

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: April 25, 2019

S. Previdi
Huawei Technologies
K. Talaulikar
C. Filsfils
Cisco Systems, Inc.
H. Gredler
RtBrick Inc.
M. Chen
Huawei Technologies
October 22, 2018

BGP Link-State extensions for Segment Routing
draft-ietf-idr-bgp-ls-segment-routing-ext-11

Abstract

Segment Routing (SR) allows for a flexible definition of end-to-end paths by encoding paths as sequences of topological sub-paths, called "segments". These segments are advertised by routing protocols e.g. by the link state routing protocols (IS-IS, OSPFv2 and OSPFv3) within IGP topologies.

This draft defines extensions to the BGP Link-state address-family in order to carry segment routing information via BGP.

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 <https://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 April 25, 2019.

Internet-Draft BGP LS extensions for Segment Routing October 2018

Copyright Notice

Copyright (c) 2018 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 (<https://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.	BGP-LS Extensions for Segment Routing	5
2.1.	Node Attributes TLVs	5
2.1.1.	SID/Label Sub-TLV	6
2.1.2.	SR-Capabilities TLV	7
2.1.3.	SR-Algorithm TLV	8
2.1.4.	SR Local Block TLV	8
2.1.5.	SRMS Preference TLV	9
2.2.	Link Attribute TLVs	10
2.2.1.	Adjacency SID TLV	11
2.2.2.	LAN Adjacency SID TLV	12
2.2.3.	L2 Bundle Member	13
2.3.	Prefix Attribute TLVs	14
2.3.1.	Prefix-SID TLV	15
2.3.2.	Prefix Attribute Flags TLV	16
2.3.3.	Source Router Identifier (Source Router-ID) TLV	17
2.3.4.	Range TLV	17
2.4.	Equivalent IS-IS Segment Routing TLVs/Sub-TLVs	19
2.5.	Equivalent OSPFv2/OSPFv3 Segment Routing TLVs/Sub-TLVs	20
3.	Implementation Status	21
4.	IANA Considerations	21
4.1.	TLV/Sub-TLV Code Points Summary	21
5.	Manageability Considerations	22
6.	Security Considerations	23
7.	Contributors	23
8.	Acknowledgements	24

9.	References	24
9.1.	Normative References	24
9.2.	Informative References	25
9.3.	URIs	26
	Authors' Addresses	28

[1.](#) Introduction

Segment Routing (SR) allows for a flexible definition of end-to-end paths by combining sub-paths called "segments". A segment can represent any instruction, topological or service-based. A segment can have a local semantic to an SR node or global within a domain. Within IGP topologies an SR path is encoded as a sequence of topological sub-paths, called "IGP segments". These segments are advertised by the link-state routing protocols (IS-IS, OSPFv2 and OSPFv3).

Two types of IGP segments are defined, Prefix segments and Adjacency segments. Prefix segments, by default, represent an ECMP-aware shortest-path to a prefix, 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. [\[RFC8402\]](#).

When Segment Routing is enabled in a IGP domain, segments are advertised in the form of Segment Identifiers (SIDs). The IGP link-state routing protocols have been extended to advertise SIDs and other SR-related information. IGP extensions are described in: IS-IS [\[I-D.ietf-isis-segment-routing-extensions\]](#), OSPFv2 [\[I-D.ietf-ospf-segment-routing-extensions\]](#) and OSPFv3 [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#). Using these extensions, Segment Routing can be enabled within an IGP domain.

