

Internet Engineering Task Force (IETF)
Request for Comments: 9085
Category: Standards Track
ISSN: 2070-1721

S. Previdi
Huawei Technologies
K. Talaulikar, Ed.
C. Filsfils
Cisco Systems, Inc.
H. Gredler
RtBrick Inc.
M. Chen
Huawei Technologies
August 2021

Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing

Abstract

Segment Routing (SR) allows for a flexible definition of end-to-end paths by encoding paths as sequences of topological subpaths, 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 document defines extensions to the Border Gateway Protocol - Link State (BGP-LS) address family in order to carry SR information via BGP.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9085>.

Copyright Notice

Copyright (c) 2021 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
 - 1.1. Requirements Language
2. BGP-LS Extensions for Segment Routing
 - 2.1. Node Attribute TLVs
 - 2.1.1. SID/Label TLV
 - 2.1.2. SR Capabilities TLV
 - 2.1.3. SR-Algorithm TLV
 - 2.1.4. SR Local Block TLV
 - 2.1.5. SRMS Preference TLV

- 2.2. Link Attribute TLVs
 - 2.2.1. Adjacency SID TLV
 - 2.2.2. LAN Adjacency SID TLV
 - 2.2.3. L2 Bundle Member Attributes TLV
 - 2.3. Prefix Attribute TLVs
 - 2.3.1. Prefix-SID TLV
 - 2.3.2. Prefix Attribute Flags TLV
 - 2.3.3. Source Router Identifier TLV
 - 2.3.4. Source OSPF Router-ID TLV
 - 2.3.5. Range TLV
 - 2.4. Equivalent IS-IS Segment Routing TLVs/Sub-TLVs
 - 2.5. Equivalent OSPFv2/OSPFv3 Segment Routing TLVs/Sub-TLVs
 - 3. IANA Considerations
 - 3.1. TLV/Sub-TLV Code Points Summary
 - 4. Manageability Considerations
 - 5. Security Considerations
 - 6. References
 - 6.1. Normative References
 - 6.2. Informative References
- Acknowledgements
 Contributors
 Authors' Addresses

1. Introduction

Segment Routing (SR) allows for a flexible definition of end-to-end paths by combining subpaths 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 semantic within a domain. Within IGP topologies, an SR path is encoded as a sequence of topological subpaths, called "IGP segments". These segments are advertised by the link-state routing protocols (IS-IS, OSPFv2, and OSPFv3).

[RFC8402] defines the link-state IGP segments -- prefix, node, anycast, 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. Node and anycast segments are variations of the prefix segment with their specific characteristics.

When SR is enabled in an 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 for: IS-IS [RFC8667], OSPFv2 [RFC8665], and OSPFv3 [RFC8666]. Using these extensions, SR can be enabled within an IGP domain.

SR allows advertisement of single or multi-hop paths. The flooding scope for the IGP extensions for SR 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; therefore, by using the IGP alone, it is not enough to construct segments across multiple IGP area or Autonomous System (AS) boundaries.

In order to address the need for applications that require topological visibility across IGP areas, or even across ASes, the BGP-LS address family / subaddress 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.

```

+-----+
| Consumer |
+-----+
^

```

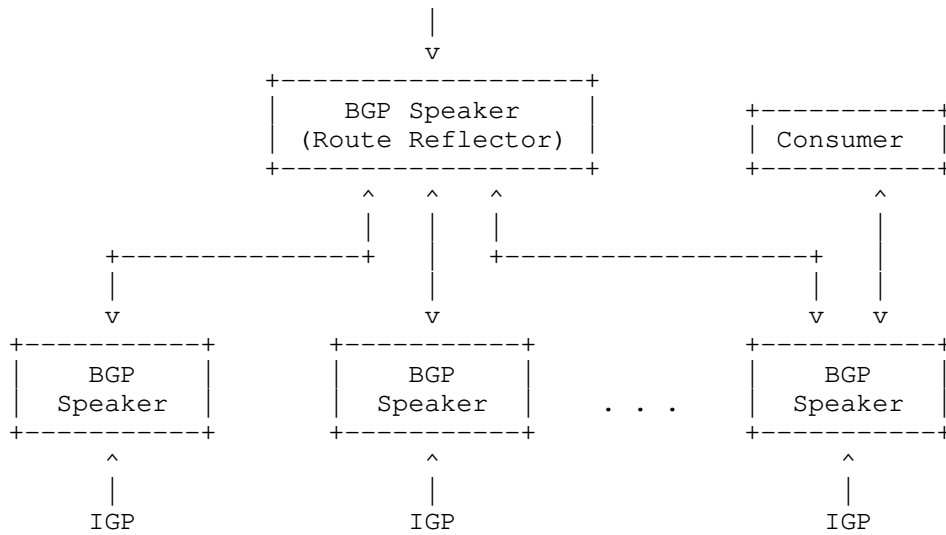


Figure 1: Link-State Information Collection

Figure 1 denotes a typical deployment scenario. In each IGP area, one or more nodes are configured with BGP-LS. These BGP speakers form an Internal BGP (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 External BGP (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) as described in [RFC7752].

