Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: May 3, 2017

S. Previdi, Ed.
P. Psenak
C. Filsfils
Cisco Systems, Inc.
H. Gredler
RtBrick Inc.
M. Chen
Huawei Technologies
J. Tantsura
Individual
October 30, 2016

BGP Link-State extensions for Segment Routing
draft-gredler-idr-bgp-ls-segment-routing-ext-04

Abstract

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

This draft defines extensions to the BGP Link-state address-family in order to carry segment 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 http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on May 3, 2017.

Copyright Notice

Table of Contents

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

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.

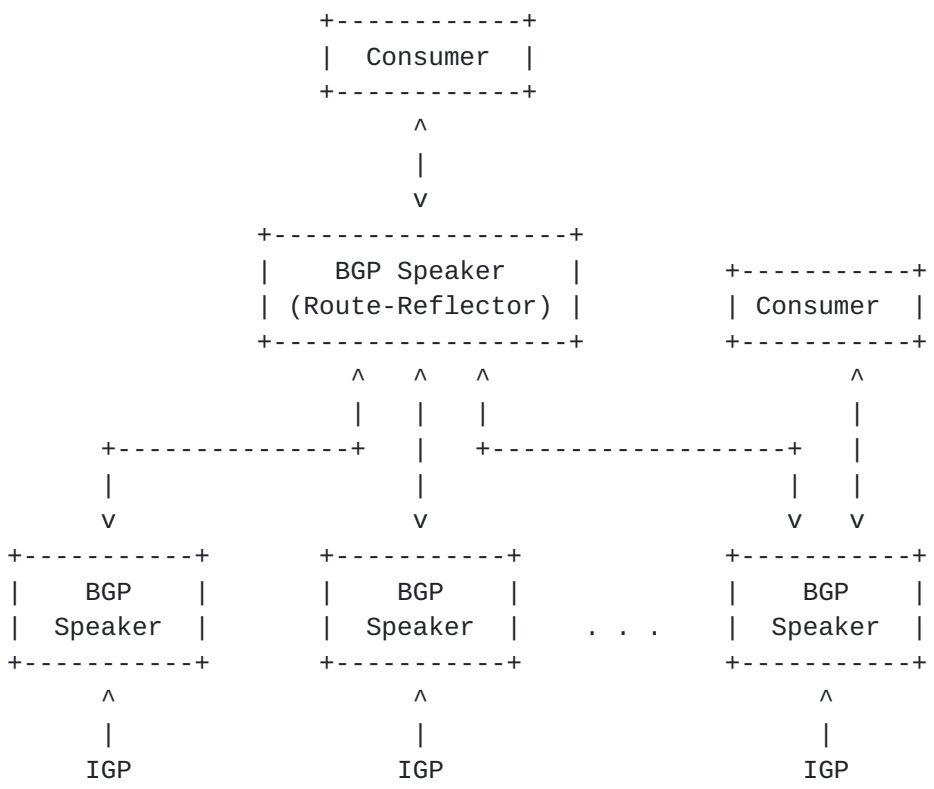

Figure 1: Link State info collection

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 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 in the "northbound" direction across IGP areas or ASes 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.

2. BGP-LS Extensions for Segment Routing

This document defines IGP SR extensions BGP-LS TLVs and Sub-TLVs. Section 2.4 and Section 2.5 illustrates the equivalent TLVs and Sub-TLVs in IS-IS, OSPF and OSPFv3 protocols.

BGP-LS [RFC7752] defines the BGP-LS NLRI that can be a Node NLRI, a Link NLRI or a Prefix NLRI. The corresponding BGP-LS attribute is a Node Attribute, a Link Attribute or a Prefix Attribute. BGP-LS [RFC7752] defines the TLVs that map link-state information to BGP-LS NLRI and the BGP-LS attribute. This document adds additional BGP-LS attribute TLVs in order to encode SR information.

2.1. Node Attributes TLVs

The following Node Attribute TLVs are defined:

TLV Code Point	Description	Length	Section
1034	SR Capabilities	variable	Section 2.1.1
1035	SR Algorithm	variable	Section 2.1.2
1036	SR Local Block	variable	Section 2.1.3
1037	SRMS Preference	variable	Section 2.1.4

Table 1: Node Attribute TLVs

These TLVs can ONLY be added to the Node Attribute associated with the Node NLRI that originates the corresponding SR TLV.

2.1.1. SR-Capabilities TLV

The SR Capabilities sub-TLV has following format:

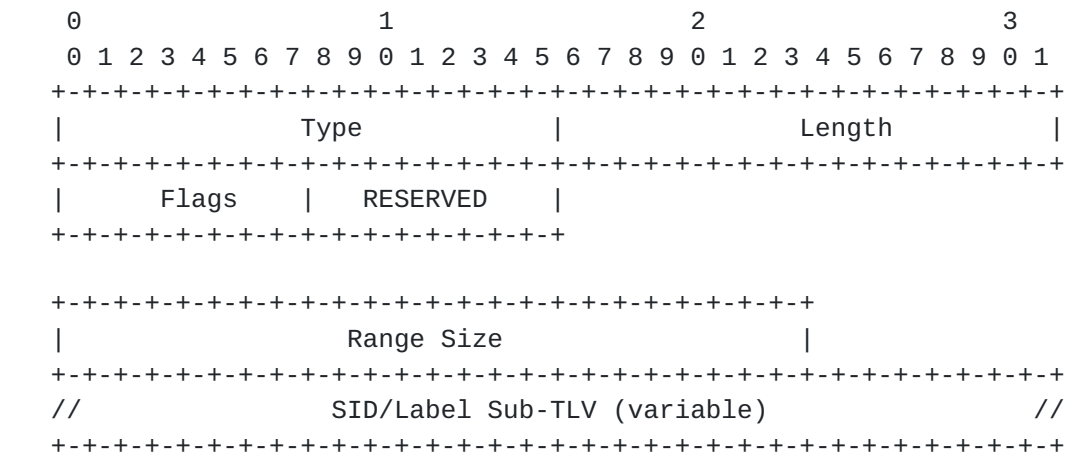

Type: TBD, suggested value 1034.

Length: Variable.

Flags: 1 octet of flags as defined in [I-D.ietf-isis-segment-routing-extensions] and [I-D.ietf-ospf-ospfv3-segment-routing-extensions].

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 2.3.7.2).