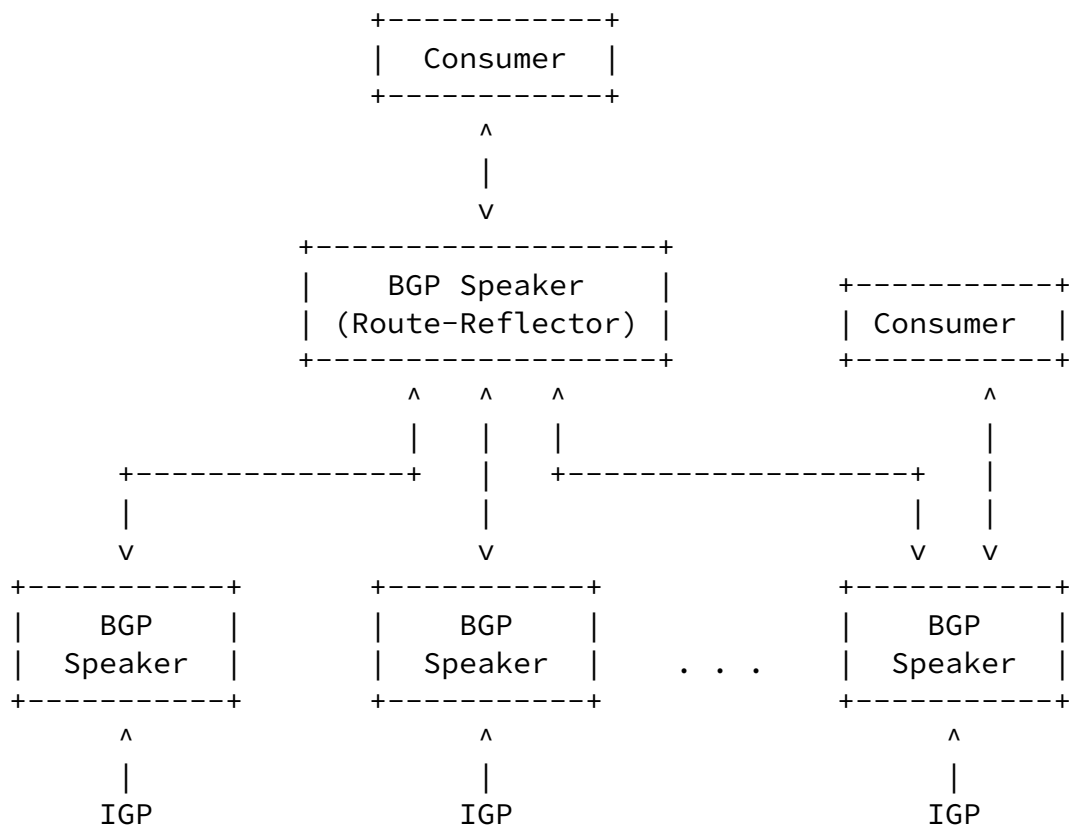


Figure 1: Link State info collection

Segment Routing (SR) allows advertisement of single or multi-hop paths. The flooding scope for the IGP extensions for Segment routing is IGP area-wide. Consequently, the contents of a Link State Database (LSDB) or a Traffic Engineering Database (TED) has the scope

of an IGP area and therefore, by using the IGP alone it is not enough to construct segments across multiple IGP Area or AS boundaries.

In order to address the need for applications that require topological visibility across IGP areas, or even across Autonomous Systems (AS), the BGP-LS address-family/sub-address-family have been defined to allow BGP to carry Link-State information. The BGP Network Layer Reachability Information (NLRI) encoding format for BGP-LS and a new BGP Path Attribute called the BGP-LS attribute are defined in [[RFC7752](#)]. The identifying key of each Link-State object, namely a node, link, or prefix, is encoded in the NLRI and the properties of the object are encoded in the BGP-LS attribute. Figure 1 describes a typical deployment scenario. In each IGP area, one or more nodes are configured with BGP-LS. These BGP speakers form an IBGP mesh by connecting to one or more route-reflectors. This way, all BGP speakers (specifically the route-reflectors) obtain Link-State information from all IGP areas (and from other ASes from EBGP peers). An external component connects to the route-reflector

to obtain this information (perhaps moderated by a policy regarding what information is or isn't advertised to the external component).

This document describes extensions to BGP-LS to advertise the SR information. An external component (e.g., a controller) then can collect SR information from across an SR domain and construct the end-to-end path (with its associated SIDs) that need to be applied to an incoming packet to achieve the desired end-to-end forwarding. Here the SR domain is defined as a single administrative domain that may be comprised of a single AS or multiple ASes under consolidated global SID administration.

[2.](#) BGP-LS Extensions for Segment Routing

This document defines SR extensions to BGP-LS and specifies the TLVs and sub-TLVs for advertising SR information within the BGP-LS Attribute. [Section 2.4](#) and [Section 2.5](#) illustrates the equivalent TLVs and sub-TLVs in IS-IS, OSPFv2 and OSPFv3 protocols.

BGP-LS [[RFC7752](#)] defines the BGP-LS NLRI that can be a Node NLRI, a Link NLRI or a Prefix NLRI. BGP-LS [[RFC7752](#)] defines the TLVs that map link-state information to BGP-LS NLRI within the BGP-LS

Attribute. This document adds additional BGP-LS Attribute TLVs in order to encode SR information. It does not introduce any changes to the encoding of the BGP-LS NLRIs.

Some of the TLVs defined in this document contain fields (e.g. flags) whose semantics need to be interpreted accordingly to the respective underlying IS-IS, OSPFv2 or OSPFv3 protocol. The receiver of the BGP-LS update for any of the NLRIs MUST check the Protocol-ID of the NLRI and refer to the underlying protocol specification in order to parse such fields. The individual field descriptions in the sub-sections below point to the relevant underlying protocol specifications for such fields.

[2.1.](#) Node Attributes TLVs

The following Node Attribute TLVs are defined:

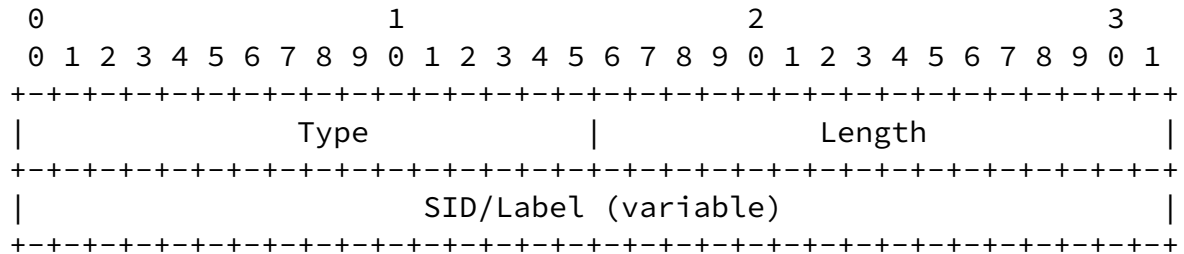
Description	Length	Section
SID/Label	variable	Section 2.1.1
SR Capabilities	variable	Section 2.1.2
SR Algorithm	variable	Section 2.1.3
SR Local Block	variable	Section 2.1.4
SRMS Preference	variable	Section 2.1.5

Table 1: Node Attribute TLVs

These TLVs can ONLY be added to the BGP-LS Attribute associated with the Node NLRI that originates the corresponding underlying IGP TLV/sub-TLV described below.

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

The SID/Label TLV is used as sub-TLV by the SR-Capabilities ([Section 2.1.2](#)) and SRLB ([Section 2.1.4](#)) TLVs and has the following format:



where:

Type: TBD, see [Section 4](#).

Length: Variable, 3 or 4.