This document describes extensions to BGP-LS to advertise the SR information. An external component (e.g., a controller) can collect SR information from across an SR domain (as described in [RFC8402]) and construct the end-to-end path (with its associated SIDs) that needs to be applied to an incoming packet to achieve the desired end-to-end forwarding. SR operates within a trusted domain consisting of a single AS or multiple ASes managed by the same administrative entity, e.g., within a single provider network.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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. Sections 2.4 and 2.5 list the equivalent TLVs and sub-TLVs in the 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, and it 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.

2.1. Node Attribute TLVs

The following Node Attribute TLVs are defined:

Type	Description	Section
1161	SID/Label	Section 2.1.1

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

Table 1: Node Attribute TLVs

These TLVs should only be added to the BGP-LS Attribute associated with the Node NLRI that describes the IGP node that is originating the corresponding IGP TLV/sub-TLV described below.

2.1.1.1. SID/Label TLV

The SID/Label TLV is used as a sub-TLV by the SR Capabilities (Section 2.1.2) and Segment Routing Local Block (SRLB) (Section 2.1.4) TLVs. This information is derived from the protocol-specific advertisements.

- * IS-IS, as defined by the SID/Label Sub-TLV in Section 2.3 of [RFC8667].
- * OSPFv2/OSPFv3, as defined by the SID/Label Sub-TLV in Section 2.1 of [RFC8665] and Section 3.1 of [RFC8666].

The TLV has the following format:




Figure 2: SID/Label TLV Format

Where:

Type: 1161

Length: Variable. Either 3 or 4 octets depending on whether the value is encoded as a label or as an index/SID.


SID/Label: If the length is set to 3, then the 20 rightmost bits represent a label (the total TLV size is 7), and the 4 leftmost bits are set to 0. If the length is set to 4, then the value represents a 32-bit SID (the total TLV size is 8).

2.1.1.2. 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 the SR-MPLS forwarding plane. This information is derived from the protocol-specific advertisements.

- * IS-IS, as defined by the SR-Capabilities Sub-TLV in Section 3.1 of [RFC8667].
- * OSPFv2/OSPFv3, as defined by the SID/Label Range TLV in Section 3.2 of [RFC8665]. OSPFv3 leverages the same TLV as defined for OSPFv2.

The SR Capabilities TLV has the following format:



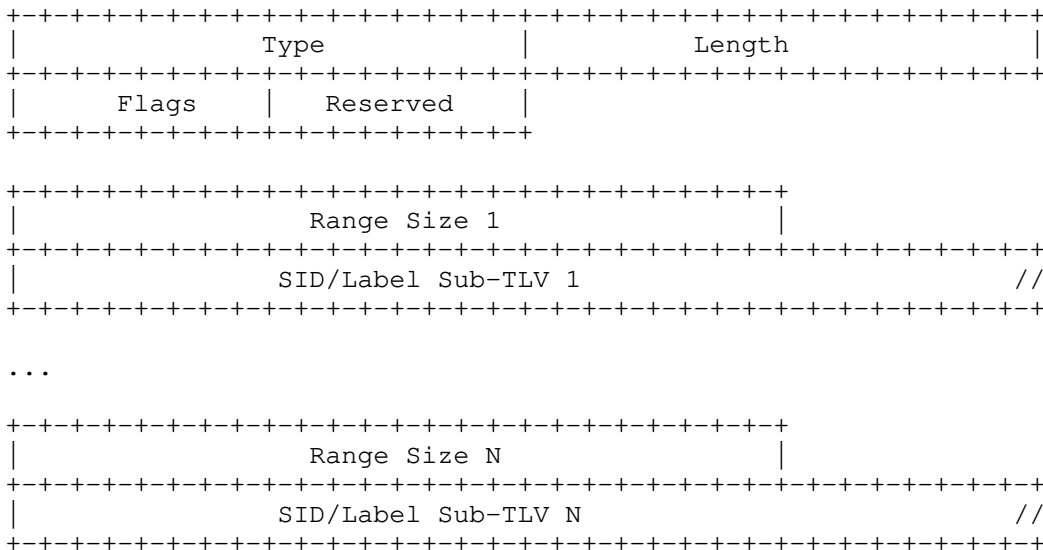


Figure 3: SR Capabilities TLV Format

Where:

Type: 1034

Length: Variable. The minimum length is 12 octets.

Flags: 1 octet of flags as defined in Section 3.1 of [RFC8667] for IS-IS. The flags are not currently defined for OSPFv2 and OSPFv3 and MUST be set to 0 and ignored on receipt.

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

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

Range Size: 3 octets with a non-zero value indicating the number of labels in the range.

SID/Label TLV: (as defined in Section 2.1.1) used as a sub-TLV, which encodes the first label in the range. Since the SID/Label TLV is used to indicate the first label of the SRGB range, only label encoding is valid under the SR Capabilities TLV.