2.1.2. SR-Algorithm TLV

The SR-Algorithm TLV has the following format:

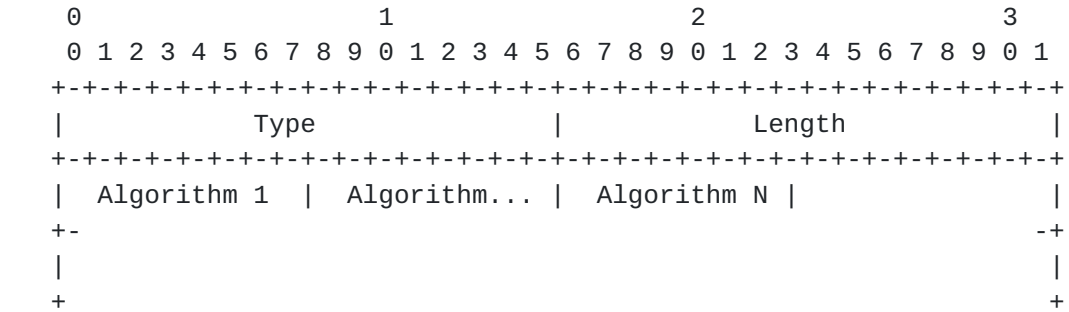

where:

Type: TBD, suggested value 1035.

Length: Variable.

Algorithm: 1 octet identifying the algorithm.

2.1.3. SR Local Block TLV

The SR Local Block (SRLB) Sub-TLV contains the range of labels the node has reserved for local SIDs. Local SIDs are used, e.g., in IGP (IS-IS, OSPF) for Adjacency-SIDs, and may also be allocated by other components 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.

The SRLB TLV has the following format:

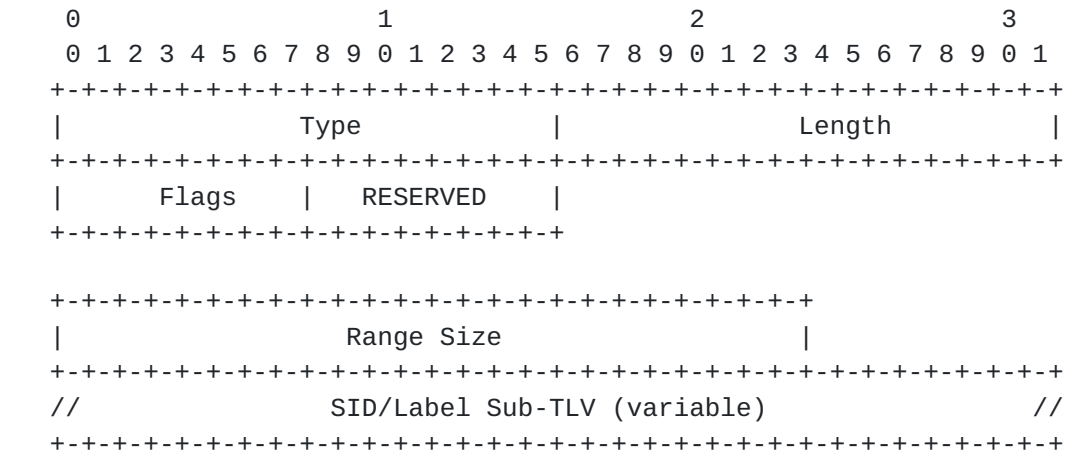

Type: TBD, suggested value 1036.

Length: Variable.

Flags: 1 octet of flags. None are defined at this stage.

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 2.3.7.2).

2.1.4. SRMS Preference TLV

The Segment Routing Mapping Server (SRMS) Preference sub-TLV is used in order to associate a preference with SRMS advertisements from a particular source.

The SRMS Preference sub-TLV has following format:

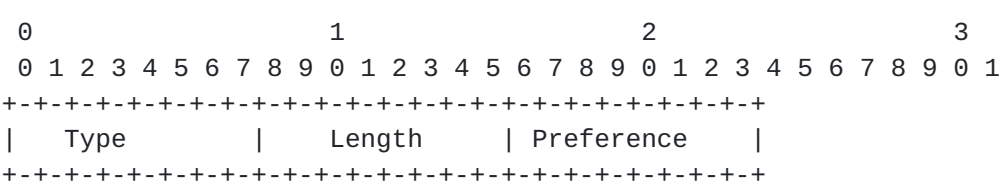

Type: TBD, suggested value 1037.

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].

2.2. Link Attribute TLVs

The following Link Attribute TLVs are are defined:

TLV Code Point	Description	Length	Section
1099	Adjacency Segment Identifier (Adj-SID) TLV	variable	Section 2.2.1
1100	LAN Adjacency Segment Identifier (Adj-SID) TLV	variable	Section 2.2.2

Table 2: Link Attribute TLVs

These TLVs can ONLY be added to the Link Attribute associated with the link whose local node originates the corresponding 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)[I-D.ietf-isis-segment-routing-extensions]. The LAN Adjecency 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:

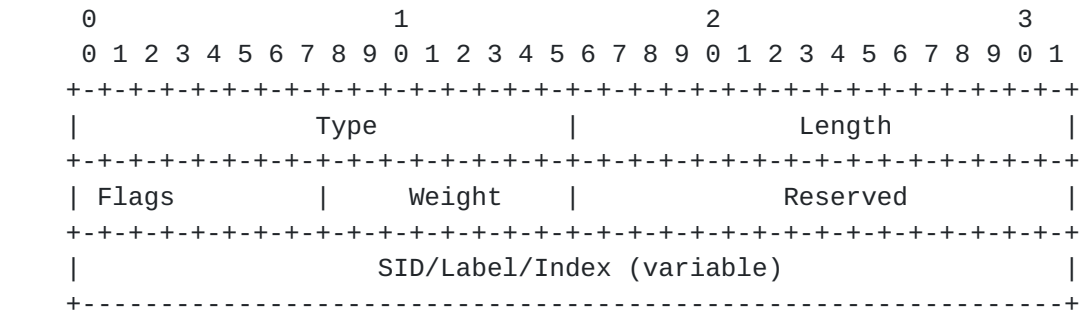

where:

Type: TBD, suggested value 1099.

Length: Variable.

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.

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].