SID/Label: If length is set to 3, then the 20 rightmost bits represent a label (the total TLV size is 7). If length is set to 4, then the value represents a 32 bit SID (the total TLV size is 8).

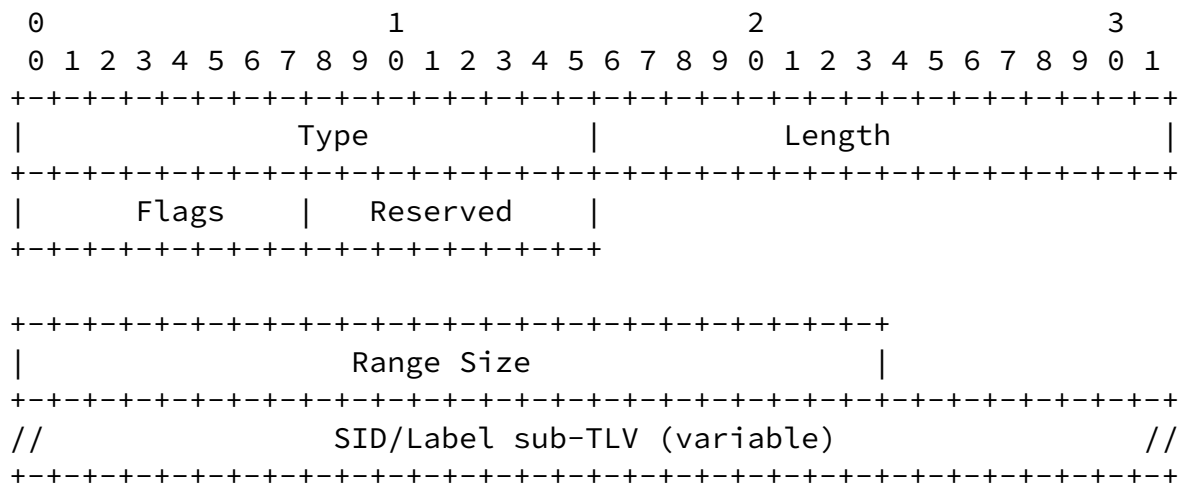
The receiving router MUST ignore the SID/Label sub-TLV if the length is other than 3 or 4.

[2.1.2.](#) SR-Capabilities TLV

The SR-Capabilities TLV is used in order to advertise the node's SR Capabilities including its Segment Routing Global Base (SRGB) range(s). In the case of IS-IS, the capabilities also include the IPv4 and IPv6 support for SR-MPLS forwarding plane. This information is derived from the protocol specific advertisements.

- o IS-IS, as defined by the SR-Capabilities sub-TLV in [\[I-D.ietf-isis-segment-routing-extensions\]](#).
- o OSPFv2/OSPFv3, as defined by the SID/Label Range TLV in [\[I-D.ietf-ospf-segment-routing-extensions\]](#) and [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#).

The SR Capabilities TLV has following format:



Type: TBD, see [Section 4](#).

Length: Variable.

Flags: 1 octet of flags as defined in [\[I-D.ietf-isis-segment-routing-extensions\]](#).

Reserved: 1 octet that SHOULD be set to 0 and MUST be ignored on receipt.

One or more entries, each of which have the following format:

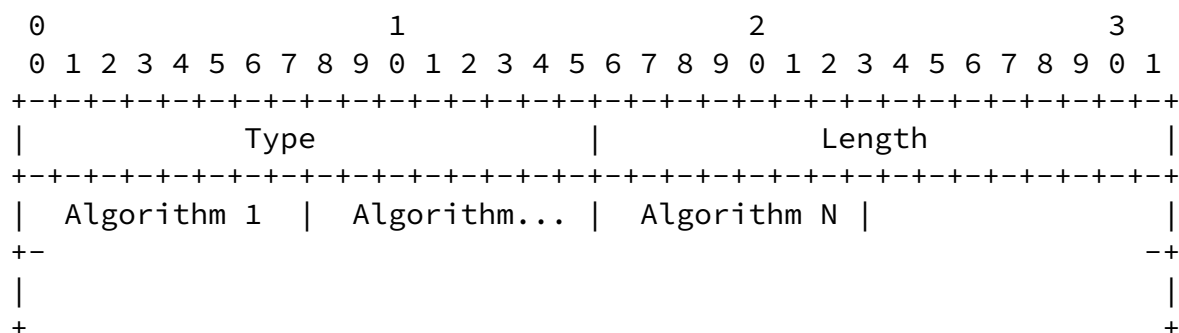
Range Size: 3 octet value indicating the number of labels in the range.

SID/Label sub-TLV (as defined in [Section 2.1.1](#)) which encodes the first label in the range.

The SR-Algorithm TLV is used in order to advertise the SR Algorithms supported by the node. This information is derived from the protocol specific advertisements.

- o IS-IS, as defined by the SR-Algorithm sub-TLV in [\[I-D.ietf-isis-segment-routing-extensions\]](#).
- o OSPFv2/OSPFv3, as defined by the SR-Algorithm TLV in [\[I-D.ietf-ospf-segment-routing-extensions\]](#) and [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#).

The SR-Algorithm TLV has the following format:



where:

Type: TBD, see [Section 4](#).

Length: Variable.

Algorithm: 1 octet identifying the algorithm.

[2.1.4](#). SR Local Block TLV

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

This information is derived from the protocol specific advertisements.

- The SRLB TLV has the following format:

The Segment Routing Mapping Server (SRMS) Preference TLV is used in order to associate a preference with SRMS advertisements from a particular source. [[I-D.ietf-spring-segment-routing-ldp-interop](#)] specifies the SRMS functionality along with SRMS preference of the

This information is derived from the protocol specific advertisements.

