Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: November 9, 2019

S. Previdi
Huawei Technologies
K. Talaulikar, Ed.
C. Filsfils
Cisco Systems, Inc.
H. Gredler
RtBrick Inc.
M. Chen
Huawei Technologies
May 8, 2019

BGP Link-State extensions for Segment Routing draft-ietf-idr-bgp-ls-segment-routing-ext-14

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", "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.

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 November 9, 2019.

Table of Contents

1. Introduction 3
2. BGP-LS Extensions for Segment Routing 4
2.1. Node Attributes TLVs 5
2.1.1. SID/Label Sub-TLV 5
2.1.2. SR Capabilities TLV 6
2.1.3. SR Algorithm TLV 7
2.1.4. SR Local Block TLV 8
2.1.5. SRMS Preference TLV 10
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 Attribute TLV 14
2.3. Prefix Attribute TLVs 16
2.3.1. Prefix SID TLV 17
2.3.2. Prefix Attribute Flags TLV 18
2.3.3. Source Router Identifier (Source Router-ID) TLV . . . 19
2.3.4. Range TLV 19
2.4. Equivalent IS-IS Segment Routing TLVs/Sub-TLVs 21
2.5. Equivalent OSPFv2/OSPFv3 Segment Routing TLVs/Sub-TLVs . 22
3. IANA Considerations 25
3.1. TLV/Sub-TLV Code Points Summary 25
4. Manageability Considerations 25
5. Security Considerations 26
6. Contributors 27
7. Acknowledgements 28
8. References 28
8.1. Normative References 28
8.2. Informative References 30
Authors' Addresses 30

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).

[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 Segment Routing 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 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.

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.

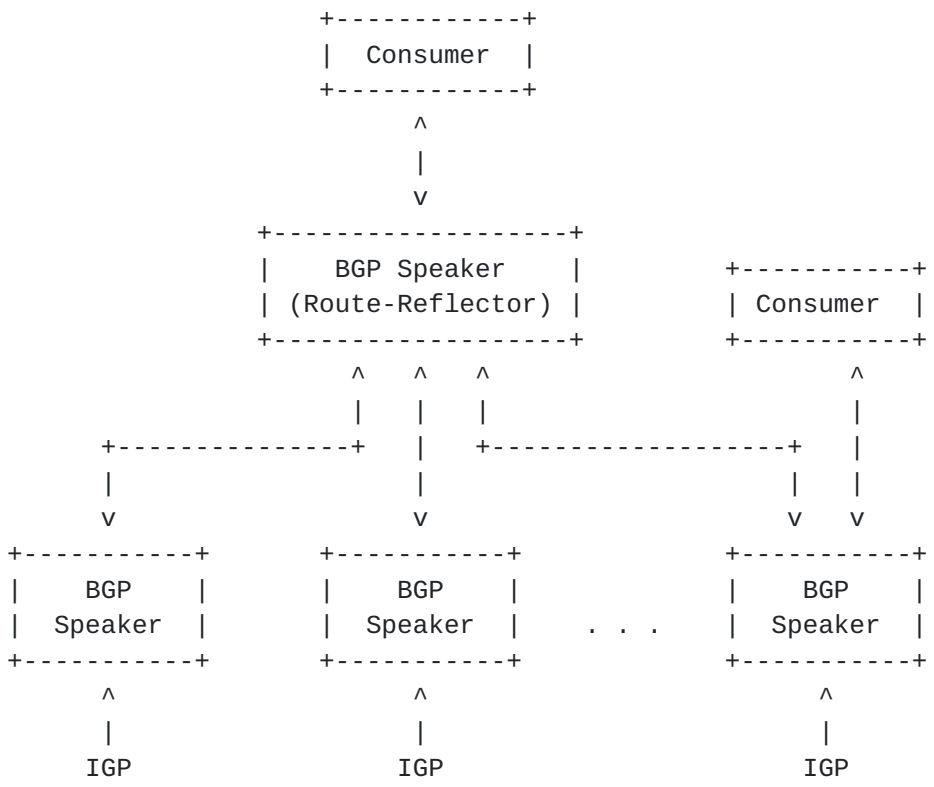

Figure 1: Link State info collection

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) as described in [RFC7752].

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 (as described in [RFC8402]) 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. The SR domain may be comprised of a single AS or multiple ASes.

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 lists 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.

2.1. Node Attributes 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 describing the IGP node that is originating the corresponding IGP TLV/sub-TLV described below.