2.2.2. LAN Adjacency SID TLV

The LAN Adjacency SID (LAN-Adj-SID-SID) has the following format:

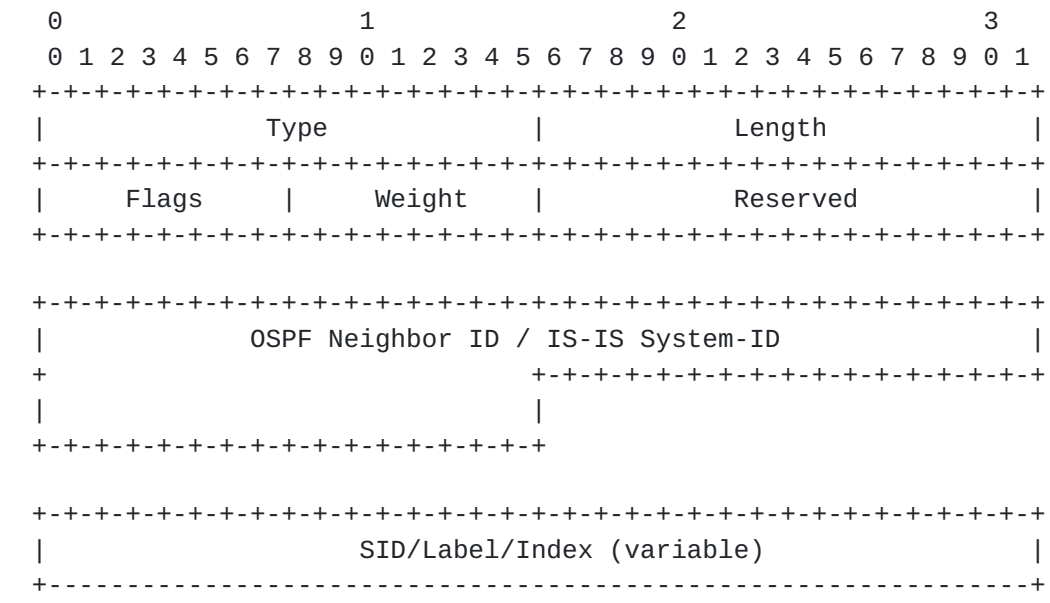

where:

Type: TBD, suggested value 1100.

Length: Variable.

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.

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].

2.3. Prefix Attribute TLVs

The following Prefix Attribute TLVs and Sub-TLVs are defined:

TLV Code Point	Description	Length	Section
1158	Prefix SID	variable	Section 2.3.1
1159	Range	variable	Section 2.3.5
1160	Binding SID	variable	Section 2.3.6
1169	IPv6 Prefix SID	variable	Section 2.3.2
1170	IGP Prefix Attributes	variable	Section 2.3.3
1171	Source Router-ID	variable	Section 2.3.4

Table 3: Prefix Attribute TLVs

TLV Code Point	Description	Length	Section
1161	SID/Label TLV	variable	Section 2.3.7.2
1162	ERO Metric TLV	4 octets	Section 2.3.7.3
1163	IPv4 ERO TLV	8 octets	Section 2.3.7.4
1164	IPv6 ERO TLV	20 octets	Section 2.3.7.5
1165	Unnumbered Interface ID ERO TLV	12	Section 2.3.7.6
1166	IPv4 Backup ERO TLV	8 octets	Section 2.3.7.7
1167	IPv6 Backup ERO TLV	10 octets	Section 2.3.7.8
1168	Unnumbered Interface ID Backup ERO TLV	12	Section 2.3.7.9

Table 4: Prefix Attribute - Binding SID Sub-TLVs

2.3.1. Prefix-SID TLV

The Prefix-SID TLV can ONLY be added to the Prefix Attribute whose local node in the corresponding Prefix NLRI is the node that originates the corresponding SR TLV.

The Prefix-SID has the following format:

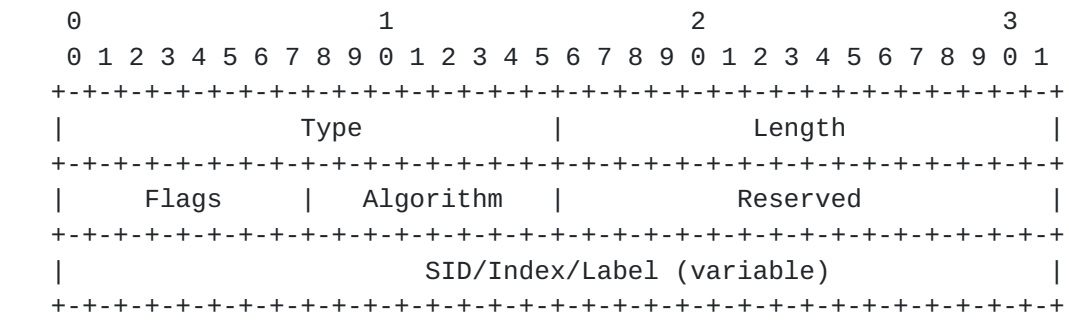

where:

Type: TBD, suggested value 1158.

Length: Variable

Algorithm: 1 octet value identify the algorithm.

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].

The Prefix-SID TLV includes a Flags field. In the context of BGP-LS, the Flags field format and the semantic of each individual flag MUST be taken from the corresponding source protocol (i.e.: the protocol of origin of the Prefix-SID being advertised in BGP-LS).

IS-IS Prefix-SID flags are defined in [I-D.ietf-isis-segment-routing-extensions] section 2.1.

OSPF 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.

2.3.2. IPv6 Prefix-SID TLV

The IPv6 Prefix-SID TLV can ONLY be added to the Prefix Attribute whose local node in the corresponding Prefix NLRI is the node that originates the corresponding SR TLV.

The IPv6 Prefix-SID has the following format:

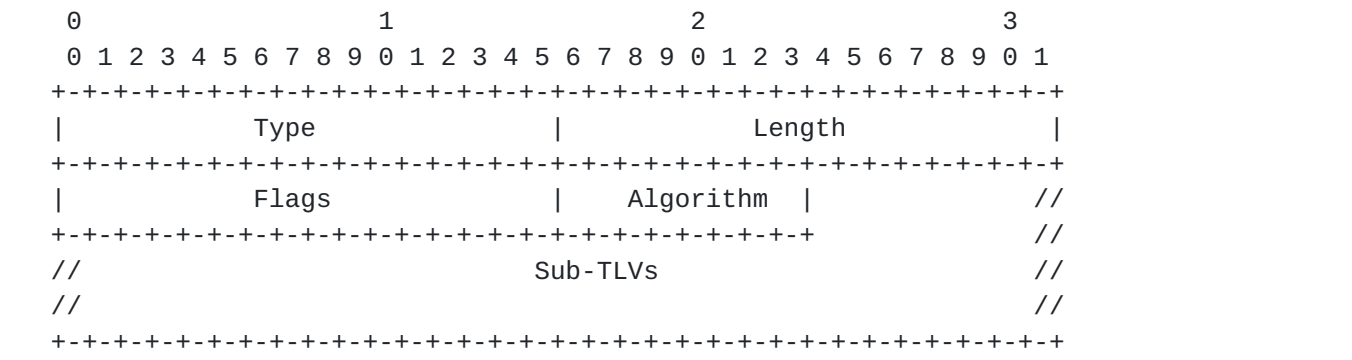

where:

Type: TBD, suggested value 1169.