- o IS-IS, as defined by the SRMS Preference sub-TLV in [\[I-D.ietf-isis-segment-routing-extensions\]](#).
- o OSPFv2/OSPFv3, as defined by the SRMS Preference TLV in [\[I-D.ietf-ospf-segment-routing-extensions\]](#) and [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#).

The SRMS Preference TLV has following format:

[illegible]

Type: TBD, see [Section 4](#).

Length: 1.

Preference: 1 octet. Unsigned 8 bit SRMS preference.

The use of the SRMS Preference TLV is defined in [\[I-D.ietf-isis-segment-routing-extensions\]](#), [\[I-D.ietf-ospf-segment-routing-extensions\]](#) and [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#).

2.2. Link Attribute TLVs

The following Link Attribute TLVs are defined:

Description	Length	Section
Adjacency Segment Identifier (Adj-SID)	variable	Section 2.2.1

TLV			
LAN Adjacency Segment Identifier (Adj-	variable	Section 2.2.2	
SID) TLV			
L2 Bundle Member TLV	variable	Section 2.2.3	
+-----+	+-----+	+-----+	+-----+

Table 2: Link Attribute TLVs

These TLVs can ONLY be added to the BGP-LS Attribute associated with the Link NLRI whose local node originates the corresponding underlying IGP TLV/sub-TLV described below.

For a LAN, normally a node only announces its adjacency to the IS-IS pseudo-node (or the equivalent OSPF Designated and Backup Designated Routers)[[I-D.ietf-isis-segment-routing-extensions](#)]. The LAN Adjacency Segment TLV allows a node to announce adjacencies to all other nodes attached to the LAN in a single instance of the BGP-LS Link NLRI. Without this TLV, the corresponding BGP-LS link NLRI would need to be originated for each additional adjacency in order to advertise the SR TLVs for these neighbor adjacencies.

[2.2.1.](#) Adjacency SID TLV

The Adjacency SID (Adj-SID) 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								
+-----																																							

where:

Type: TBD, see [Section 4](#).

Length: Variable, 7 or 8 depending on Label or Index encoding of

Length: Variable. For IS-IS it would be 13 or 14 depending on Label or Index encoding of the SID. For OSPF it would be 11 or 12 depending on Label or Index encoding of the SID.

Flags. 1 octet field of following flags as defined in [\[I-D.ietf-isis-segment-routing-extensions\]](#), [\[I-D.ietf-ospf-segment-routing-extensions\]](#) and [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#).

Weight: Weight used for load-balancing purposes.

Reserved: 2 octets that SHOULD be set to 0 and MUST be ignored on receipt.

SID/Index/Label: Label or index value depending on the flags setting as defined in [[I-D.ietf-isis-segment-routing-extensions](#)], [[I-D.ietf-ospf-segment-routing-extensions](#)] and [[I-D.ietf-ospf-ospfv3-segment-routing-extensions](#)].

2.2.3. L2 Bundle Member

The L2 Bundle Member Attribute TLV identifies an L2 Bundle Member link which in turn is associated with a parent L3 link. The L3 link is described by the Link NLRI defined in [\[RFC7752\]](#) and the L2 Bundle Member Attribute TLV is associated with the Link NLRI. The TLV MAY include sub-TLVs which describe attributes associated with the bundle member. The identified bundle member represents a unidirectional path from the originating router to the neighbor specified in the parent L3 Link. Multiple L2 Bundle Member Attribute TLVs MAY be associated with a Link NLRI.

The L2 Bundle Member Attribute TLV has the following format:

Type																Length															
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

```

+-----+
|
|               L2 Bundle Member Descriptor               |
|
+-----+
//               Link attribute sub-TLVs(variable)         //
+-----+

```

where:

Type: TBD, see [Section 4](#).

Length: Variable.

L2 Bundle Member Descriptor: A Link Local Identifier as defined in [\[RFC4202\]](#).

Link attributes for L2 Bundle Member Links are advertised as sub-TLVs of the L2Bundle Member Attribute TLV. The sub-TLVs are identical to existing BGP-LS TLVs as identified in the table below.

TLV Code Point	Description	Reference Document
1088	Administrative group (color)	[RFC7752]
1089	Maximum link bandwidth	[RFC7752]
1090	Max. reservable link bandwidth	[RFC7752]
1091	Unreserved bandwidth	[RFC7752]
1092	TE default metric	[RFC7752]

1093	Link protection type	[RFC7752]
1099	Adjacency Segment Identifier (Adj-SID) TLV	Section 2.2.1
1100	LAN Adjacency Segment Identifier (Adj-SID) TLV	Section 2.2.2
1114	Unidirectional link delay	[I-D.ietf-idr-te-pm-bgp]
1115	Min/Max Unidirectional link delay	[I-D.ietf-idr-te-pm-bgp]
1116	Unidirectional Delay Variation	[I-D.ietf-idr-te-pm-bgp]
1117	Unidirectional packet loss	[I-D.ietf-idr-te-pm-bgp]
1118	Unidirectional residual bandwidth	[I-D.ietf-idr-te-pm-bgp]
1119	Unidirectional available bandwidth	[I-D.ietf-idr-te-pm-bgp]
1120	Unidirectional bandwidth utilization	[I-D.ietf-idr-te-pm-bgp]

Table 3: BGP-LS Attribute TLVs also used as sub-TLVs of L2 Bundle Member Attribute TLV

[2.3.](#) Prefix Attribute TLVs

The following Prefix Attribute TLVs are defined:

Description	Length	Section
Prefix SID	variable	Section 2.3.1
Range	variable	Section 2.3.4
Prefix Attribute Flags	variable	Section 2.3.2
Source Router-ID	variable	Section 2.3.3

Table 4: Prefix Attribute TLVs

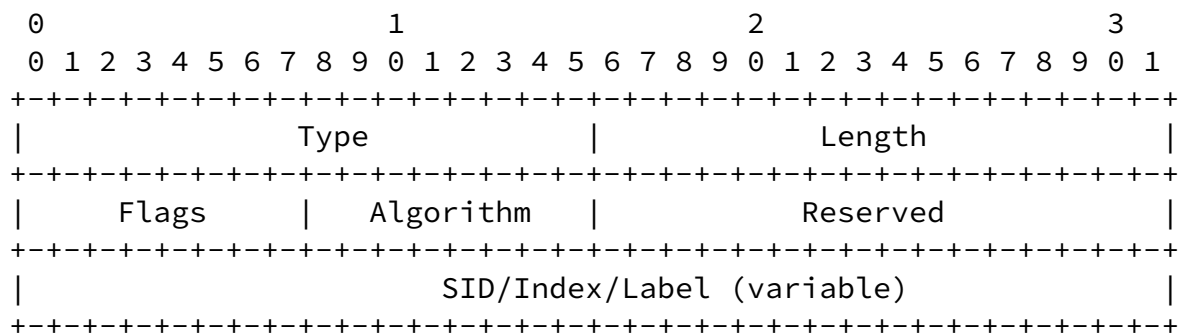
These TLVs can ONLY be added to the BGP-LS Attribute associated with the Prefix NLRI whose local node originates the corresponding underlying IGP TLV/sub-TLV described below.

2.3.1. Prefix-SID TLV

The Prefix-SID TLV is used in order to advertise information related to a Prefix-SID. This information is originated in:

- o IS-IS, as defined by the Prefix-SID TLV in [\[I-D.ietf-isis-segment-routing-extensions\]](#).
- o OSPFv2 and OSPFv3, as defined by the Prefix-SID TLV in [\[I-D.ietf-ospf-segment-routing-extensions\]](#) and [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#) respectively.

The Prefix-SID has the following format:



where:

Type: TBD, see [Section 4](#).

Length: Variable, 7 or 8 depending on Label or Index encoding of the SID

Flags: 1 octet value which should be parsed as:

- * IS-IS Prefix-SID flags are defined in [\[I-D.ietf-isis-segment-routing-extensions\]](#) [section 2.1](#).
- * OSPFv2 Prefix-SID flags are defined in [\[I-D.ietf-ospf-segment-routing-extensions\]](#) [section 5](#).
- * OSPFv3 Prefix-SID flags are defined in [\[I-D.ietf-ospf-segment-routing-extensions\]](#) [section 5](#).

Algorithm: 1 octet value identify the algorithm.

Reserved: 2 octets that SHOULD be set to 0 and MUST be ignored on receipt.

SID/Index/Label:

- * IS-IS: Label or index value as defined in [\[I-D.ietf-isis-segment-routing-extensions\]](#),
- * OSPFv2: Label or index value as defined in [\[I-D.ietf-ospf-segment-routing-extensions\]](#),
- * OSPFv3: Label or index value as defined in [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#),

[2.3.2.](#) Prefix Attribute Flags TLV

The Prefix Attribute Flags TLV carries IPv4/IPv6 prefix attribute flags information. These flags are defined for OSPFv2 in [\[RFC7684\]](#), for OSPFv3 in [\[RFC5340\]](#) and for IS-IS in [\[RFC7794\]](#).

The Prefix Attribute Flags 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 (variable)                                     //
+-----+-----+-----+-----+-----+-----+-----+-----+

```

where:

Type: TBD, see [Section 4](#).

Length: variable.

Flags: a variable length flag field (according to the length field). Flags are routing protocol specific and are to be parsed as below:

- * IS-IS flags are defined in [\[RFC7794\]](#)
- * OSPFv2 flags are defined in [\[RFC7684\]](#)
- * OSPFv3 flags map to the Prefix Options field defined in [\[RFC7794\]](#) and extended via [\[RFC8362\]](#)

[2.3.3.](#) Source Router Identifier (Source Router-ID) TLV

The Source Router-ID TLV contains the IPv4 or IPv6 Router-ID of the originator of the Prefix. For IS-IS protocol this is as defined in [\[RFC7794\]](#). The Source Router-ID TLV may be used to carry the OSPF Router-ID of the prefix originator.

The Source Router-ID 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                    |
+-----+-----+-----+-----+-----+-----+-----+-----+
//                                4 or 6 octet Router-ID                                //
+-----+-----+-----+-----+-----+-----+-----+-----+

```

where:

Type: TBD, see [Section 4](#).

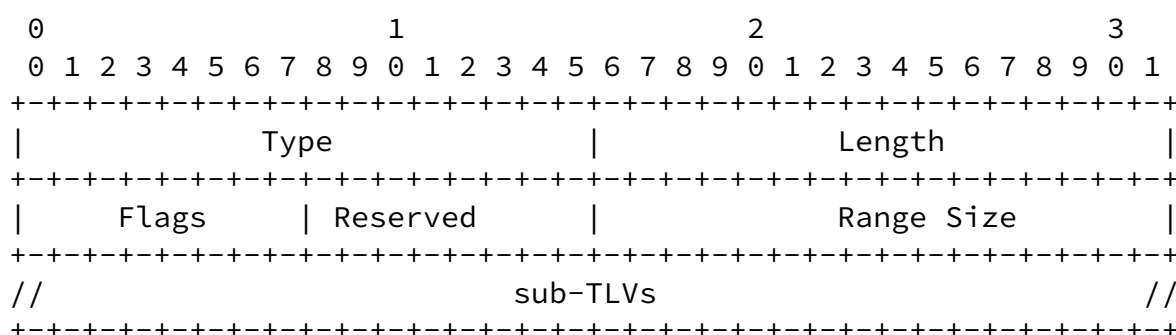
Length: 4 or 16.

