

Internet Engineering Task Force (IETF)
Request for Comments: 9352
Updates: 7370
Category: Standards Track
ISSN: 2070-1721

P. Psenak, Ed.
C. Filsfils
A. Bashandy
Cisco Systems
B. Decraene
Orange
Z. Hu
Huawei Technologies
February 2023

IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane

Abstract

The Segment Routing (SR) architecture allows a flexible definition of the end-to-end path by encoding it as a sequence of topological elements called "segments". It can be implemented over the MPLS or the IPv6 data plane. This document describes the IS-IS extensions required to support SR over the IPv6 data plane.

This document updates RFC 7370 by modifying an existing registry.

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/rfc9352>.

Copyright Notice

Copyright (c) 2023 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction
 - 1.1. Requirements Language
2. SRv6 Capabilities Sub-TLV
3. Advertising Supported Algorithms
4. Advertising Maximum SRv6 SID Depths
 - 4.1. Maximum Segments Left MSD Type
 - 4.2. Maximum End Pop MSD Type
 - 4.3. Maximum H.Encaps MSD Type

4.4.	Maximum End D MSD Type
5.	SRv6 SIDs and Reachability
6.	Advertising Anycast Property
7.	Advertising Locators and End SIDs
7.1.	SRv6 Locator TLV Format
7.2.	SRv6 End SID Sub-TLV
8.	Advertising SRv6 Adjacency SIDs
8.1.	SRv6 End.X SID Sub-TLV
8.2.	SRv6 LAN End.X SID Sub-TLV
9.	SRv6 SID Structure Sub-Sub-TLV
10.	Advertising Endpoint Behaviors
11.	IANA Considerations
11.1.	SRv6 Locator TLV
11.1.1.	SRv6 End SID Sub-TLV
11.1.2.	IS-IS Sub-TLVs for TLVs Advertising Prefix Reachability Registry
11.2.	SRv6 Capabilities Sub-TLV
11.3.	IS-IS Sub-Sub-TLVs for the SRv6 Capabilities Sub-TLV Registry
11.4.	SRv6 End.X SID and SRv6 LAN End.X SID Sub-TLVs
11.5.	MSD Types
11.6.	IS-IS Sub-Sub-TLVs for SRv6 SID Sub-TLVs Registry
11.7.	Prefix Attribute Flags Sub-TLV
11.8.	IS-IS SRv6 Capabilities Sub-TLV Flags Registry
11.9.	IS-IS SRv6 Locator TLV Flags Registry
11.10.	IS-IS SRv6 End SID Sub-TLV Flags Registry
11.11.	IS-IS SRv6 Adjacency SID Sub-TLVs Flags Registry
12.	Security Considerations
13.	References
13.1.	Normative References
13.2.	Informative References
	Acknowledgements
	Contributors
	Authors' Addresses

1. Introduction

With Segment Routing (SR) [RFC8402], a node steers a packet through an ordered list of instructions, which are called segments.

Segments are identified through Segment Identifiers (SIDs).

SR can be directly instantiated on the IPv6 data plane through the use of the Segment Routing Header (SRH) defined in [RFC8754]. SRv6 refers to this SR instantiation on the IPv6 data plane.

The network programming paradigm [RFC8986] is central to SRv6. It describes how any behavior can be bound to a SID and how any network program can be expressed as a combination of SIDs.

This document specifies IS-IS extensions that allow the IS-IS protocol to encode some of these SIDs and their behaviors.

Familiarity with the network programming paradigm [RFC8986] is necessary to understand the extensions specified in this document.

The new SRv6 Locator top-level TLV announces SRv6 Locators -- a form of summary address for the set of topology-/algorithm-specific SIDs instantiated at the node.

The SRv6 Capabilities sub-TLV announces the ability to support SRv6.

Several new sub-TLVs are defined to advertise various SRv6 Maximum SID Depths (MSDs).

The SRv6 End SID sub-TLV, the SRv6 End.X SID sub-TLV, and the SRv6

LAN End.X SID sub-TLV are used to advertise which SIDs are instantiated at a node and what Endpoint behavior is bound to each instantiated SID.

This document updates [RFC7370] by modifying an existing registry (Section 11.1.2).