2.1.1. SID/Label Sub-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 [I-D.ietf-isis-segment-routing-extensions].

- OSPFv2/OSPFv3, as defined by the SID/Label sub-TLV in [I-D.ietf-ospf-segment-routing-extensions] and [I-D.ietf-ospf-ospfv3-segment-routing-extensions].

The TLV and has the following format:

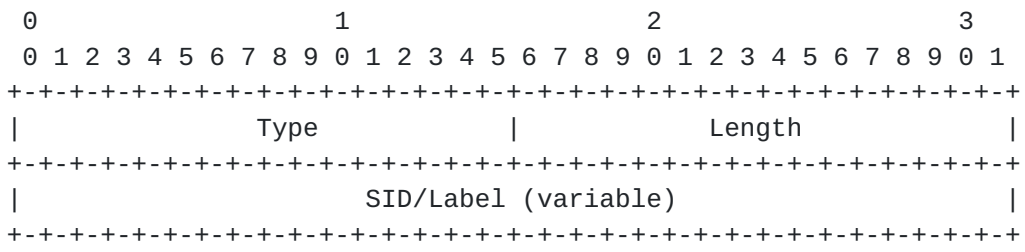

Figure 2: SID/Label sub-TLV Format

Where:

Type: 1161

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

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).

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

- OSPFv2/OSPFv3, as defined by the SID/Label Range TLV in [I-D.ietf-ospf-segment-routing-extensions]. OSPFv3 leverages the same TLV as defined for OSPFv2.

The SR Capabilities TLV has the following format:

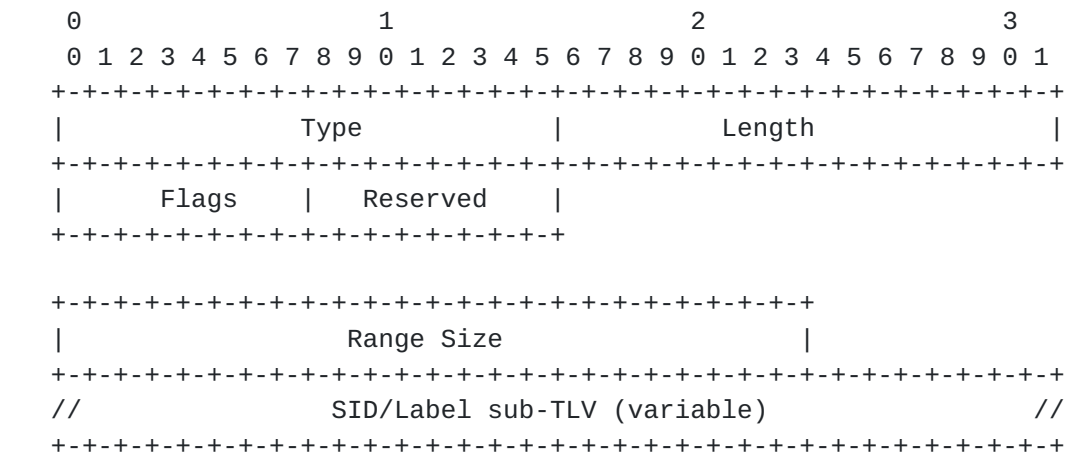

Figure 3: SR Capabilities TLV Format

Where:

Type: 1034

Length: Variable. Minimum length is 12.

Flags: 1 octet of flags as defined in [I-D.ietf-isis-segment-routing-extensions] for IS-IS. The flags are not currently defined for OSPFv2 and OSPFv3 and SHOULD be set to 0 and MUST be ignored on receipt.

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 with a non-zero 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. Since the SID/Label sub-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 [I-D.ietf-isis-segment-routing-extensions].

- OSPFv2/OSPFv3, as defined by the SR-Algorithm TLV in [I-D.ietf-ospf-segment-routing-extensions]. OSPFv3 leverages the same TLV as defined for OSPFv2.