IPv4/IPv6 Address: 4 octet IPv4 address or 16 octet IPv6 address.

[2.3.4.](#) Range TLV

The range TLV is used in order to advertise a range of prefix-to-SID mappings as part of the Segment Routing Mapping Server functionality [\[I-D.ietf-spring-segment-routing-ldp-interop\]](#), as defined in the respective underlying IGP SR extensions [\[I-D.ietf-ospf-segment-routing-extensions\]](#), [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#) and [\[I-D.ietf-isis-segment-routing-extensions\]](#). The Prefix-NLRI to which the Range TLV is attached MUST be advertised as a non-routing prefix where no IGP metric TLV (TLV 1095) is attached.

The format of the Range TLV is as follows:



where:

Figure 2: Range TLV format

Type: TBD, see [Section 4](#).

Length: variable

Flags: as defined in [[I-D.ietf-ospf-segment-routing-extensions](#)], [[I-D.ietf-ospf-ospfv3-segment-routing-extensions](#)] and [[I-D.ietf-isis-segment-routing-extensions](#)].

Reserved: 1 octet that SHOULD be set to 0 and MUST be ignored on receipt.

Range Size: 2 octets as defined in [[I-D.ietf-ospf-segment-routing-extensions](#)].

Within the Range TLV, the prefix-to-SID mappings are advertised using sub-TLVs as below:

Range TLV

Prefix-SID TLV (used as a sub-TLV in this context)

where:

- o The Range TLV is defined in [Section 2.3.4](#).
- o The Prefix-SID TLV (used as sub-TLV in this context) is defined in [Section 2.3.1](#).

[2.3.4.1](#). Advertisement Procedure for OSPF

The OSPFv2/OSPFv3 Extended Prefix Range TLV is encoded in the Range TLV. The flags of the Range TLV have the semantic mapped to the definition in [[I-D.ietf-ospf-segment-routing-extensions](#)] [section 4](#) or [[I-D.ietf-ospf-ospfv3-segment-routing-extensions](#)] [section 4](#).

Then the prefix-to-SID mapping from the OSPF Prefix SID sub-TLV is encoded using the BGP-LS Prefix-SID TLV as defined in [Section 2.3.1](#) with the flags set according to the definition in [[I-D.ietf-ospf-segment-routing-extensions](#)] [section 5](#) or [[I-D.ietf-ospf-ospfv3-segment-routing-extensions](#)] [section 5](#).

[2.3.4.2](#). Advertisement Procedure for IS-IS

The IS-IS SID/Label Binding TLV, when used to signal mapping server label bindings, is encoded in the Range TLV. The flags of the Range TLV have the semantic mapped to the definition in [[I-D.ietf-isis-segment-routing-extensions](#)] [section 2.4.1](#).

Then the prefix-to-SID mappings from the IS-IS Prefix SID sub-TLV is encoded using the BGP-LS Prefix-SID TLV as defined in [Section 2.3.1](#) with the flags set according to the definition in [[I-D.ietf-isis-segment-routing-extensions](#)] [section 2.4.4.1](#).

[2.4](#). Equivalent IS-IS Segment Routing TLVs/Sub-TLVs

This section illustrate the IS-IS Segment Routing Extensions TLVs and sub-TLVs mapped to the ones defined in this document.

The following table, illustrates for each BGP-LS TLV, its equivalence in IS-IS.

Description	Length	IS-IS TLV/sub-TLV
SR Capabilities	variable	2 [1]
SR Algorithm	variable	19 [2]
SR Local Block	variable	22 [3]
SRMS Preference	1	19 [4]
Adjacency Segment Identifier (Adj-SID)	variable	31 [5]
LAN Adjacency Segment Identifier (LAN-Adj-SID)	variable	32 [6]
Prefix SID	variable	3 [7]
Range	variable	149 [8]
SID/Label TLV	variable	1 [9]
Prefix Attribute Flags	variable	4 [10]
Source Router ID	variable	11/12 [11]
L2 Bundle Member TLV	variable	25 [12]

Table 5: IS-IS Segment Routing Extensions TLVs/Sub-TLVs

[2.5.](#) Equivalent OSPFv2/OSPFv3 Segment Routing TLVs/Sub-TLVs

This section illustrate the OSPFv2 and OSPFv3 Segment Routing Extensions TLVs and sub-TLVs mapped to the ones defined in this document.

The following table, illustrates for each BGP-LS TLV, its equivalence in OSPFv2 and OSPFv3.

Description	Length	OSPFv2 TLV/sub-TLV
SR Capabilities	variable	9 [13]
SR Algorithm	variable	8 [14]
SR Local Block	variable	14 [15]
SRMS Preference	1	15 [16]
Adjacency Segment Identifier (Adj-SID)	variable	2 [17]

LAN Adjacency Segment Identifier (Adj-SID)	variable	3	[18]
Prefix SID	variable	2	[19]
Range	variable	2	[20]
SID/Label TLV	variable	1	[21]
Prefix Attribute Flags	variable	4	[22]

Table 6: OSPF Segment Routing Extensions TLVs/Sub-TLVs

Description	Length	OSPFv3 TLV/sub-TLV
SR Capabilities	variable	9 [23]
SR Algorithm	variable	8 [24]
SR Local Block	variable	14 [25]
SRMS Preference	1	15 [26]
Adjacency Segment Identifier (Adj-SID)	variable	5 [27]
LAN Adjacency Segment Identifier (Adj-SID)	variable	6 [28]
Prefix SID	variable	4 [29]
Range	variable	9 [30]
SID/Label TLV	variable	7 [31]
Prefix Attribute Flags	variable	4 [32]

Table 7: OSPFv3 Segment Routing Extensions TLVs/Sub-TLVs

[3.](#) Implementation Status

Note to RFC Editor: Please remove this section prior to publication, as well as the reference to [RFC 7942](#).