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. SRv6 Capabilities Sub-TLV

A node indicates that it supports the SR Segment Endpoint Node functionality as specified in [RFC8754] by advertising a new SRv6 Capabilities sub-TLV of the Router Capability TLV [RFC7981].

The SRv6 Capabilities sub-TLV may contain optional sub-sub-TLVs. No sub-sub-TLVs are currently defined.

The SRv6 Capabilities sub-TLV has the following format:

0										1										2										3																															
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																														
Type																																Length										Flags																			
optional sub-sub-TLVs...																																																													

Type: 25. Single octet, as defined in Section 9 of [ISO10589].

Length: Single octet, as defined in Section 9 of [ISO10589]. The length value is 2 + length of sub-sub-TLVs.

Flags: 2 octets. The following flags are defined:

0										1									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5				
O										Reserved									

where:

O-flag: If set, the router supports use of the O-bit in the SRH, as defined in [RFC9259].

The remaining bits, including bit 0, are reserved for future use. They MUST be set to zero on transmission and MUST be ignored on receipt.

3. Advertising Supported Algorithms

An SRv6-capable router indicates one or more supported algorithms by advertising the Segment Routing Algorithm sub-TLV, as defined in [RFC8667].

4. Advertising Maximum SRv6 SID Depths

[RFC8491] defines the means to advertise node-/link-specific values for MSDs of various types. Node MSDs are advertised in a sub-TLV of the Router Capability TLV [RFC7981]. Link MSDs are advertised in a sub-TLV of TLVs 22, 23, 25, 141, 222, and 223.

This document defines the relevant SRv6 MSDs and requests MSD type assignments in the "IGP MSD-Types" registry created by [RFC8491].

4.1. Maximum Segments Left MSD Type

The Maximum Segments Left MSD Type signals the maximum value of the "Segments Left" field [RFC8754] in the SRH of a received packet before applying the Endpoint behavior associated with a SID.

SRH Max Segments Left Type: 41

If no value is advertised, the supported value is 0.

4.2. Maximum End Pop MSD Type

The Maximum End Pop MSD Type signals the maximum number of SIDs in the SRH to which the router can apply "Penultimate Segment Pop (PSP) of the SRH" or "Ultimate Segment Pop (USP) of the SRH" behavior, as defined in "Flavors" (Section 4.16 of [RFC8986]).

SRH Max End Pop Type: 42

If the advertised value is zero or no value is advertised, then the router cannot apply PSP or USP flavors.

4.3. Maximum H.Encaps MSD Type

The Maximum H.Encaps MSD Type signals the maximum number of SIDs that can be added to the segment list of an SRH as part of the "H.Encaps" behavior, as defined in [RFC8986].

SRH Max H.encaps Type: 44

If the advertised value is zero or no value is advertised, then the headend can apply an SR Policy that only contains one segment without inserting any SRH header.

A non-zero SRH Max H.encaps MSD indicates that the headend can insert an SRH up to the advertised number of SIDs.

4.4. Maximum End D MSD Type

The Maximum End D MSD Type specifies the maximum number of SIDs present in an SRH when performing decapsulation. As specified in [RFC8986], the permitted SID types include, but are not limited to, End.DX6, End.DT4, End.DT46, End with USD, and End.X with USD.

SRH Max End D Type: 45

If the advertised value is zero or no value is advertised, then the router cannot apply any behavior that results in decapsulation and forwarding of the inner packet if the outer IPv6 header contains an SRH.

5. SRv6 SIDs and Reachability

As discussed in [RFC8986], an SRv6 Segment Identifier (SID) is 128 bits and consists of locator, function, and argument parts.

A node is provisioned with topology-/algorithm-specific locators for each of the topology/algorithm pairs supported by that node. Each locator is a covering prefix for all SIDs provisioned on that node that have the matching topology/algorithm.

Locators MUST be advertised in the SRv6 Locator TLV (see

Section 7.1). Forwarding entries for the locators advertised in the SRv6 Locator TLV MUST be installed in the forwarding plane of receiving SRv6-capable routers when the associated topology/algorithm is supported by the receiving node. The processing of the prefix advertised in the SRv6 Locator TLV, the calculation of its reachability, and the installation in the forwarding plane follows the process defined for the Prefix Reachability TLV 236 [RFC5308] or TLV 237 [RFC5120].