2.1.3. SR-Algorithm TLV

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.

- * IS-IS, as defined by the SR-Algorithm Sub-TLV in Section 3.2 of [RFC8667].
- * OSPFv2/OSPFv3, as defined by the SR-Algorithm TLV in Section 3.1 of [RFC8665]. OSPFv3 leverages the same TLV as defined for OSPFv2.

The SR-Algorithm TLV has the following format:




Figure 4: SR-Algorithm TLV Format

Where:

Type: 1035

Length: Variable. The minimum length is 1 octet and the maximum can be 256.

Algorithm: One or more fields of 1 octet each that identifies the algorithm.

2.1.4. SR Local Block TLV

The 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, the node is required to advertise its SRLB.

This information is derived from the protocol-specific advertisements.

- * IS-IS, as defined by the SRLB Sub-TLV in Section 3.3 of [RFC8667].
- * OSPFv2/OSPFv3, as defined by the SR Local Block TLV in Section 3.3 of [RFC8665]. OSPFv3 leverages the same TLV as defined for OSPFv2.

The SRLB TLV has the following format:

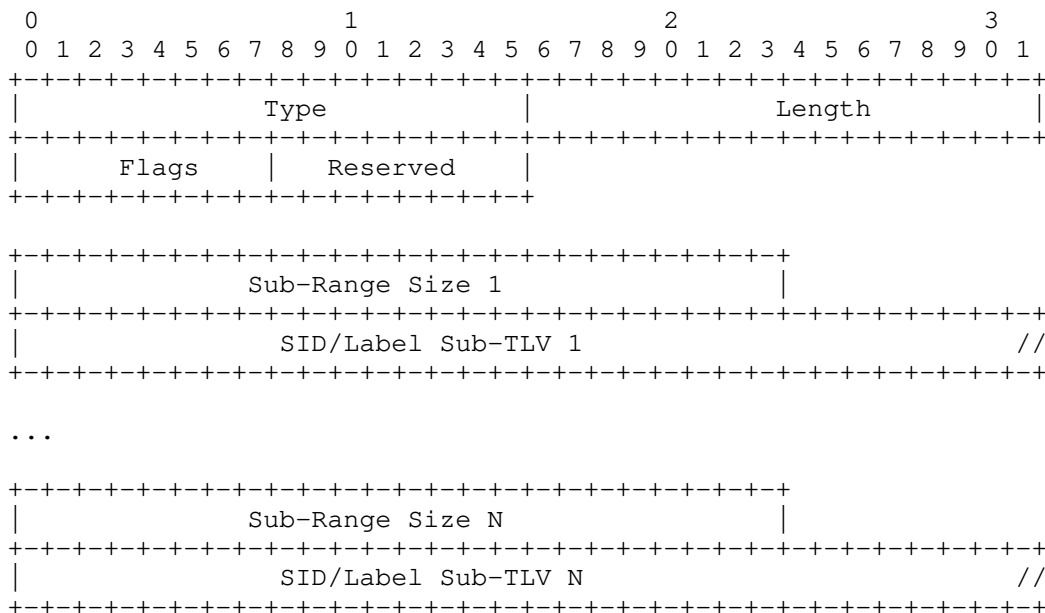


Figure 5: SRLB TLV Format

Where:

Type: 1036

Length: Variable. The minimum length is 12 octets.

Flags: 1 octet of flags. The flags are as defined in Section 3.3 of [RFC8667] for IS-IS. The flags are not currently defined for OSPFv2 and OSPFv3 and MUST be set to 0 and ignored on receipt.

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

One or more entries corresponding to a sub-range(s), each of which have the following format:

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

SID/Label TLV: (as defined in Section 2.1.1) used as a sub-TLV, which encodes the first label in the sub-range. Since the SID/Label TLV is used to indicate the first label of the SRLB sub-range, only label encoding is valid under the SR Local Block TLV.

2.1.5. SRMS Preference TLV

The Segment Routing Mapping Server (SRMS) Preference TLV is used in order to associate a preference with SRMS advertisements from a particular source. [RFC8661] specifies the SRMS functionality along with the SRMS preference of the node advertising the SRMS Prefix-to-SID mapping ranges.

This information is derived from the protocol-specific advertisements.

- * IS-IS, as defined by the SRMS Preference Sub-TLV in Section 3.4 of [RFC8667].
- * OSPFv2/OSPFv3, as defined by the SRMS Preference TLV in Section 3.4 of [RFC8665]. OSPFv3 leverages the same TLV as defined for OSPFv2.

The SRMS Preference TLV has the following format:




Figure 6: SRMS Preference TLV Format

Where:

Type: 1037

Length: 1 octet

Preference: 1 octet carrying an unsigned 8-bit SRMS preference.