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [\[RFC7942\]](#). The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort

has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [[RFC7942](#)], "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

Several early implementations exist and will be reported in detail in a forthcoming version of this document. For purposes of early interoperability testing, when no FCFS code point was available, implementations have made use of the values described in Table 8.

It will ease implementation interoperability and deployment if the value could be preserved also due to the large amount of codepoints this draft requires. However, when IANA-assigned values are available, implementations will be updated to use them.

[4.](#) IANA Considerations

This document requests assigning code-points from the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" based on table Table 8. The column "IS-IS TLV/Sub-TLV" defined in the registry does not require any value and should be left empty.

[4.1.](#) TLV/Sub-TLV Code Points Summary

This section contains the global table of all TLVs/sub-TLVs defined in this document.

TLV Code Point	Description	Reference
1034	SR Capabilities	Section 2.1.2

1035	SR Algorithm	Section 2.1.3
1036	SR Local Block	Section 2.1.4
1037	SRMS Preference	Section 2.1.5
1099	Adjacency Segment Identifier (Adj-SID) TLV	Section 2.2.1
1100	LAN Adjacency Segment Identifier (Adj-SID) TLV	Section 2.2.2
1158	Prefix SID	Section 2.3.1
1159	Range	Section 2.3.4
1161	SID/Label TLV	Section 2.1.1
1170	Prefix Attribute Flags	Section 2.3.2
1171	Source Router-ID	Section 2.3.3
1172	L2 Bundle Member TLV	Section 2.2.3

Table 8: Summary Table of TLV/Sub-TLV Codepoints

5. Manageability Considerations

This section is structured as recommended in [\[RFC5706\]](#).

The new protocol extensions introduced in this document augment the existing IGP topology information that was distributed via [\[RFC7752\]](#). Procedures and protocol extensions defined in this document do not affect the BGP protocol operations and management other than as discussed in the Manageability Considerations section of [\[RFC7752\]](#). Specifically, the malformed attribute tests for syntactic checks in the Fault Management section of [\[RFC7752\]](#) now encompass the new BGP-LS Attribute TLVs defined in this document. The semantic or content checking for the TLVs specified in this document and their association with the BGP-LS NLRI types or their BGP-LS Attribute is left to the consumer of the BGP-LS information (e.g. an application or a controller) and not the BGP protocol.

A consumer of the BGP-LS information is retrieving this information from a BGP protocol component that is doing the signaling over a BGP-LS session, via some APIs or a data model (refer [Section 1](#) and 2 of [\[RFC7752\]](#)). The handling of semantic or content errors by the consumer would be dictated by the nature of its application usage and hence is beyond the scope of this document. This document only introduces new Attribute TLVs and an error in them would result in only that specific attribute being discarded with an error log.

The extensions, specified in this document, do not introduce any new configuration or monitoring aspects in BGP or BGP-LS other than as discussed in [[RFC7752](#)]. The manageability aspects of the underlying SR features are covered by [[I-D.ietf-spring-sr-yang](#)], [[I-D.ietf-isis-sr-yang](#)] and [[I-D.ietf-ospf-sr-yang](#)].

[6.](#) Security Considerations

The new protocol extensions introduced in this document augment the existing IGP topology information that was distributed via [[RFC7752](#)]. The Security Considerations section of [[RFC7752](#)] also applies to these extensions. The procedures and new TLVs defined in this document, by themselves, do not affect the BGP-LS security model discussed in [[RFC7752](#)].

BGP-LS SR extensions enable traffic engineering use-cases within the Segment Routing domain. SR operates within a trusted domain (refer Security Considerations section in [[RFC8402](#)] for more detail) and its security considerations also apply to BGP-LS sessions when carrying SR information. The SR traffic engineering policies using the SIDs advertised via BGP-LS are expected to be used entirely within this trusted SR domain (e.g. between multiple AS/domains within a single provider network). Therefore, precaution is necessary to ensure that the SR information collected via BGP-LS is limited to specific controllers or applications in a secure manner within this SR domain.

The isolation of BGP-LS peering sessions is also required to ensure that BGP-LS topology information (including the newly added SR information) is not advertised to an external BGP peering session outside an administrative domain.

[7.](#) Contributors

The following people have substantially contributed to the editing of this document:

Peter Psenak
Cisco Systems
Email: ppsenak@cisco.com

Les Ginsberg
Cisco Systems
Email: ginsberg@cisco.com

Acee Lindem
Cisco Systems
Email: acee@cisco.com

Internet-Draft BGP LS extensions for Segment Routing October 2018

Saikat Ray
Individual
Email: raysaikat@gmail.com

Jeff Tantsura
Apstra Inc.
Email: jefftant.ietf@gmail.com

[8.](#) Acknowledgements

The authors would like to thank Jeffrey Haas, Aijun Wang and Robert Raszuk for their review of this document and their comments.