Locators associated with algorithms 0 and 1 (for all supported topologies) SHOULD also be advertised in a Prefix Reachability TLV (236 or 237) so that legacy routers (i.e., routers that do not support SRv6) will install a forwarding entry for algorithms 0 and 1 SRv6 traffic.

In cases where the same prefix with the same prefix length, Multi-Topology Identifier (MTID), and algorithm is received in both a Prefix Reachability TLV and an SRv6 Locator TLV, the Prefix Reachability advertisement MUST be preferred when installing entries in the forwarding plane. This is to prevent inconsistent forwarding entries between SRv6-capable and SRv6-incapable routers. Such preference of Prefix Reachability advertisement does not have any impact on the rest of the data advertised in the SRv6 Locator TLV.

Locators associated with Flexible Algorithms (see Section 4 of [RFC9350]) SHOULD NOT be advertised in Prefix Reachability TLVs (236 or 237). Advertising the Flexible Algorithm locator in a regular Prefix Reachability TLV (236 or 237) would make the forwarding for it follow the algorithm 0 path.

SRv6 SIDs are advertised as sub-TLVs in the SRv6 Locator TLV, except for SRv6 SIDs that are associated with a specific neighbor/link and are therefore advertised as sub-TLVs in TLVs 22, 23, 25, 141, 222, and 223.

SRv6 SIDs received from other nodes are not directly routable and MUST NOT be installed in the forwarding plane. Reachability to SRv6 SIDs depends upon the existence of a covering locator.

Adherence to the rules defined in this section will ensure that SRv6 SIDs associated with a supported topology/algorithm pair will be forwarded correctly, while SRv6 SIDs associated with an unsupported topology/algorithm pair will be dropped. NOTE: The drop behavior depends on the absence of a default/summary route covering a given locator.

In order for forwarding to work correctly, the locator associated with SRv6 SID advertisements must be the longest match prefix installed in the forwarding plane for those SIDs. In order to ensure correct forwarding, network operators should take steps to make sure that this requirement is not compromised. For example, the following situations should be avoided:

- * Another locator associated with a different topology/algorithm is the longest match.
- * Another prefix advertisement (i.e., from TLV 236 or 237) is the longest match.

6. Advertising Anycast Property

Both prefixes and SRv6 Locators may be configured as anycast; as such, the same value can be advertised by multiple routers. It is useful for other routers to know that the advertisement is for an anycast identifier.

A new flag in the Prefix Attribute Flags sub-TLV [RFC7794] is defined to advertise the anycast property:

Bit #: 4
Name: Anycast Flag (A-flag)

When the prefix/SRv6 Locator is configured as anycast, the A-flag SHOULD be set. Otherwise, this flag MUST be clear.

The A-flag MUST be preserved when the advertisement is leaked between levels.

The A-flag and the N-flag MUST NOT both be set. If both the N-flag and the A-flag are set in the prefix/SRv6 Locator advertisement, the receiving routers MUST ignore the N-flag.

The same prefix/SRv6 Locator can be advertised by multiple routers. If at least one of them sets the A-flag in its advertisement, the prefix/SRv6 Locator SHOULD be considered as anycast.

A prefix/SRv6 Locator that is advertised by a single node and without an A-flag MUST be considered node specific.

All the nodes advertising the same anycast locator MUST instantiate the exact same set of SIDs under that anycast locator. Failure to do so may result in traffic being dropped or misrouted.

The Prefix Attribute Flags sub-TLV can be carried in the SRv6 Locator TLV as well as the Prefix Reachability TLVs. When a router originates both the Prefix Reachability TLV and the SRv6 Locator TLV for a given prefix, it SHOULD advertise the Prefix Attribute Flags sub-TLV, if used, in both TLVs and use the same flags. However, unlike TLVs 236 [RFC5308] and 237 [RFC5120], the X-flag in the Prefix Attributes Flags sub-TLV is valid when sent in the SRv6 Locator TLV. When included in the Locator TLV, the state of the X-flag in the Prefix Attributes Flags sub-TLV MUST match the setting of the embedded "X-bit" in any advertisement for the same prefix in TLVs 236 [RFC5308] and 237 [RFC5120]. In case of any inconsistency between the Prefix Attribute Flags advertised in the Locator TLV and in the Prefix Reachability TLV, the ones advertised in the Prefix Reachability TLV MUST be preferred.