2.2. Link Attribute TLVs

The following Link Attribute TLVs are defined:

Type	Description	Section
1099	Adjacency SID TLV	Section 2.2.1
1100	LAN Adjacency SID TLV	Section 2.2.2
1172	L2 Bundle Member Attributes TLV	Section 2.2.3

Table 2: Link Attribute TLVs

These TLVs should only be added to the BGP-LS Attribute associated with the Link NLRI that describes the link of the IGP node that is originating the corresponding IGP TLV/sub-TLV described below.

2.2.1. Adjacency SID TLV

The Adjacency SID TLV is used in order to advertise information related to an Adjacency SID. This information is derived from the Adj-SID Sub-TLV of IS-IS (Section 2.2.1 of [RFC8667]), OSPFv2 (Section 6.1 of [RFC8665]), and OSPFv3 (Section 7.1 of [RFC8666]).

The Adjacency SID TLV has the following format:




Figure 7: Adjacency SID TLV Format

Where:

Type: 1099

Length: Variable. Either 7 or 8 octets depending on the label or index encoding of the SID.

Flags: 1-octet value that should be set as:

- * IS-IS Adj-SID flags as defined in Section 2.2.1 of [RFC8667].
- * OSPFv2 Adj-SID flags as defined in Section 6.1 of [RFC8665].
- * OSPFv3 Adj-SID flags as defined in Section 7.1 of [RFC8666].

Weight: 1 octet carrying the weight used for load-balancing purposes. The use of weight is described in Section 3.4 of [RFC8402].

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

SID/Index/Label:

IS-IS: Label or index value as defined in Section 2.2.1 of [RFC8667].

OSPFv2: Label or index value as defined in Section 6.1 of [RFC8665].

OSPFv3: Label or index value as defined in Section 7.1 of [RFC8666].

The Flags and, as an extension, the SID/Index/Label fields of this TLV are interpreted according to the respective underlying IS-IS, OSPFv2, or OSPFv3 protocol. The Protocol-ID of the BGP-LS Link NLRI is used to determine the underlying protocol specification for parsing these fields.

2.2.2. LAN Adjacency SID TLV

For a LAN, normally a node only announces its adjacency to the IS-IS pseudonode (or the equivalent OSPF Designated and Backup Designated Routers). The LAN Adjacency SID 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.

This information is derived from the LAN-Adj-SID Sub-TLV of IS-IS (Section 2.2.2 of [RFC8667]), the LAN Adj-SID Sub-TLV of OSPFv2 (Section 6.2 of [RFC8665]), and the LAN Adj-SID Sub-TLV of OSPFv3 (Section 7.2 of [RFC8666]).

The LAN Adjacency SID TLV has the following format:

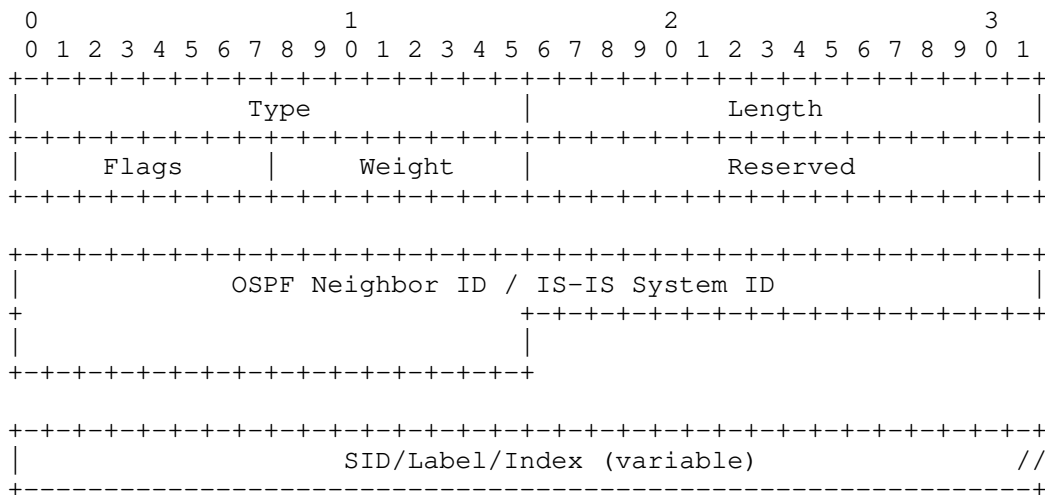


Figure 8: LAN Adjacency SID TLV Format

Where:

Type: 1100

Length: Variable. For IS-IS, it would be 13 or 14 octets depending on the label or index encoding of the SID. For OSPF, it would be 11 or 12 octets depending on the label or index encoding of the SID.

Flags: 1-octet value that should be set as:

- * IS-IS LAN Adj-SID flags as defined in Section 2.2.2 of [RFC8667].
- * OSPFv2 LAN Adj-SID flags as defined in Section 6.2 of [RFC8665].
- * OSPFv3 LAN Adj-SID flags as defined in Section 7.2 of [RFC8666].