Length: 3 + length of Sub-TLVs.

Flags: 2 octet field of flags. None of them is defined at this stage.

Algorithm: 1 octet value identify the algorithm as defined in [I-D.previdi-isis-ipv6-prefix-sid].

Sub-TLVs: additional information encoded into the IPv6 Prefix-SID Sub-TLV as defined in [I-D.previdi-isis-ipv6-prefix-sid].

The IPv6 Prefix-SID TLV is defined in [I-D.previdi-isis-ipv6-prefix-sid].

2.3.3. IGP Prefix Attributes TLV

The IGP Prefix Attribute TLV carries IPv4/IPv6 prefix attribute flags as defined in [RFC7684] and [RFC7794].

The IGP Prefix Attribute TLV has the following format:

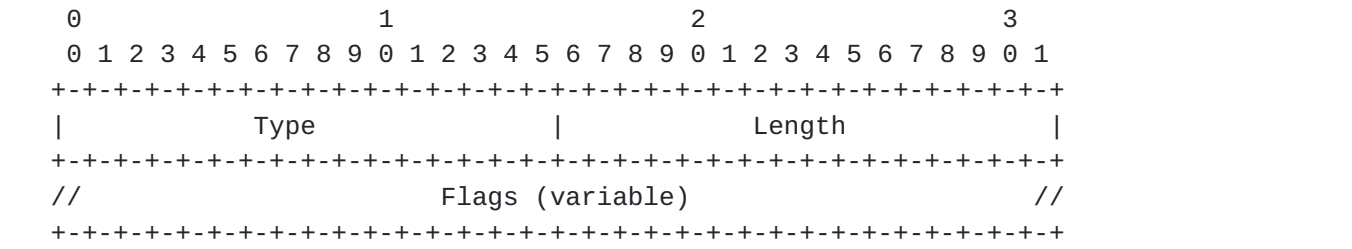

where:

Type: TBD, suggested value 1170.

Length: variable.

Flags: a variable length flag field (according to the length field). Flags are routing protocol specific (OSPF and IS-IS). OSPF flags are defined in [RFC7684] and IS-IS flags are defined in [RFC7794]. The receiver of the BGP-LS update, when inspecting the IGP Prefix Attribute TLV, MUST check the Protocol-ID of the NLRI and refer to the protocol specification in order to parse the flags.

2.3.4. Source Router Identifier (Source Router-ID) TLV

The Source Router-ID TLV contains the IPv4 or IPv6 Router-ID of the originator as defined in [RFC7794]. While defined in the IS-IS protocol, 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:

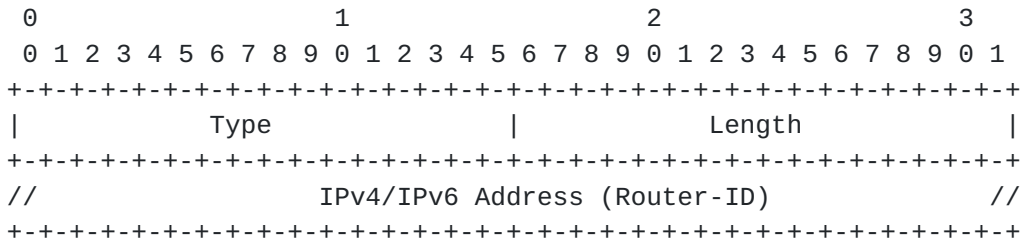

where:

Type: TBD, suggested value 1171.

Length: 4 or 16.

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

The semantic of the Source Router-ID TLV is defined in [RFC7794].

2.3.5. Range TLV

The Range TLV can ONLY be added to the Prefix Attribute whose local node in the corresponding Prefix NLRI is the node that originates the corresponding SR TLV.

When the range TLV is used in order to advertise a path to a prefix or a range of prefix-to-SID mappings, the Prefix-NLRI the Range TLV is attached to 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:

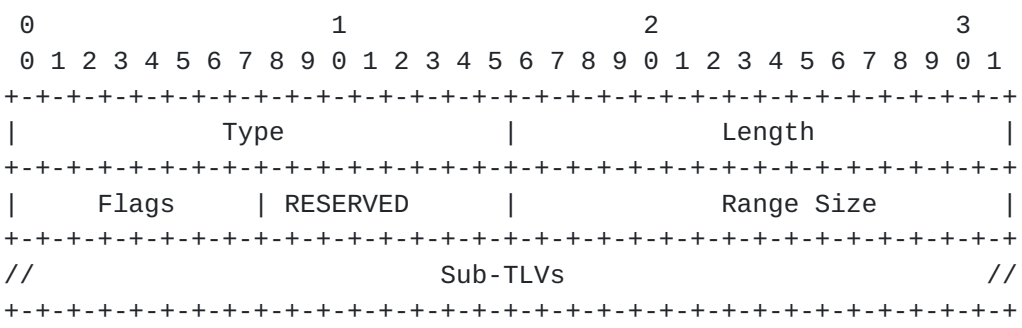

where:

Figure 2: Range TLV format

Type: 1159

Length is 4.

Flags: Only used when the source protocol is OSPF and defined in [I-D.ietf-ospf-segment-routing-extensions] section 4 and [I-D.ietf-ospf-ospfv3-segment-routing-extensions] section 4.

Range Size: 2 octets as defined in [I-D.ietf-ospf-segment-routing-extensions] section 4.

Within the Range TLV, the following SubTLVs are may be present:

Binding SID TLV, defined in Section 2.3.6

Prefix-SID TLV, defined in Section 2.3.1

SID/Label TLV, defined in Section 2.3.7.2

2.3.6. Binding SID TLV

The Binding SID TLV can be used in two ways:

- as a sub-TLV of the Range TLV
- as a Prefix Attribute TLV

The format of the Binding SID TLV is as follows:

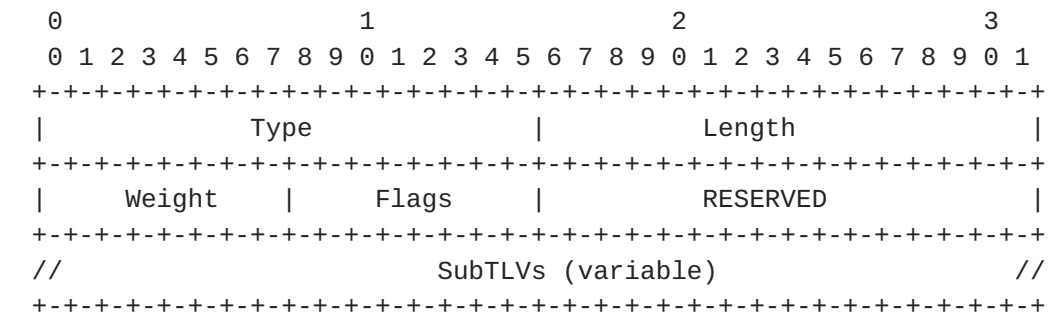

where:

Figure 3: Binding SID Sub-TLV format

Type is 1160

Length is variable

Weight and Flags are mapped to Weight and Flags defined in [I-D.ietf-isis-segment-routing-extensions] section 2.4, [I-D.ietf-ospf-segment-routing-extensions] section 4 and [I-D.ietf-ospf-ospfv3-segment-routing-extensions] section 4.