7. Advertising Locators and End SIDs

The SRv6 Locator TLV is introduced to advertise SRv6 Locators and End SIDs associated with each locator.

This new TLV shares the sub-TLV space defined for TLVs 135, 235, 236, and 237.

7.1. SRv6 Locator TLV Format

The SRv6 Locator TLV has the following format:

0										1										2										3															
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9						
Type										Length										R				R				R				R				MTID									
Locator Entries . . .																																													

Type: 27. Single octet, as defined in Section 9 of [ISO10589].

Length: Single octet, as defined in Section 9 of [ISO10589]. The length value is variable.

R Bits: Reserved for future use. They MUST be set to zero on transmission and MUST be ignored on receipt.

MTID: Multi-Topology Identifier, as defined in [RFC5120]. Note that the value 0 is legal.

The SRv6 Locator TLV is followed by one or more locator entries of the form:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Metric                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Flags      | Algorithm  | Loc Size  |                                     |
+-----+-----+-----+-----+-----+-----+-----+

//                               Locator (continued, variable)                               //
+-----+-----+-----+-----+-----+-----+-----+-----+
| Sub-TLV-len | Sub-TLVs (variable) . . . |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

Metric: 4 octets, as described in Section 4 of [RFC5305].

Flags: 1 octet. The following flags are defined:

```

      0
      0 1 2 3 4 5 6 7
+-----+
|D|   Reserved   |
+-----+
```

D-flag: "up/down bit" as described in Section 4.1 of [RFC5305].

The remaining bits are reserved for future use. They MUST be set to zero on transmission and MUST be ignored on receipt.

Algorithm: 1 octet, as defined in the "IGP Algorithm Types" registry [RFC8665].

Loc-Size: 1 octet. Number of bits in the SRv6 Locator field, which MUST be from the range (1-128). The entire TLV MUST be ignored if the Loc-Size is outside this range.

Locator: 1-16 octets. This field encodes the advertised SRv6 Locator. The SRv6 Locator is encoded in the minimal number of octets for the given number of bits. Trailing bits MUST be set to zero and ignored when received.

Sub-TLV-length: 1 octet. Number of octets used by sub-TLVs.

Optional Sub-TLVs: Supported sub-TLVs are specified in Section 11.1.2. Any sub-TLV that is not allowed in the SRv6 Locator TLV MUST be ignored.

The Prefix Attribute Flags sub-TLV [RFC7794] SHOULD be included in the Locator TLV.

The Prefix Attribute Flags sub-TLV MUST be included in the Locator TLV when it is leaked upwards in the hierarchy or originated as a result of the redistribution from another protocol or another IS-IS instance. If the Prefix Attribute Flags sub-TLV is not included in these cases, receivers will be unable to determine the correct source of the advertisement. The receivers will be unable to detect the violation.


7.2. SRv6 End SID Sub-TLV

The SRv6 End SID sub-TLV is introduced to advertise SRv6 SIDs with Endpoint behaviors that do not require a particular neighbor in order to be correctly applied. SRv6 SIDs associated with a neighbor are advertised using the sub-TLVs defined in Section 8.

Supported behavior values, together with parent TLVs in which they are advertised, are specified in Section 10 of this document. Please note that not all behaviors defined in [RFC8986] are defined in this document, e.g., End.T is not.

This new sub-TLV is advertised in the SRv6 Locator TLV defined in the previous section. SRv6 End SIDs inherit the topology/algorithm from the parent locator.

The SRv6 End SID sub-TLV has the following format:



Type: 5. Single octet, as defined in Section 9 of [ISO10589].

Length: Single octet, as defined in Section 9 of [ISO10589]. The length value is variable.

Flags: 1 octet. No flags are currently defined. All bits are reserved for future use. They MUST be set to zero on transmission and MUST be ignored on receipt.

Endpoint Behavior: 2 octets, as defined in [RFC8986]. Supported behavior values for this sub-TLV are defined in Section 10 of this document. Unsupported or unrecognized behavior values are ignored by the receiver.