Weight: 1 octet carrying the weight used for load-balancing purposes. The use of weight is described in Section 3.4 of [RFC8402].

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

Neighbor ID: 6 octets for IS-IS for the System ID, and 4 octets for OSPF for the OSPF Router-ID of the neighbor.

SID/Index/Label:

IS-IS: Label or index value as defined in Section 2.2.2 of [RFC8667].

OSPFv2: Label or index value as defined in Section 6.2 of [RFC8665].

OSPFv3: Label or index value as defined in Section 7.2 of [RFC8666].

The Neighbor ID, Flags, and, as an extension, the SID/Index/Label fields of this TLV are interpreted according to the respective underlying IS-IS, OSPFv2, or OSPFv3 protocol. The Protocol-ID of the BGP-LS Link NLRI is used to determine the underlying protocol specification for parsing these fields.

2.2.3. L2 Bundle Member Attributes TLV

The L2 Bundle Member Attributes 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 Attributes TLV is associated with the Link NLRI. The TLV MAY include sub-TLVs that 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 Attributes TLVs MAY be associated with a Link NLRI.

This information is derived from L2 Bundle Member Attributes TLV of IS-IS (Section 2 of [RFC8668]). The equivalent functionality has not been specified as yet for OSPF.

The L2 Bundle Member Attributes TLV has the following format:

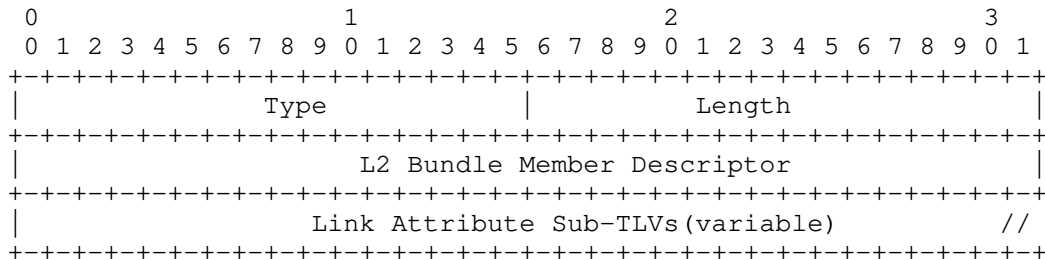


Figure 9: L2 Bundle Member Attributes TLV Format

Where:

Type: 1172

Length: Variable.

L2 Bundle Member Descriptor: 4-octet field that carries a link-local identifier as defined in [RFC4202].

Link attributes for L2 Bundle Member links are advertised as sub-TLVs of the L2 Bundle Member Attributes 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	[RFC8571]
1115	Min/Max Unidirectional link delay	[RFC8571]
1116	Unidirectional Delay Variation	[RFC8571]
1117	Unidirectional Link Loss	[RFC8571]

1118	Unidirectional residual bandwidth	[RFC8571]
1119	Unidirectional available bandwidth	[RFC8571]
1120	Unidirectional Utilized Bandwidth	[RFC8571]

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

2.3. Prefix Attribute TLVs

The following Prefix Attribute TLVs are defined:

Type	Description	Section
1158	Prefix-SID	Section 2.3.1
1159	Range	Section 2.3.5
1170	Prefix Attribute Flags	Section 2.3.2
1171	Source Router Identifier	Section 2.3.3
1174	Source OSPF Router-ID	Section 2.3.4

Table 4: Prefix Attribute TLVs

These TLVs should only be added to the BGP-LS Attribute associated with the Prefix NLRI that describes the prefix of the IGP node that is originating the corresponding 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 derived from the Prefix-SID Sub-TLV of IS-IS (Section 2.1 of [RFC8667]), the Prefix-SID Sub-TLV of OSPFv2 (Section 5 of [RFC8665]), and the Prefix-SID Sub-TLV of OSPFv3 (Section 6 of [RFC8666]).

The Prefix-SID TLV has the following format:




Figure 10: Prefix-SID TLV Format

Where:

Type: 1158

Length: Variable. 7 or 8 octets depending on the label or index encoding of the SID.

Flags: 1-octet value that should be set as:

- * IS-IS Prefix-SID flags as defined in Section 2.1.1 of [RFC8667].

* OSPFv2 Prefix-SID flags as defined in Section 5 of [RFC8665].

* OSPFv3 Prefix-SID flags as defined in Section 6 of [RFC8665].

Algorithm: 1-octet value identifies the algorithm. The semantics of the algorithm are described in Section 3.1.1 of [RFC8402].

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

SID/Index/Label:

IS-IS: Label or index value as defined in Section 2.1 of [RFC8667].

OSPFv2: Label or index value as defined in Section 5 of [RFC8665].

OSPFv3: Label or index value as defined in Section 6 of [RFC8666].

The Flags and, as an extension, the SID/Index/Label fields of this TLV are interpreted according to the respective underlying IS-IS, OSPFv2, or OSPFv3 protocol. The Protocol-ID of the BGP-LS Prefix NLRI is used to determine the underlying protocol specification for parsing these fields.