[9.](#) References

[9.1.](#) Normative References

[I-D.ietf-idr-te-pm-bgp]

Ginsberg, L., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP-LS Advertisement of IGP Traffic Engineering Performance Metric Extensions", [draft-ietf-idr-te-pm-bgp-14](#) (work in progress), October 2018.

[I-D.ietf-isis-segment-routing-extensions]

Previdi, S., Ginsberg, L., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., Decraene, B., and J. Tantsura, "IS-IS Extensions for Segment Routing", [draft-ietf-isis-segment-routing-extensions-19](#) (work in progress), July 2018.

[I-D.ietf-ospf-ospfv3-segment-routing-extensions]

Psenak, P. and S. Previdi, "OSPFv3 Extensions for Segment Routing", [draft-ietf-ospf-ospfv3-segment-routing-extensions-16](#) (work in progress), October 2018.

[I-D.ietf-ospf-segment-routing-extensions]

Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", [draft-ietf-ospf-segment-routing-extensions-25](#) (work in progress), April 2018.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate

Requirement Levels", [BCP 14](#), [RFC 2119](#),
DOI 10.17487/RFC2119, March 1997,
<<https://www.rfc-editor.org/info/rfc2119>>.

- [RFC4202] Kompella, K., Ed. and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4202](#), DOI 10.17487/RFC4202, October 2005, <<https://www.rfc-editor.org/info/rfc4202>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", [RFC 5340](#), DOI 10.17487/RFC5340, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.
- [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, <<https://www.rfc-editor.org/info/rfc7684>>.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", [RFC 7752](#), DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.
- [RFC7794] Ginsberg, L., Ed., Decraene, B., Previdi, S., Xu, X., and U. Chunduri, "IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability", [RFC 7794](#), DOI 10.17487/RFC7794, March 2016, <<https://www.rfc-editor.org/info/rfc7794>>.
- [RFC8362] Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", [RFC 8362](#), DOI 10.17487/RFC8362, April 2018, <<https://www.rfc-editor.org/info/rfc8362>>.

[9.2](#). Informative References

- [I-D.ietf-isis-sr-yang]
Litkowski, S., Qu, Y., Sarkar, P., Chen, I., and J.
Tantsura, "YANG Data Model for IS-IS Segment Routing",

[draft-ietf-isis-sr-yang-04](#) (work in progress), June 2018.

[I-D.ietf-ospf-sr-yang]

Yeung, D., Qu, Y., Zhang, Z., Chen, I., and A. Lindem, "Yang Data Model for OSPF SR (Segment Routing) Protocol", [draft-ietf-ospf-sr-yang-05](#) (work in progress), July 2018.

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

Bashandy, A., Filsfils, C., Previdi, S., Decraene, B., and S. Litkowski, "Segment Routing interworking with LDP", [draft-ietf-spring-segment-routing-ldp-interop-15](#) (work in progress), September 2018.

Previdi, et al.

Expires April 25, 2019

[Page 25]

Internet-Draft BGP LS extensions for Segment Routing October 2018

[I-D.ietf-spring-sr-yang]

Litkowski, S., Qu, Y., Sarkar, P., and J. Tantsura, "YANG Data Model for Segment Routing", [draft-ietf-spring-sr-yang-09](#) (work in progress), June 2018.

[RFC5706] Harrington, D., "Guidelines for Considering Operations and Management of New Protocols and Protocol Extensions", [RFC 5706](#), DOI 10.17487/RFC5706, November 2009, <<https://www.rfc-editor.org/info/rfc5706>>.

[RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", [BCP 205](#), [RFC 7942](#), DOI 10.17487/RFC7942, July 2016, <<https://www.rfc-editor.org/info/rfc7942>>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

[9.3](#). URIs

[1] <http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-16#section-3.1>

[2] <http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-16#section-3.2>

- [3] <http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-16#section-3.3>
- [4] <http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-3.2>
- [5] <http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-16#section-2.2.1>
- [6] <http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-16#section-2.2.2>
- [7] <http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-16#section-2.1>
- [8] <http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-16#section-2.4>
- [9] <http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-16#section-2.3>

- [10] <http://tools.ietf.org/html/RFC7794>
- [11] <http://tools.ietf.org/html/RFC7794>
- [12] <http://tools.ietf.org/html/draft-ietf-isis-l2bundles-07>
- [13] <http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-25#section-3.2>
- [14] <http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-25#section-3.1>
- [15] <http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-25#section-3.3>
- [16] <http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-25#section-3.4>
- [17] <http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-25#section-6.1>

- [18] <http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-25#section-6.2>
- [19] <http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-25#section-5>
- [20] <http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-25#section-4>
- [21] <http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-25#section-2.1>
- [22] <http://tools.ietf.org/html/RFC7684#section-2.1>
- [23] <http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-12#section-3.2>
- [24] <http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-12#section-3.1>
- [25] <http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-12#section-3.3>
- [26] <http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-12#section-3.4>

- [27] <http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-12#section-6.1>
- [28] <http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-12#section-6.2>
- [29] <http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-12#section-5>
- [30] <http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-12#section-4>
- [31] <http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-12#section-2.1>

[32] <http://tools.ietf.org/html/RFC8362#section-3.1>

Authors' Addresses

Stefano Previdi
Huawei Technologies
Rome
Italy

Email: stefano@previdi.net

Ketan Talaulikar
Cisco Systems, Inc.
India

Email: ketant@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Brussels
Belgium

Email: cfilsfil@cisco.com

Hannes Gredler
RtBrick Inc.

Email: hannes@rtbrick.com

Mach(Guoyi) Chen
Huawei Technologies
Huawei Building, No. 156 Beiqing Rd.
Beijing 100095
China

Email: mach.chen@huawei.com