SID: 16 octets. This field encodes the advertised SRv6 SID.

Sub-sub-TLV-length: 1 octet. Number of octets used by sub-sub-TLVs.

Optional Sub-sub-TLVs: Supported sub-sub-TLVs are specified in Section 11.6. Any sub-sub-TLV that is not allowed in the SRv6 End SID sub-TLV MUST be ignored.

The SRv6 End SID MUST be allocated from its associated locator. SRv6 End SIDs that are not allocated from the associated locator MUST be ignored.

Multiple SRv6 End SIDs MAY be associated with the same locator. In

cases where the number of SRv6 End SID sub-TLVs exceeds the capacity of a single TLV, multiple Locator TLVs for the same locator MAY be advertised. For a given MTID/Locator, the algorithm MUST be the same in all TLVs. If this restriction is not met, all TLVs for that MTID/Locator MUST be ignored.

8. Advertising SRv6 Adjacency SIDs

Certain SRv6 Endpoint behaviors [RFC8986] are associated with a particular adjacency.

This document defines two new sub-TLVs of TLVs 22, 23, 25, 141, 222, and 223 -- namely "SRv6 End.X SID sub-TLVs" and "SRv6 LAN End.X SID sub-TLVs".

IS-IS neighbor advertisements are topology specific but not algorithm specific. SIDs advertised in SRv6 End.X SID and SRv6 LAN End.X SID sub-TLVs therefore inherit the topology from the associated neighbor advertisement, but the algorithm is specified in the individual SID.

All SIDs advertised in SRv6 End.X SID and SRv6 LAN End.X SID sub-TLVs MUST be a subnet of a Locator with matching topology and algorithm that are advertised by the same node in an SRv6 Locator TLV. SIDs that do not meet this requirement MUST be ignored. This ensures that the node advertising these SIDs is also advertising its corresponding Locator with the algorithm that will be used for computing paths destined to the SID.

8.1. SRv6 End.X SID Sub-TLV

This sub-TLV is used to advertise an SRv6 SID associated with a point-to-point adjacency. Multiple SRv6 End.X SID sub-TLVs MAY be associated with the same adjacency.

The SRv6 End.X SID sub-TLV has the following format:

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Type      |      Length      |   Flags   |   Algorithm   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Weight    |      Endpoint Behavior      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| SID (128 bits) . . . |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| SID (cont . . .) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| SID (cont . . .) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| SID (cont . . .) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Sub-sub-tilv-len|      Sub-sub-TLVs (variable) . . . |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type: 43. Single octet, as defined in Section 9 of [IS010589].

Length: Single octet, as defined in Section 9 of [IS010589]. The length value is variable.

Flags: 1 octet.

```

0 1 2 3 4 5 6 7
+---+---+---+---+---+---+
|B|S|P|Reserved |

```

+-----+

where:

B-Flag: Backup flag. If set, the SID is eligible for protection, e.g., using IP Fast Reroute (IPFRR) [RFC5286], as described in [RFC8355].

S-Flag: Set flag. When set, the S-flag indicates that the SID refers to a set of adjacencies (and therefore MAY be assigned to other adjacencies as well).

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

Reserved bits: Reserved bits MUST be zero when originated and MUST be ignored when received.

Algorithm: 1 octet, as defined in the "IGP Algorithm Types" registry [RFC8665].

Weight: 1 octet. The value represents the weight of the SID for the purpose of load balancing. The use of the weight is defined in [RFC8402].

Endpoint Behavior: 2 octets, as defined in [RFC8986]. Supported behavior values for this sub-TLV are defined in Section 10 of this document. Unsupported or unrecognized behavior values are ignored by the receiver.

SID: 16 octets. This field encodes the advertised SRv6 SID.

Sub-sub-TLV-length: 1 octet. Number of octets used by sub-sub-TLVs.

Optional Sub-sub-TLVs: Supported sub-sub-TLVs are specified in Section 11.6. Any sub-sub-TLV that is not allowed in SRv6 End.X SID sub-TLV MUST be ignored.

Note that multiple TLVs for the same neighbor may be required in order to advertise all the SRv6 SIDs associated with that neighbor.