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 Section 2.1 of [RFC7684], OSPFv3 in Appendix A.4.1.1 of [RFC5340], and IS-IS in Section 2.1 of [RFC7794].

The Prefix Attribute Flags TLV has the following format:




Figure 11: Prefix Attribute Flags TLV Format

Where:

Type: 1170

Length: Variable.

Flags: a variable-length Flag field (according to the Length field). Flags are routing protocol specific and are to be set as below:

- * IS-IS flags correspond to the IPv4/IPv6 Extended Reachability Attribute Flags defined in Section 2.1 of [RFC7794]. In the case of the X-flag when associated with IPv6 prefix reachability, the setting corresponds to the setting of the X-flag in the fixed format of IS-IS TLVs 236 [RFC5308] and 237 [RFC5120].
- * OSPFv2 flags correspond to the Flags field of the OSPFv2 Extended Prefix TLV defined in Section 2.1 of [RFC7684].
- * OSPFv3 flags map to the Prefix Options field defined in Appendix A.4.1.1 of [RFC5340] and extended in Section 3.1 of [RFC8362].

The Flags field of this TLV is interpreted according to the respective underlying IS-IS, OSPFv2, or OSPFv3 protocol. The

Protocol-ID of the BGP-LS Prefix NLRI is used to determine the underlying protocol specification for parsing this field.

2.3.3. Source Router Identifier TLV

The Source Router Identifier TLV contains the IPv4 or IPv6 Router Identifier of the originator of the prefix. For the IS-IS protocol, this is derived from the IPv4/IPv6 Source Router ID Sub-TLV as defined in Section 2.2 of [RFC7794]. For the OSPF protocol, this is derived from the Prefix Source Router Address Sub-TLV as defined in Section 2.2 of [RFC9084].

The Source Router Identifier TLV has the following format:




Figure 12: Source Router Identifier TLV Format

Where:

Type: 1171

Length: Variable. 4 or 16 octets for the IPv4 or IPv6 prefix, respectively.

Router-ID: the IPv4 or IPv6 Router-ID in the case of IS-IS, and the IPv4 or IPv6 Router Address in the case of OSPF.

2.3.4. Source OSPF Router-ID TLV

The Source OSPF Router-ID TLV is applicable only for the OSPF protocol and contains the OSPF Router-ID of the originator of the prefix. It is derived from the Prefix Source OSPF Router-ID Sub-TLV as defined in Section 2.1 of [RFC9084].

The Source OSPF Router-ID TLV has the following format:




Figure 13: Source OSPF Router-ID TLV Format

Where:

Type: 1174

Length: 4 octets

OSPF Router-ID: the OSPF Router-ID of the node originating the prefix.

2.3.5. Range TLV

The Range TLV is used in order to advertise a range of prefix-to-SID mappings as part of the SRMS functionality [RFC8661], as defined in the respective underlying IGP SR extensions: Section 4 of [RFC8665], Section 5 of [RFC8666], and Section 2.4 of [RFC8667]. The information advertised in the Range TLV is derived from the SID/Label Binding TLV in the case of IS-IS and the OSPFv2/OSPFv3 Extended Prefix Range TLV in the case of OSPFv2/OSPFv3.

A Prefix NLRI, that has been advertised with a Range TLV, is considered a normal routing prefix (i.e., prefix reachability) only when there is also an IGP metric TLV (TLV 1095) associated it. Otherwise, it is considered only as the first prefix in the range for prefix-to-SID mapping advertisement.

The format of the Range TLV is as follows:




Figure 14: Range TLV Format

Where:

Type: 1159

Length: Variable. 11 or 12 octets depending on the label or index encoding of the SID.

Flags: 1-octet value that should be set as:

- * IS-IS SID/Label Binding TLV flags as defined in Section 2.4.1 of [RFC8667].
- * OSPFv2 OSPF Extended Prefix Range TLV flags as defined in Section 4 of [RFC8665].
- * OSPFv3 Extended Prefix Range TLV flags as defined in Section 5 of [RFC8666].

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

Range Size: 2 octets that carry the number of prefixes that are covered by the advertisement.

The Flags field of this TLV is interpreted according to the respective underlying IS-IS, OSPFv2, or OSPFv3 protocol. The Protocol-ID of the BGP-LS Prefix NLRI is used to determine the underlying protocol specification for parsing this field.

The prefix-to-SID mappings are advertised using sub-TLVs as below:

IS-IS:

- SID/Label Range TLV
- Prefix-SID Sub-TLV

OSPFv2/OSPFv3:

- OSPFv2/OSPFv3 Extended Prefix Range TLV
- Prefix-SID Sub-TLV

BGP-LS:

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

The prefix-to-SID mapping information for the BGP-LS Prefix-SID TLV (used as a sub-TLV in this context) is encoded as described in Section 2.3.1.