Sub-TLVs are defined in the following sections.

2.3.7. Binding SID SubTLVs

This section defines the Binding SID Sub-TLVs in BGP-LS to encode the equivalent Sub-TLVs 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].

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

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

2.3.7.1. Binding SID Prefix-SID Sub-TLV

When encoding IS-IS Mapping Server entries as defined in [I-D.ietf-isis-segment-routing-extensions] the Prefix-SID TLV defined in Section 2.3.1 is used as Sub-TLV in the Binding TLV.

2.3.7.2. SID/Label Sub-TLV

The SID/Label TLV has following format:

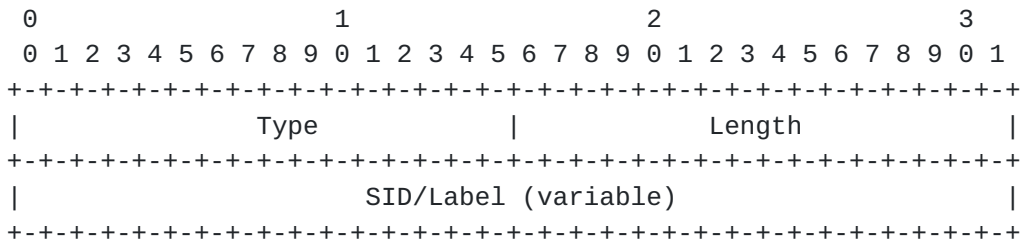

where:

Type: TBD, suggested value 1161.

Length: Variable, 3 or 4 bytes

SID/Label: If length is set to 3, then the 20 rightmost bits represent a label. If length is set to 4, then the value represents a 32 bit SID.

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

2.3.7.3. ERO Metric Sub-TLV

The ERO Metric Sub-TLV has following format:

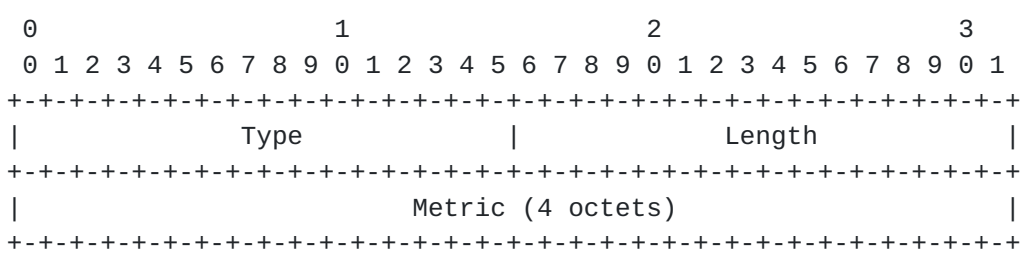

ERO Metric Sub-TLV format

where:

Type: TBD, suggested value 1162.

Length: Always 4

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

2.3.7.4. IPv4 ERO Sub-TLV

The ERO Sub-TLV has following format:

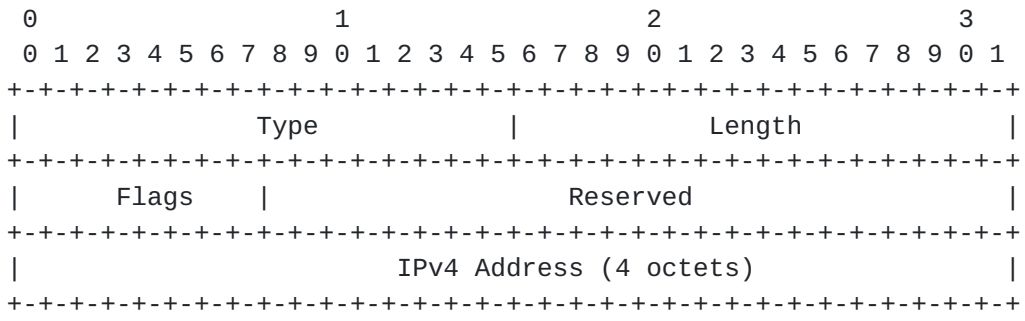

IPv4 ERO Sub-TLV format

where:

Type: TBD, suggested value 1163

Length: 8 octets

Flags: 1 octet of 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].

IPv4 Address - the address of the explicit route hop.

2.3.7.5. IPv6 ERO Sub-TLV

The IPv6 ERO Sub-TLV has following format:

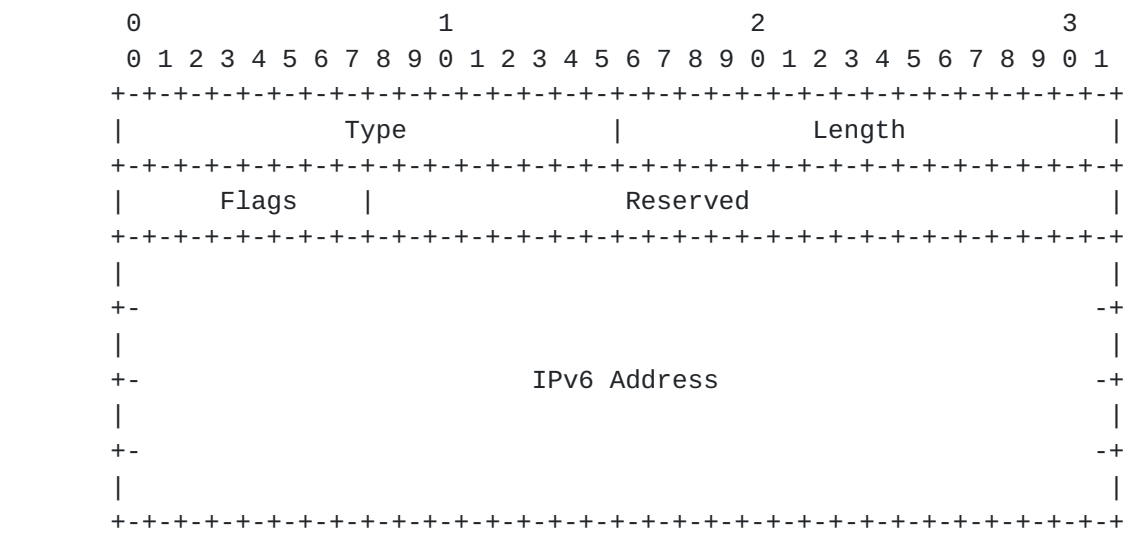

IPv6 ERO Sub-TLV format

where:

Type: TBD, suggested value 1164

Length: 20 octets

Flags: 1 octet of 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].