The SR Algorithm TLV has the following format:

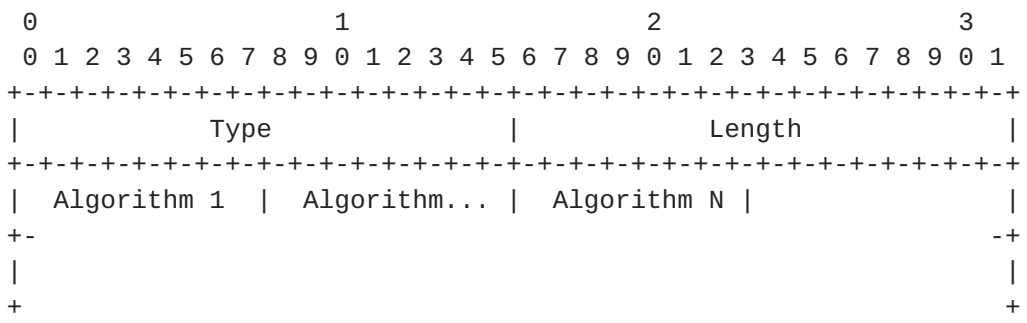

Figure 4: SR Algorithm TLV Format

Where:

Type: 1035

Length: Variable. Minimum length is 1 and maximum can be 256.

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.

- IS-IS, as defined by the SR Local Block sub-TLV in [I-D.ietf-isis-segment-routing-extensions].

o OSPFv2/OSPFv3, as defined by the SR Local Block TLV in [I-D.ietf-ospf-segment-routing-extensions]. OSPFv3 leverages the same TLV as defined for OSPFv2.

The SRLB TLV has the following format:

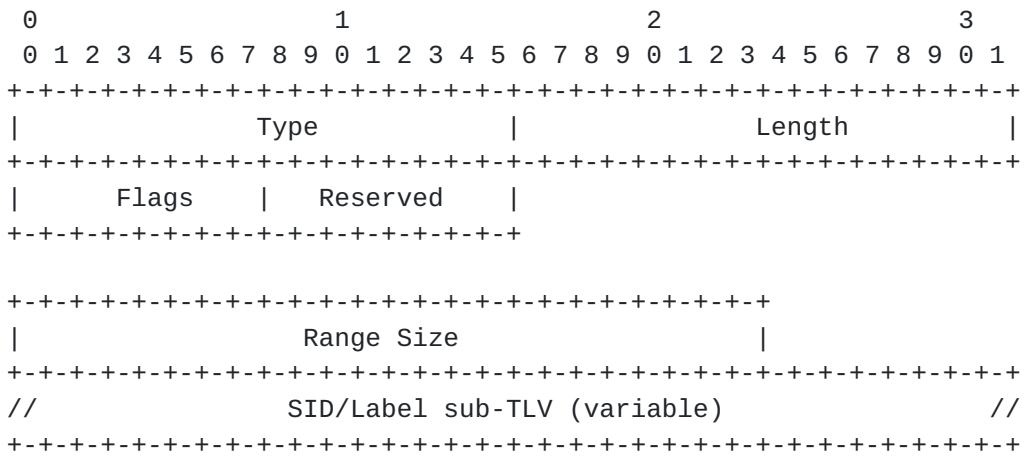

Figure 5: SRLB TLV Format

Where:

Type: 1036

Length: Variable. Minimum length is 12.

Flags: 1 octet of flags. The flags are as defined in [I-D.ietf-isis-segment-routing-extensions] for IS-IS. The flags are not currently defined for OSPFv2 and OSPFv3 and SHOULD be set to 0 and MUST be ignored on receipt.

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. Since the SID/Label sub-TLV is used to indicate the first label of the SRLB 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. [I-D.ietf-spring-segment-routing-ldp-interop] specifies the SRMS functionality along with 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 [I-D.ietf-isis-segment-routing-extensions].
- OSPFv2/OSPFv3, as defined by the SRMS Preference TLV in [I-D.ietf-ospf-segment-routing-extensions]. OSPFv3 leverages the same TLV as defined for OSPFv2.

The SRMS Preference TLV has the following format:

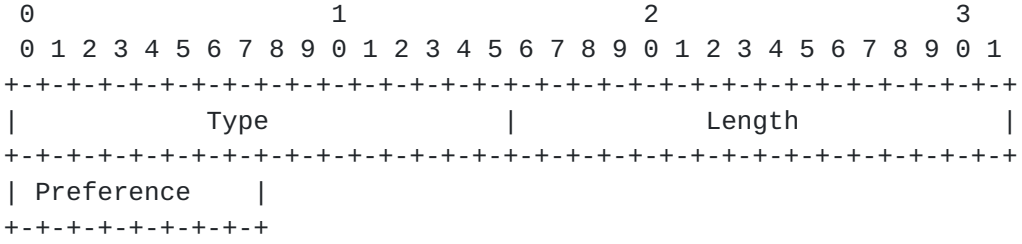

Figure 6: SRMS Preference TLV Format

Where:

Type: 1037

Length: 1.