2.4. Equivalent IS-IS Segment Routing TLVs/Sub-TLVs

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

For each BGP-LS TLV, the following table illustrates its equivalence in IS-IS.

Description	IS-IS TLV/sub-TLV	Reference
SR Capabilities	SR-Capabilities Sub-TLV (2)	[RFC8667]
SR Algorithm	SR-Algorithm Sub-TLV (19)	[RFC8667]
SR Local Block	SR Local Block Sub-TLV (22)	[RFC8667]
SRMS Preference	SRMS Preference Sub-TLV (19)	[RFC8667]
Adjacency SID	Adj-SID Sub-TLV (31)	[RFC8667]
LAN Adjacency SID	LAN-Adj-SID Sub-TLV (32)	[RFC8667]
Prefix-SID	Prefix-SID Sub-TLV (3)	[RFC8667]
Range	SID/Label Binding TLV (149)	[RFC8667]
SID/Label	SID/Label Sub-TLV (1)	[RFC8667]
Prefix Attribute Flags	Prefix Attribute Flags Sub-TLV (4)	[RFC7794]
Source Router Identifier	IPv4/IPv6 Source Router ID Sub-TLV (11/12)	[RFC7794]
L2 Bundle Member Attributes	L2 Bundle Member Attributes TLV (25)	[RFC8668]

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

2.5. Equivalent OSPFv2/OSPFv3 Segment Routing TLVs/Sub-TLVs

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

For each BGP-LS TLV, the following tables illustrate its equivalence in OSPFv2 and OSPFv3.

Description	OSPFv2 TLV/sub-TLV	Reference
SR Capabilities	SID/Label Range TLV (9)	[RFC8665]
SR Algorithm	SR-Algorithm TLV (8)	[RFC8665]
SR Local Block	SR Local Block TLV (14)	[RFC8665]
SRMS Preference	SRMS Preference TLV (15)	[RFC8665]
Adjacency SID	Adj-SID Sub-TLV (2)	[RFC8665]
LAN Adjacency SID	LAN Adj-SID Sub-TLV (3)	[RFC8665]
Prefix-SID	Prefix-SID Sub-TLV (2)	[RFC8665]
Range	OSPF Extended Prefix Range TLV (2)	[RFC8665]
SID/Label	SID/Label Sub-TLV (1)	[RFC8665]
Prefix Attribute Flags	Flags of OSPFv2 Extended Prefix TLV (1)	[RFC7684]

Source Router Identifier	Prefix Source Router Address Sub-TLV (5)	[RFC9084]
Source OSPF Router-ID	Prefix Source OSPF Router-ID Sub-TLV (4)	[RFC9084]

Table 6: OSPFv2 Segment Routing Extensions TLVs/Sub-TLVs

Description	OSPFv3 TLV/sub-TLV	Reference
SR Capabilities	SID/Label Range TLV (9)	[RFC8665]
SR Algorithm	SR-Algorithm TLV (8)	[RFC8665]
SR Local Block	SR Local Block TLV (14)	[RFC8665]
SRMS Preference	SRMS Preference TLV (15)	[RFC8665]
Adjacency SID	Adj-SID Sub-TLV (5)	[RFC8666]
LAN Adjacency SID	LAN Adj-SID Sub-TLV (6)	[RFC8666]
Prefix-SID	Prefix-SID Sub-TLV (4)	[RFC8666]
Range	OSPFv3 Extended Prefix Range TLV (9)	[RFC8666]
SID/Label	SID/Label Sub-TLV (7)	[RFC8666]
Prefix Attribute Flags	Prefix Option Fields of Prefix TLV types 3,5,6	[RFC8362]
Source OSPF Router Identifier	Prefix Source Router Address Sub-TLV (28)	[RFC9084]
Source OSPF Router-ID	Prefix Source OSPF Router-ID Sub-TLV (27)	[RFC9084]

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

3. IANA Considerations

IANA has registered the following code points in the "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" registry under the "Border Gateway Protocol - Link State (BGP-LS) Parameter" registry based on Table 8. The column "IS-IS TLV/Sub-TLV" defined in the registry does not require any value and should be left empty.

3.1. 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 SID	Section 2.2.1
1100	LAN Adjacency SID	Section 2.2.2

1158	Prefix-SID	Section 2.3.1
1159	Range	Section 2.3.5
1161	SID/Label	Section 2.1.1
1170	Prefix Attribute Flags	Section 2.3.2
1171	Source Router Identifier	Section 2.3.3
1172	L2 Bundle Member Attributes	Section 2.2.3
1174	Source OSPF Router-ID	Section 2.3.4

Table 8: Summary of TLV/Sub-TLV Code Points

4. 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 is 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 retrieves this information over a BGP-LS session (refer to Sections 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 any syntactic error in them would result in the BGP-LS Attribute being discarded with an error log. The SR information introduced in BGP-LS by this specification may be used by BGP-LS consumer applications like an SR Path Computation Engine (PCE) to learn the SR capabilities of the nodes in the topology and the mapping of SR segments to those nodes. This can enable the SR PCE to perform path computations based on SR for traffic engineering use cases and to steer traffic on paths different from the underlying IGP-based distributed best-path computation. Errors in the encoding or decoding of the SR information may result in the unavailability of such information to the SR PCE or incorrect information being made available to it. This may result in the SR PCE not being able to perform the desired SR-based optimization functionality or to perform it in an unexpected or inconsistent manner. The handling of such errors by applications like SR PCE may be implementation specific and out of scope of this document.