IPv6 Address - the address of the explicit route hop.

2.3.7.6. Unnumbered Interface ID ERO Sub-TLV

The Unnumbered Interface-ID ERO Sub-TLV has following format:

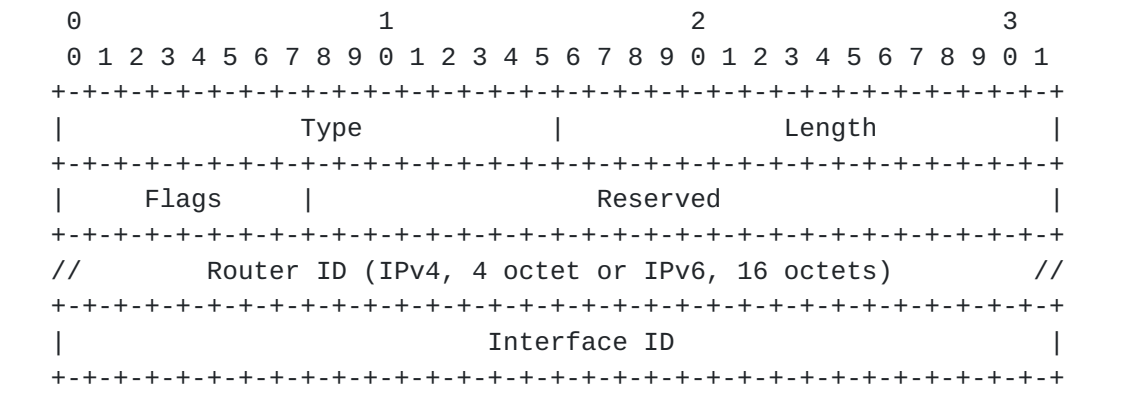

where:

Unnumbered Interface ID ERO Sub-TLV format

Type: TBD, suggested value 1165.

Length: Variable (12 for IPv4 Router-ID or 24 for IPv6 Router-ID).

Flags: 1 octet of 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].

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

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

2.3.7.7. IPv4 Backup ERO Sub-TLV

The IPv4 Backup ERO Sub-TLV has following format:

IPv4 Backup ERO Sub-TLV format

where:

Type: TBD, suggested value 1166.

Length: 8 octets

Flags: 1 octet of 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].

IPv4 Address: Address of the explicit route hop.

2.3.7.8. IPv6 Backup ERO Sub-TLV

The IPv6 Backup ERO Sub-TLV has following format:

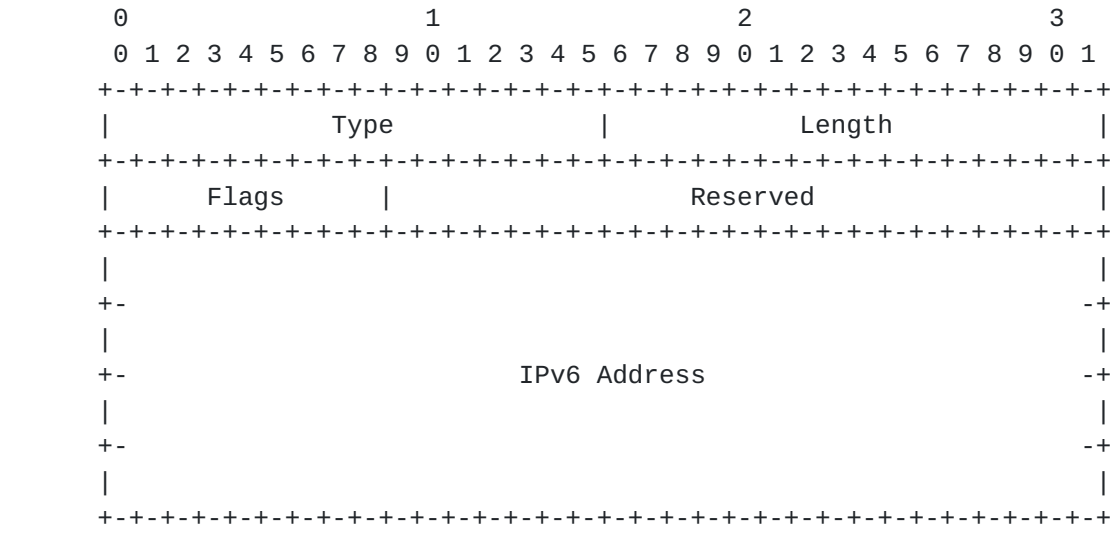

IPv6 Backup ERO Sub-TLV format

where:

Type: TBD, suggested value 1167.

Length: 8 octets

Flags: 1 octet of 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].

IPv6 Address: Address of the explicit route hop.

2.3.7.9. Unnumbered Interface ID Backup ERO Sub-TLV

The Unnumbered Interface-ID Backup ERO Sub-TLV has following format:

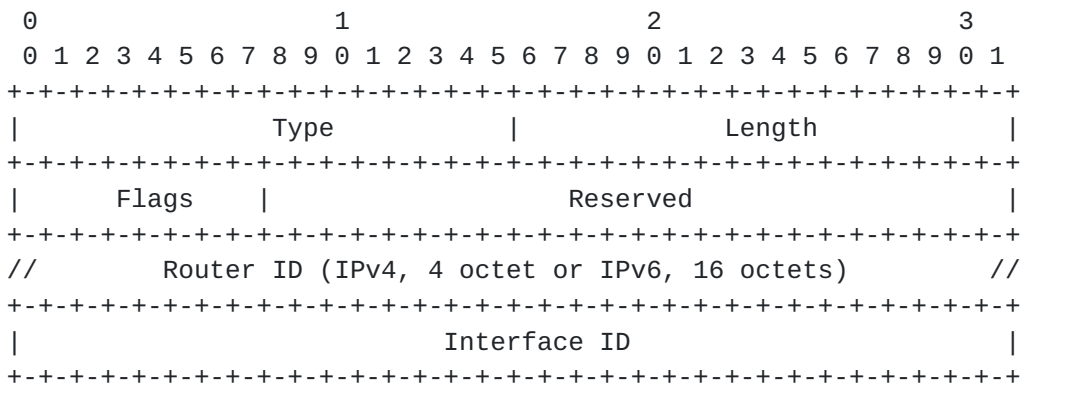

Unnumbered Interface ID Backup ERO Sub-TLV format

where:

Type: TBD, suggested value 1168.

