We have implemented all sections and have support for the latest draft.

Licensing: Part of the software package that needs to be purchased.

Implementation experience: Great spec. We also performed inter-operability testing with Cisco's OSPF Segment Routing implementation.

Contact information: wim.henderickx@nokia.com

Responses from Cisco Systems:

Link to a web page describing the implementation:

www.segment-routing.net/home/tutorial

The implementation's level of maturity: Production.

Coverage: All sections, except the [section 6](#) (SID/Label Binding Sub-TLV) have been implemented according to the latest draft.

Licensing: Part of a commercial software package.

Implementation experience: Many aspects of the draft are result of the actual implementation experience, as the draft evolved from its initial version to the current one. Interoperability testing with Alcatel-Lucent was performed, which confirmed the draft's ability to serve as a reference for the implementors.

Contact information: ppsenak@cisco.com

Responses from Juniper:

The implementation's name and/or a link to a web page describing the implementation:

Feature name is OSPF SPRING

The implementation's level of maturity: To be released in 16.2 (second half of 2016)

Coverage: All sections implemented except Sections [4](#), and [6](#).

Licensing: JUNOS Licensing needed.

Implementation experience: NA

Contact information: shraddha@juniper.net

[11.](#) Security Considerations

Implementations must assure that malformed TLV and Sub-TLV permutations do not result in errors which cause hard OSPF failures.

[12.](#) Contributors

The following people gave a substantial contribution to the content of this document: Acee Lindem, Ahmed Bashandy, Martin Horneffer, Bruno Decraene, Stephane Litkowski, Igor Milojevic, Rob Shakir and Saku Ytti.

[13.](#) Acknowledgements

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 definition of the "Binding / MPLS Label TLV".

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

[14.](#) References

[14.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/RFC2328, April 1998, <<http://www.rfc-editor.org/info/rfc2328>>.

- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001, <<http://www.rfc-editor.org/info/rfc3209>>.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", [RFC 3477](#), DOI 10.17487/RFC3477, January 2003, <<http://www.rfc-editor.org/info/rfc3477>>.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), DOI 10.17487/RFC3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", [RFC 4915](#), DOI 10.17487/RFC4915, June 2007, <<http://www.rfc-editor.org/info/rfc4915>>.
- [RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", [RFC 5250](#), DOI 10.17487/RFC5250, July 2008, <<http://www.rfc-editor.org/info/rfc5250>>.
- [RFC6845] Sheth, N., Wang, L., and J. Zhang, "OSPF Hybrid Broadcast and Point-to-Multipoint Interface Type", [RFC 6845](#), DOI 10.17487/RFC6845, January 2013, <<http://www.rfc-editor.org/info/rfc6845>>.
- [RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", [RFC 7684](#), DOI 10.17487/RFC7684, November 2015, <<http://www.rfc-editor.org/info/rfc7684>>.
- [RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", [RFC 7770](#), DOI 10.17487/RFC7770, February 2016, <<http://www.rfc-editor.org/info/rfc7770>>.

[14.2. 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.ietf-spring-conflict-resolution]

Ginsberg, L., Psenak, P., Previdi, S., and M. Pilka, "Segment Routing Conflict Resolution", [draft-ietf-spring-conflict-resolution-01](#) (work in progress), June 2016.

[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-00](#) (work in progress), December 2014.

[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.

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
RtBrick Inc.

Email: hannes@rtbrick.com

Rob Shakir
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: robjs@google.com

Wim Henderickx
Nokia
Copernicuslaan 50
Antwerp 2018
BE

Email: wim.henderickx@nokia.com

Jeff Tantsura
Individual
US

Email: jefftant.ietf@gmail.com