Preference: 1 octet. Unsigned 8 bit SRMS preference.

2.2. Link Attribute TLVs

The following Link Attribute TLVs are 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 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 describing 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 Adj-SID sub-TLV of 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].

The Adjacency SID TLV has the following format:

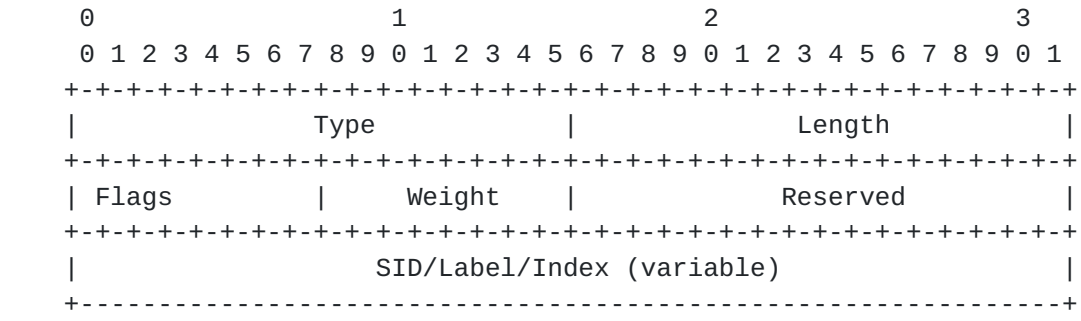

Figure 7: Adjacency SID TLV Format

Where:

Type: 1099

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

Flags. 1 octet value which should be set as:

* IS-IS Adj-SID flags are defined in [I-D.ietf-isis-segment-routing-extensions].

* OSPFv2 Adj-SID flags are defined in [I-D.ietf-ospf-segment-routing-extensions].

* OSPFv3 Adj-SID flags are defined in [I-D.ietf-ospf-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:

* 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].

The Flags and, as an extension, the SID/Index/Label fields of this TLV need to be interpreted accordingly to the respective underlying IS-IS, OSPFv2 or OSPFv3 protocol. The Protocol-ID of the BGP-LS Link NLRI should be 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 pseudo-node (or the equivalent OSPF Designated and Backup Designated Routers). 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.

This information is derived from LAN-Adj-SID sub-TLV of IS-IS [I-D.ietf-isis-segment-routing-extensions] and LAN Adj-SID sub-TLV of OSPFv2 [I-D.ietf-ospf-segment-routing-extensions] and OSPFv3 [I-D.ietf-ospf-ospfv3-segment-routing-extensions].

The LAN Adjacency SID TLV has the following format:

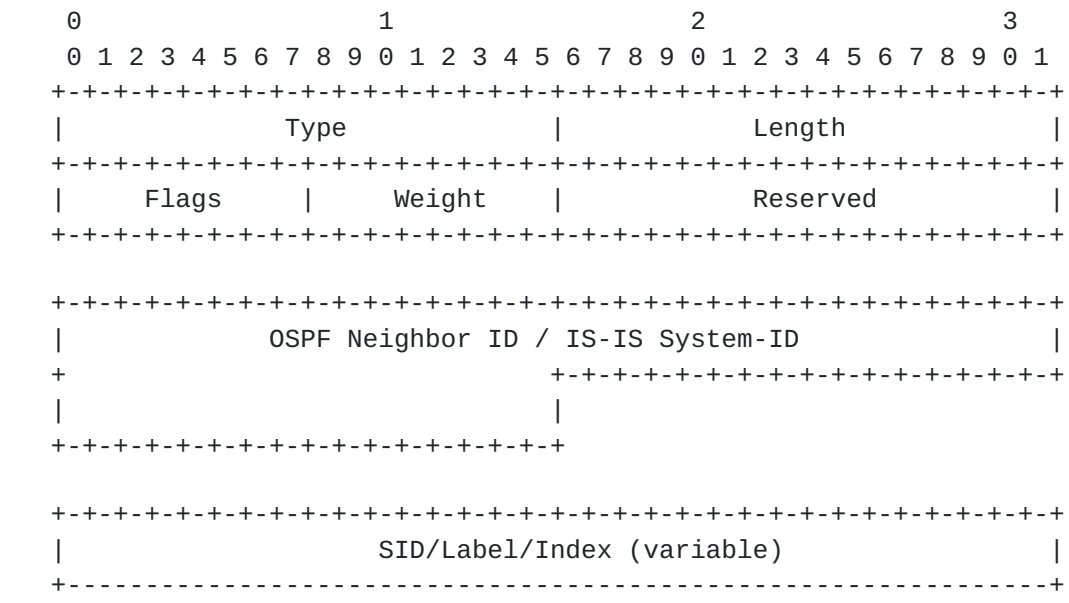

Figure 8: LAN Adjacency SID TLV Format

Where:

Type: 1100

Length: Variable. For IS-IS it would be 13 or 14 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 value which should be set as:

* IS-IS LAN Adj-SID flags are defined in [I-D.ietf-isis-segment-routing-extensions].

* OSPFv2 LAN Adj-SID flags are defined in [I-D.ietf-ospf-segment-routing-extensions].

* OSPFv3 LAN Adj-SID flags are defined in [I-D.ietf-ospf-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.

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 [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].

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

2.2.3. L2 Bundle Member Attribute TLV

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.

This information is derived from L2 Bundle Member Attributes TLV of IS-IS [I-D.ietf-isis-l2bundles]. The equivalent functionality has not been specified as yet for OSPF.

The L2 Bundle Member Attribute TLV has the following format:

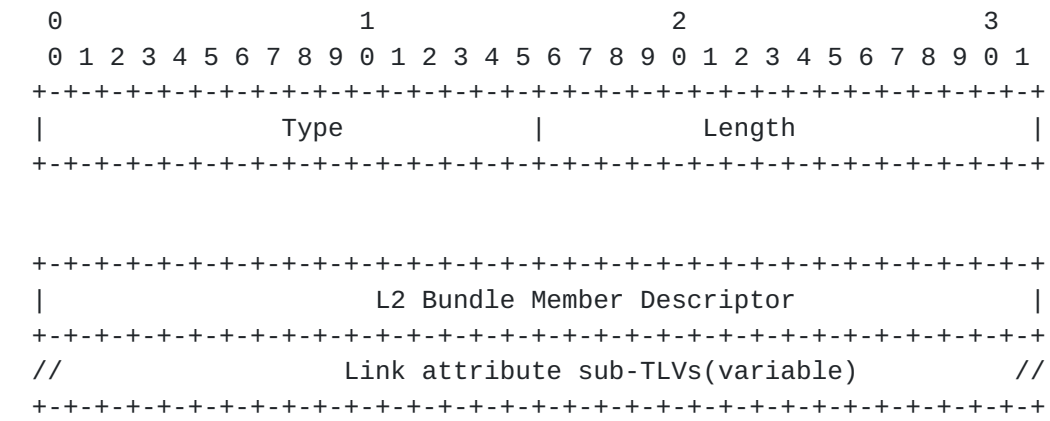

Figure 9: L2 Bundle Member Attributes TLV Format

Where:

Type: 1172

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 L2 Bundle 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	[RFC8571]
1115	Min/Max Unidirectional link delay	[RFC8571]
1116	Unidirectional Delay Variation	[RFC8571]
1117	Unidirectional packet loss	[RFC8571]
1118	Unidirectional residual bandwidth	[RFC8571]
1119	Unidirectional available bandwidth	[RFC8571]
1120	Unidirectional bandwidth utilization	[RFC8571]

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:

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

Table 4: Prefix Attribute TLVs

These TLVs should only be added to the BGP-LS Attribute associated with the Prefix NLRI describing 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 Prefix-SID sub-TLV of IS-IS [I-D.ietf-isis-segment-routing-extensions] and the Prefix SID sub-TLV of OSPFv2 [I-D.ietf-ospf-segment-routing-extensions] and OSPFv3 [I-D.ietf-ospf-ospfv3-segment-routing-extensions].

The Prefix SID TLV has the following format:

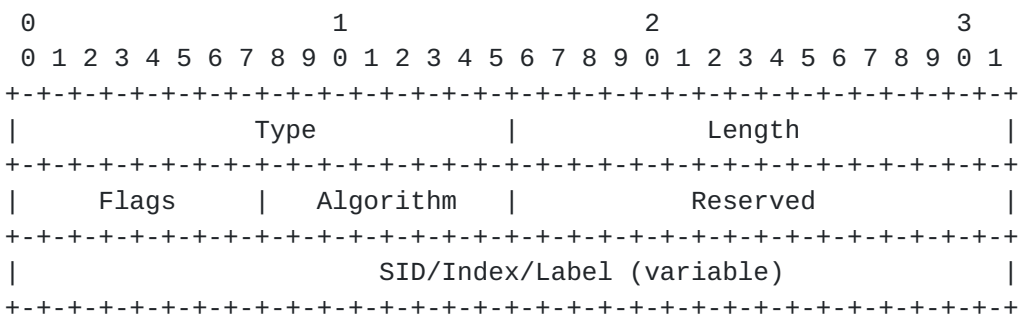

Figure 10: Prefix SID TLV Format

Where:

Type: 1158

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

Flags: 1 octet value which should be set as:

* IS-IS Prefix SID flags are defined in [I-D.ietf-isis-segment-routing-extensions].

* OSPFv2 Prefix SID flags are defined in [I-D.ietf-ospf-segment-routing-extensions].

* OSPFv3 Prefix SID flags are defined in [I-D.ietf-ospf-segment-routing-extensions].

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].

The Flags and, as an extension, the SID/Index/Label fields of this TLV need to be interpreted accordingly to the respective underlying IS-IS, OSPFv2 or OSPFv3 protocol. The Protocol-ID of the BGP-LS Prefix NLRI should be 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 [RFC7684], for OSPFv3 in [RFC5340] and for IS-IS in [RFC7794].

The Prefix Attribute Flags TLV has the following format:

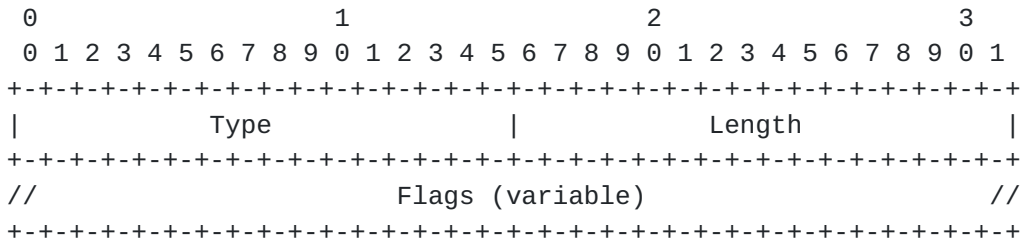

Figure 11: Prefix Attribute Flags TLV Format

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 [RFC7794]

* OSPFv2 flags correspond to the Flags field of the OSPFv2 Extended Prefix TLV defined in [RFC7684]

* OSPFv3 flags map to the Prefix Options field defined in [RFC5340] and extended via [RFC8362]

The Flags field of this TLV need to be interpreted accordingly to the respective underlying IS-IS, OSPFv2 or OSPFv3 protocol. The Protocol-ID of the BGP-LS Prefix NLRI should be used to determine the underlying protocol specification for parsing these fields.

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 the IS-IS protocol this is derived from the IPv4/IPv6 Source Router ID sub-TLV as defined in [RFC7794]. For the OSPF protocol, this is derived from the Prefix Source Router-ID sub-TLV as defined in [I-D.ietf-lsr-ospf-prefix-originator].

The Source Router-ID TLV has the following format:

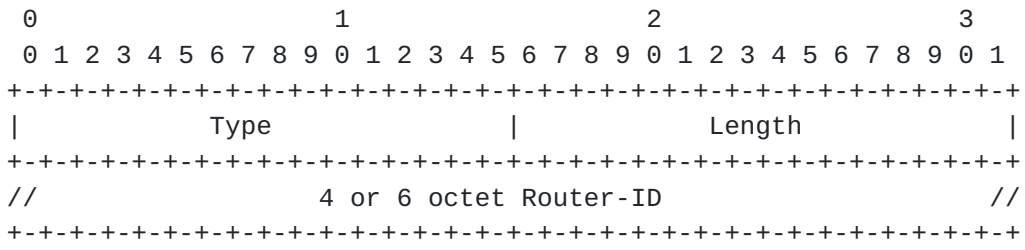

Figure 12: Source Router-ID TLV Format

Where:

Type: 1171

Length: 4 or 16 in case of IS-IS and 4 in case of OSPF.

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

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 (SRMS) 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 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 been advertised with a Range TLV, is considered as a normal routing prefix (i.e. prefix reachability) unless there is also an IGP metric TLV (TLV 1095) attached to it.