Length: Variable (12 for IPv4 Router-ID or 24 for IPv6 Router-ID).

Flags: 1 octet of 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].

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

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

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.

TLV Code Point	Description	Length	IS-IS TLV /Sub-TLV
1034	SR Capabilities	variable	2 [1]
1035	SR Algorithm	variable	19 [2]
1099	Adjacency Segment Identifier (Adj-SID) TLV	variable	31 [3]
1100	LAN Adjacency Segment Identifier (LAN-Adj-SID) TLV	variable	32 [4]
1158	Prefix SID	variable	3 [5]
1160	Binding SID	variable	149 [6]
1161	SID/Label TLV	variable	1 [7]
1162	ERO Metric TLV	4 octets	10 [8]
1163	IPv4 ERO TLV	5 octets	11 [9]
1164	IPv6 ERO TLV	17 octets	12 [10]
1165	Unnumbered Interface ID ERO TLV	variable	13 [11]
1166	IPv4 Backup ERO TLV	5 octets	14 [12]
1167	IPv6 Backup ERO TLV	17 octets	15 [13]
1168	Unnumbered Interface ID Backup ERO TLV	variable	16 [14]
1169	IPv6 Prefix SID	variable	5 [15]
1170	IGP Prefix Attributes	variable	4 [16]
1171	Source Router ID	variable	11/12 [17]

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

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

This section illustrate the OSPF 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 OSPF and OSPFv3.

TLV Code Point	Description	Length	OSPF TLV /Sub-TLV
1034	SR Capabilities	variable	9 [18]
1035	SR Algorithm	variable	8 [19]
1099	Adjacency Segment Identifier (Adj-SID) TLV	variable	2 [20]
1100	LAN Adjacency Segment Identifier (Adj-SID) TLV	variable	3 [21]
1158	Prefix SID	variable	2 [22]
1161	SID/Label TLV	variable	1 [23]
1162	ERO Metric TLV	4 octets	8 [24]
1163	IPv4 ERO TLV	8 octets	4 [25]
1165	Unnumbered Interface ID ERO TLV	12 octets	5 [26]
1166	IPv4 Backup ERO TLV	8 octets	6 [27]
1167	Unnumbered Interface ID Backup ERO TLV	12 octets	7 [28]
1167	Unnumbered Interface ID Backup ERO TLV	12 octets	7 [29]

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

TLV Code Point	Description	Length	OSPFv3 TLV /Sub-TLV
1034	SR Capabilities	variable	9 [30]
1035	SR Algorithm	variable	8 [31]
1099	Adjacency Segment Identifier (Adj-SID) TLV	variable	5 [32]
1100	LAN Adjacency Segment Identifier (Adj-SID) TLV	variable	6 [33]
1158	Prefix SID	variable	4 [34]
1161	SID/Label TLV	variable	3 [35]
1162	ERO Metric TLV	4 octets	8 [36]
1163	IPv4 ERO TLV	8 octets	9 [37]
1164	IPv6 ERO TLV	20 octets	8 [38]
1165	Unnumbered Interface ID ERO TLV	12 octets	11 [39]
1166	IPv4 Backup ERO TLV	8 octets	12 [40]
1167	IPv6 Backup ERO TLV	20 octets	13 [41]
1167	Unnumbered Interface ID Backup ERO TLV	12 octets	14 [42]

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

3. Procedures

The following sections describe the different operations for the propagation of SR TLVs into BGP-LS.

3.1. Advertisement of a IS-IS Prefix SID TLV

The advertisement of a IS-IS Prefix SID TLV has following rules:

The IS-IS Prefix-SID is encoded in the BGP-LS Prefix Attribute Prefix-SID as defined in Section 2.3.1. The flags in the Prefix-SID TLV have the semantic defined in [I-D.ietf-isis-segment-routing-extensions] section 2.1.

3.2. Advertisement of a OSPF/OSPFv3 Prefix-SID TLV

The advertisement of a OSPF/OSPFv3 Prefix-SID TLV has following rules:

The OSPF (or OSPFv3) Prefix-SID is encoded in the BGP-LS Prefix Attribute Prefix-SID as defined in Section 2.3.1. The flags in

the Prefix-SID TLV have the semantic defined in [I-D.ietf-ospf-segment-routing-extensions] section 5 or [I-D.ietf-ospf-ospfv3-segment-routing-extensions] section 5.

3.3. Advertisement of a range of prefix-to-SID mappings in OSPF

The advertisement of a range of prefix-to-SID mappings in OSPF has following rules:

The OSPF/OSPFv3 Extended Prefix Range TLV is encoded in the BGP-LS Prefix Attribute Range TLV as defined in Section 2.3.5. 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. The Prefix-SID from the original OSPF Prefix SID Sub-TLV is encoded using the BGP-LS Prefix Attribute Prefix-SID 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.

3.4. Advertisement of a range of IS-IS SR bindings

The advertisement of a range of IS-IS SR bindings has following rules:

In IS-IS the Mapping Server binding ranges are advertised using the Binding TLV. The IS-IS Binding TLV is encoded in the BGP-LS Prefix Attribute Range TLV as defined in Section 2.3.5 using the Binding Sub-TLV as defined in Section 2.3.6. The flags in the Range TLV are all set to zero on transmit and ignored on reception. The range value from the original IS-IS Binding TLV is encoded in the Range TLV "Range" field.

3.5. Advertisement of a path and its attributes from IS-IS protocol

The advertisement of a Path and its attributes is described in [I-D.ietf-isis-segment-routing-extensions] section 2.4 and has following rules:

The original Binding SID TLV (from IS-IS) is encoded into the BGP-LS Range TLV defined in Section 2.3.5 using the Binding Sub-TLV as defined in Section 2.3.6. The set of Sub-TLVs from the original IS-IS Binding TLV are encoded as Sub-TLVs of the BGP-LS Binding TLV as defined in Section 2.3.6. This includes the SID/Label TLV defined in Section 2.3.