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 [RFC9020], [ISIS-SR-YANG], and [OSPF-SR-YANG].

5. Security Considerations

The new protocol extensions introduced in this document augment the existing IGP topology information that is distributed via [RFC7752]. The advertisement of the SR link attribute information defined in this document presents similar risk as associated with the existing

set of link attribute information as described in [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].

The TLVs introduced in this document are used to propagate IGP-defined information (see [RFC8665], [RFC8666], and [RFC8667]). These TLVs represent the SR information associated with the IGP node, link, and prefix. The IGP instances originating these TLVs are assumed to support all the required security and authentication mechanisms (as described in [RFC8665], [RFC8666], and [RFC8667]) in order to prevent any security issue when propagating the TLVs into BGP-LS.

BGP-LS SR extensions enable traffic engineering use cases within the SR domain. SR operates within a trusted domain [RFC8402], 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 ASes/domains within a single provider network). Therefore, precaution is necessary to ensure that the link-state information (including SR information) advertised via BGP-LS sessions is limited to consumers in a secure manner within this trusted SR domain. BGP peering sessions for address families other than link state may be set up to routers outside the SR domain. The isolation of BGP-LS peering sessions is recommended to ensure that BGP-LS topology information (including the newly added SR information) is not advertised to an external BGP peering session outside the SR domain.

6. References

6.1. 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, <<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>>.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <<https://www.rfc-editor.org/info/rfc5308>>.
- [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>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [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>>.
- [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>>.
- [RFC8571] Ginsberg, L., Ed., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions", RFC 8571, DOI 10.17487/RFC8571, March 2019, <<https://www.rfc-editor.org/info/rfc8571>>.
- [RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <<https://www.rfc-editor.org/info/rfc8665>>.
- [RFC8666] Psenak, P., Ed. and S. Previdi, Ed., "OSPFv3 Extensions for Segment Routing", RFC 8666, DOI 10.17487/RFC8666, December 2019, <<https://www.rfc-editor.org/info/rfc8666>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.
- [RFC8668] Ginsberg, L., Ed., Bashandy, A., Filsfils, C., Nanduri, M., and E. Aries, "Advertising Layer 2 Bundle Member Link Attributes in IS-IS", RFC 8668, DOI 10.17487/RFC8668, December 2019, <<https://www.rfc-editor.org/info/rfc8668>>.
- [RFC9084] Wang, A., Lindem, A., Dong, J., Psenak, P., and K. Talaulikar, Ed., "OSPF Prefix Originator Extensions", RFC 9084, DOI 10.17487/RFC9084, August 2021, <<https://www.rfc-editor.org/info/rfc9084>>.

6.2. Informative References

- [ISIS-SR-YANG] Litkowski, S., Qu, Y., Sarkar, P., Chen, I., and J. Tantsura, "YANG Data Model for IS-IS Segment Routing", Work in Progress, Internet-Draft, draft-ietf-isis-sr-yang-10, 21 February 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-isis-sr-yang-10>>.
- [OSPF-SR-YANG] Yeung, D., Qu, Y., Zhang, J., Chen, I., and A. Lindem, "YANG Data Model for OSPF SR (Segment Routing) Protocol", Work in Progress, Internet-Draft, draft-ietf-ospf-sr-yang-15, 2 July 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-ospf-sr-yang-15>>.
- [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>>.

[RFC8661] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., and S. Litkowski, "Segment Routing MPLS Interworking with LDP", RFC 8661, DOI 10.17487/RFC8661, December 2019, <<https://www.rfc-editor.org/info/rfc8661>>.

[RFC9020] Litkowski, S., Qu, Y., Lindem, A., Sarkar, P., and J. Tantsura, "YANG Data Model for Segment Routing", RFC 9020, DOI 10.17487/RFC9020, May 2021, <<https://www.rfc-editor.org/info/rfc9020>>.

Acknowledgements

The authors would like to thank Jeffrey Haas, Aijun Wang, Robert Raszuk, and Susan Hares for their review of this document and their comments. The authors would also like to thank Alvaro Retana for his extensive review and comments, which helped correct issues and improve the document.

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

Saikat Ray
Individual

Email: raysaikat@gmail.com

Jeff Tantsura
Apstra Inc.

Email: jefftant.ietf@gmail.com

Authors' Addresses

Stefano Previdi
Huawei Technologies
Rome
Italy

Email: stefano@previdi.net

Ketan Talaulikar (editor)
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