The format of the Range TLV is as follows:

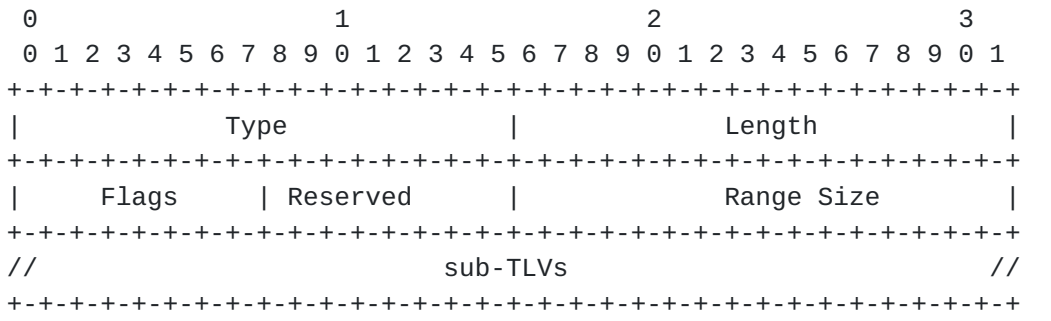

Figure 13: Range TLV Format

Where:

Type: 1159

Length: Variable, 11 or 12 depending on Label or Index encoding of the SID

Flags: 1 octet value which should be set as:

* IS-IS SID/Label Binding TLV flags are defined in [I-D.ietf-isis-segment-routing-extensions].

* OSPFv2 OSPF Extended Prefix Range TLV flags are defined in [I-D.ietf-ospf-segment-routing-extensions].

* OSPFv3 Extended Prefix Range TLV flags are defined in [I-D.ietf-ospf-ospfv3-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].

The Flags field of this TLV need to be interpreted accordingly to the respective underlying IS-IS, OSPFv2 or OSPFv3 protocol. The consumer of the BGP-LS interested in this TLV information MUST check the Protocol-ID of the BGP-LS Prefix NLRI and refer to the underlying protocol specification in order to parse 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 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 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.

Descriptio n	IS-IS TLV /sub-TLV	Reference
SR Capabil ities	SR-Capabilities sub-TLV (2)	[I-D.ietf-isis-segment-routing-exten sions]
SR Algorithm	SR-Algorithm sub-TLV (19)	[I-D.ietf-isis-segment-routing-exten sions]
SR Local Block	SR Local Block sub-TLV (22)	[I-D.ietf-isis-segment-routing-exten sions]
SRMS Preference	SRMS Preference sub-TLV (19)	[I-D.ietf-isis-segment-routing-exten sions]
Adjacency SID	Adj-SID sub-TLV (31)	[I-D.ietf-isis-segment-routing-exten sions]
LAN Adjacency SID	LAN-Adj-SID sub-TLV (32)	[I-D.ietf-isis-segment-routing-exten sions]
Prefix SID	Prefix-SID sub-TLV (3)	[I-D.ietf-isis-segment-routing-exten sions]
Range	SID/Label Binding TLV (149)	[I-D.ietf-isis-segment-routing-exten sions]
SID/Label	SID/Label sub-TLV (1)	[I-D.ietf-isis-segment-routing-exten sions]
Prefix Attribute Flags	Prefix Attributes Flags sub-TLV (4)	[RFC7794]
Source Router-ID	IPv4/IPv6 Source Router ID sub-TLV (11/12)	[RFC7794]
L2 Bundle Member Attributes	L2 Bundle Member Attributes TLV (25)	[I-D.ietf-isis-l2bundles]

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	OSPFv2 TLV /sub-TLV	Reference
SR Capabilities	SID/Label Range TLV (9)	[I-D.ietf-ospf-segment-routing-extensions]
SR Algorithm	SR-Algorithm TLV (8)	[I-D.ietf-ospf-segment-routing-extensions]
SR Local Block	SR Local Block TLV (14)	[I-D.ietf-ospf-segment-routing-extensions]
SRMS Preference	SRMS Preference TLV (15)	[I-D.ietf-ospf-segment-routing-extensions]
Adjacency SID	Adj-SID sub-TLV (2)	[I-D.ietf-ospf-segment-routing-extensions]
LAN Adjacency SID	LAN Adj-SID sub-TLV (3)	[I-D.ietf-ospf-segment-routing-extensions]
Prefix SID	Prefix SID sub-TLV (2)	[I-D.ietf-ospf-segment-routing-extensions]
Range	OSPF Extended Prefix Range TLV (2)	[I-D.ietf-ospf-segment-routing-extensions]
SID/Label	SID/Label sub-TLV (1)	[I-D.ietf-ospf-segment-routing-extensions]
Prefix Attribute Flags	Flags of OSPFv2 Extended Prefix TLV (1)	[RFC7684]
Source Router-ID	Prefix Source Router-ID sub-TLV (TBD)	[I-D.ietf-lsr-ospf-prefix-originator]