3.6. Advertisement of a path and its attributes from OSPFv2/OSPFv3 protocol

The advertisement of a Path and its attributes is described in [I-D.ietf-ospf-segment-routing-extensions] section 6 and [I-D.ietf-ospf-ospfv3-segment-routing-extensions] section 6 and has following rules:

Advertisement of a path for a single prefix: the original Binding SID TLV (from OSPFv2/OSPFv3) is encoded into the BGP-LS Prefix Attribute Binding TLV as defined in Section 2.3.6. The set of Sub-TLVs from the original OSPFv2/OSPFv3 Binding TLV are encoded as Sub-TLVs of the BGP-LS Binding TLV as defined in Section 2.3.6. This includes the SID/Label TLV defined in Section 2.3.

Advertisement of an SR path for range of prefixes: the OSPF/OSPFv3 Extended Prefix Range TLV is encoded in the BGP-LS Prefix Attribute Range TLV as defined in Section 2.3.5. The original OSPFv2/OSPFv3 Binding SID TLV is encoded into the BGP-LS Binding Sub-TLV as defined in Section 2.3.6. The set of Sub-TLVs from the original OSPFv2/OSPFv3 Binding TLV are encoded as Sub-TLVs of the BGP-LS Binding TLV as defined in Section 2.3.6. This includes the SID/Label TLV defined in Section 2.3.

4. IANA Considerations

This document requests assigning code-points from the registry for BGP-LS attribute TLVs based on table Table 8.

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	Length	Section
1034	SR Capabilities	variable	Section 2.1.1
1035	SR Algorithm	variable	Section 2.1.2
1036	SR Local Block	variable	Section 2.1.3
1037	SRMS Preference	variable	Section 2.1.4
1099	Adjacency Segment Identifier (Adj-SID) TLV	variable	Section 2.2.1
1100	LAN Adjacency Segment Identifier (Adj-SID) TLV	variable	Section 2.2.2
1158	Prefix SID	variable	Section 2.3.1
1159	Range	variable	Section 2.3.5
1160	Binding SID	variable	Section 2.3.6
1161	SID/Label TLV	variable	Section 2.3.7.2
1162	ERO Metric TLV	4 octets	1 [43]
1163	IPv4 ERO TLV	8 octets	1 [44]
1164	IPv6 ERO TLV	20 octets	1 [45]
1165	Unnumbered Interface ID ERO TLV	12 octets	1 [46]
1166	IPv4 Backup ERO TLV	8 octets	1 [47]
1167	IPv6 Backup ERO TLV	20 octets	1 [48]
1168	Unnumbered Interface ID Backup ERO TLV	12 octets	1 [49]
1169	IPv6 Prefix SID	variable	Section 2.3.2
1170	IGP Prefix Attributes	variable	Section 2.3.3
1171	Source Router-ID	variable	Section 2.3.4

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

5. Manageability Considerations

This section is structured as recommended in [RFC5706].

5.1. Operational Considerations

5.1.1. Operations

Existing BGP and BGP-LS operational procedures apply. No additional operation procedures are defined in this document.

6. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the 'Security Considerations' section of [RFC4271] for a discussion of BGP security. Also refer to [RFC4272] and [RFC6952] for analysis of security issues for BGP.

7. Contributors

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

Acee Lindem
Cisco Systems
Email: acee@cisco.com

Saikat Ray
Individual
Email: raysaikat@gmail.com

8. Acknowledgements

The authors would like to thank Les Ginsberg for the review of this document.

9. References

9.1. Normative References

[I-D.ietf-isis-segment-routing-extensions]
Previdi, S., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., Decraene, B., and j. jefftant@gmail.com, "IS-IS Extensions for Segment Routing", draft-ietf-isis-segment-routing-extensions-08 (work in progress), October 2016.

[I-D.ietf-ospf-ospfv3-segment-routing-extensions]
Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPFv3 Extensions for Segment Routing", draft-ietf-ospf-ospfv3-segment-routing-extensions-07 (work in progress), October 2016.

[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-10 (work in progress), October 2016.

[I-D.previdi-isis-ipv6-prefix-sid]
Previdi, S., Ginsberg, L., and C. Filsfils, "Segment Routing IPv6 Prefix-SID", draft-previdi-isis-ipv6-prefix-sid-02 (work in progress), May 2016.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <http://www.rfc-editor.org/info/rfc4271>.

[RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", RFC 7684, DOI 10.17487/RFC7684, November 2015, <http://www.rfc-editor.org/info/rfc7684>.

[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, <http://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, <http://www.rfc-editor.org/info/rfc7794>.

9.2. Informative References

[I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", draft-ietf-spring-segment-routing-09 (work in progress), July 2016.

[RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <http://www.rfc-editor.org/info/rfc4272>.

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

[RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, DOI 10.17487/RFC6952, May 2013, <http://www.rfc-editor.org/info/rfc6952>.

9.3. URIs

[1] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-3.1

[2] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-3.2

[3] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.2.1

[4] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.2.2

[5] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.1

[6] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4

[7] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.3

[8] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.7

[9] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.8

[10] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.9

[11] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.10

[12] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.11

[13] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.12

[14] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.13

[15] http://tools.ietf.org/html/draft-previdi-isis-ipv6-prefix-sid-01

[16] http://tools.ietf.org/html/RFC7794

[17] http://tools.ietf.org/html/RFC7794

[18] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-3.2

[19] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-3.1

[20] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-7.1

[21] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-7.2

[22] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-5

[23] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-2.1

[24] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-6.1

[25] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-6.2.1

[26] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-6.2.2

[27] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-6.2.3

[28] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-6.2.4

[29] http://tools.ietf.org/html/draft-ietf-ospf-segment-routing-extensions-05#section-6.2.4

[30] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-3.2

[31] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-3.1

[32] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-7.1

[33] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-7.2

[34] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-5

[35] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-2.1

[36] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-6.1

[37] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-6.2.1

[38] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-6.2.2

[39] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-6.2.3

[40] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-6.2.4

[41] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-6.2.5

[42] http://tools.ietf.org/html/draft-ietf-ospf-ospfv3-segment-routing-extensions-05#section-6.2.6

[43] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.7

[44] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.8

[45] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.9

[46] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.10

[47] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.11

[48] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.12

[49] http://tools.ietf.org/html/draft-ietf-isis-segment-routing-extensions-05#section-2.4.13

Authors' Addresses

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

Email: sprevidi@cisco.com

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

Email: ppsenak@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

Jeff Tantsura
Individual

Email: jefftant@gmail.com