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

Description	OSPFv3 TLV /sub-TLV	Reference
SR Capabilities	SID/Label Range TLV (9)	[I-D.ietf-ospf-segment-routing-extensions]
SR Algorithm	SR-Algorithm TLV (8)	[I-D.ietf-ospf-segment-routing-extensions]
SR Local Block	SR Local Block TLV (14)	[I-D.ietf-ospf-segment-routing-extensions]
SRMS Preference	SRMS Preference TLV (15)	[I-D.ietf-ospf-segment-routing-extensions]
Adjacency SID	Adj-SID sub-TLV (5)	[I-D.ietf-ospf-ospfv3-segment-routing-extensions]
LAN Adjacency SID	LAN Adj-SID sub-TLV (6)	[I-D.ietf-ospf-ospfv3-segment-routing-extensions]
Prefix SID	Prefix SID sub-TLV (4)	[I-D.ietf-ospf-ospfv3-segment-routing-extensions]
Range	OSPFv3 Extended Prefix Range TLV (9)	[I-D.ietf-ospf-ospfv3-segment-routing-extensions]
SID/Label	SID/Label sub-TLV (7)	[I-D.ietf-ospf-ospfv3-segment-routing-extensions]
Prefix Attribute Flags	Prefix Option Fields of Prefix TLV types 3,5,6	[RFC8362]
Source Router-ID	Prefix Source Router-ID sub-TLV (TBD)	[I-D.ietf-lsr-ospf-prefix-originator]

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

3. IANA Considerations

Early allocation of codepoints has been done by IANA for this document from the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" 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 2.3.4
1161	SID/Label	Section 2.1.1
1170	Prefix Attribute Flags	Section 2.3.2
1171	Source Router-ID	Section 2.3.3
1172	L2 Bundle Member Attributes	Section 2.2.3

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

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 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 any syntactic error in them would result in only that specific 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 a 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 [I-D.ietf-spring-sr-yang], [I-D.ietf-isis-sr-yang] and [I-D.ietf-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 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 ([I-D.ietf-isis-segment-routing-extensions], [I-D.ietf-ospf-segment-routing-extensions] and [I-D.ietf-ospf-ospfv3-segment-routing-extensions]). 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 [I-D.ietf-isis-segment-routing-extensions], [I-D.ietf-ospf-segment-routing-extensions] and

[I-D.ietf-ospf-ospfv3-segment-routing-extensions]) in order to prevent any security issue when propagating the TLVs into BGP-LS. The advertisement of the link attribute information defined in this document presents no additional risk beyond that associated with the existing set of link attribute information already supported in [RFC7752].

BGP-LS SR extensions enable traffic engineering use-cases within the Segment Routing 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 AS/ domains within a single provider network). Therefore, precaution is necessary to ensure that the 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 setup 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. 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

7. 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.

8. References

8.1. Normative References

[I-D.ietf-isis-l2bundles]
Ginsberg, L., Bashandy, A., Filsfils, C., Nanduri, M., and E. Aries, "Advertising L2 Bundle Member Link Attributes in IS-IS", draft-ietf-isis-l2bundles-07 (work in progress), May 2017.

[I-D.ietf-isis-segment-routing-extensions]
Previdi, S., Ginsberg, L., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", draft-ietf-isis-segment-routing-extensions-24 (work in progress), April 2019.

[I-D.ietf-lsr-ospf-prefix-originator]
Wang, A., Lindem, A., Dong, J., Talaulikar, K., and P. Psenak, "OSPF Extension for Prefix Originator", draft-ietf-lsr-ospf-prefix-originator-00 (work in progress), February 2019.

[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-23 (work in progress), January 2019.

[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-27 (work in progress), December 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>.

[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>.

8.2. 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-05 (work in progress), March 2019.

[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-07 (work in progress), March 2019.

[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.

[I-D.ietf-spring-sr-yang]
Litkowski, S., Qu, Y., Lindem, A., Sarkar, P., and J. Tantsura, "YANG Data Model for Segment Routing", draft-ietf-spring-sr-yang-12 (work in progress), February 2019.

[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>.

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