8.2. SRv6 LAN End.X SID Sub-TLV

This sub-TLV is used to advertise an SRv6 SID associated with a LAN adjacency. Since the parent TLV is advertising an adjacency to the Designated Intermediate System (DIS) for the LAN, it is necessary to include the System-ID of the physical neighbor on the LAN with which the SRv6 SID is associated. Given that many neighbors may exist on a given LAN, multiple SRv6 LAN END.X SID sub-TLVs may be associated with the same LAN. Note that multiple TLVs for the same DIS neighbor may be required in order to advertise all the SRv6 SIDs associated with that neighbor.

The SRv6 LAN End.X SID sub-TLV has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Type										Length																													
Neighbor System-ID (ID length octets)																																							
Flags										Algorithm										Weight																			

```

+-----+
|           Endpoint Behavior           |
+-----+
| SID (128 bits) . . .                  |
+-----+
| SID (cont . . .)                      |
+-----+
| SID (cont . . .)                      |
+-----+
| SID (cont . . .)                      |
+-----+
|Sub-sub-TLV-len|      Sub-sub-TLVs (variable) . . .
+-----+

```

Type: 44. Single octet, as defined in Section 9 of [ISO10589].

Length: Single octet, as defined in Section 9 of [ISO10589]. The length value is variable.

Neighbor System-ID: IS-IS System-ID of length "ID Length", as defined in [ISO10589].

Flags: 1 octet.

```

    0 1 2 3 4 5 6 7
+-----+
|B|S|P|Reserved |
+-----+

```

The B-, S-, and P-flags are as described in Section 8.1. Reserved bits MUST be zero when originated and MUST be ignored when received.

Algorithm: 1 octet, as defined in the "IGP Algorithm Types" registry [RFC8665].

Weight: 1 octet. The value represents the weight of the SID for the purpose of load balancing. The use of the weight is defined in [RFC8402].

Endpoint Behavior: 2 octets, as defined in [RFC8986]. Supported behavior values for this sub-TLV are defined in Section 10 of this document. Unsupported or unrecognized behavior values are ignored by the receiver.

SID: 16 octets. This field encodes the advertised SRv6 SID.

Sub-sub-TLV-length: 1 octet. Number of octets used by sub-sub-TLVs.

Optional Sub-sub-TLVs: Supported sub-sub-TLVs are specified in Section 11.6. Any sub-sub-TLV that is not allowed in SRv6 LAN End.X SID sub-TLV MUST be ignored.

Note that multiple TLVs for the same neighbor, on the same LAN, may be required in order to advertise all the SRv6 SIDs associated with that neighbor.

9. SRv6 SID Structure Sub-Sub-TLV

The SRv6 SID Structure sub-sub-TLV is an optional sub-sub-TLV of:

- * SRv6 End SID sub-TLV (Section 7.2)
- * SRv6 End.X SID sub-TLV (Section 8.1)

* SRv6 LAN End.X SID sub-TLV (Section 8.2)

The SRv6 SID Structure sub-sub-TLV is used to advertise the structure of the SRv6 SID, as defined in [RFC8986]. It has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
Type										Length																													
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
LB Length										LN Length										Fun. Length										Arg. Length									
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							

where:

Type: 1. Single octet, as defined in Section 9 of [ISO10589].

Length: Single octet, as defined in Section 9 of [ISO10589]. The length value is 4 octets.

LB Length: 1 octet. SRv6 SID Locator Block length in bits.

LN Length: 1 octet. SRv6 SID Locator Node length in bits.

Fun. Length: 1 octet. SRv6 SID Function length in bits.

Arg. Length: 1 octet. SRv6 SID Arguments length in bits.

The IS-IS SRv6 SID Structure sub-sub-TLV MUST NOT appear more than once in its parent sub-TLV. If it appears more than once in its parent sub-TLV, the parent sub-TLV MUST be ignored by the receiver.

The sum of all four sizes advertised in the IS-IS SRv6 SID Structure sub-sub-TLV MUST be less than or equal to 128 bits. If the sum of all four sizes advertised in the IS-IS SRv6 SID Structure sub-sub-TLV is larger than 128 bits, the parent sub-TLV MUST be ignored by the receiver.

The SRv6 SID Structure sub-sub-TLV is intended for informational use by the control and management planes. It MUST NOT be used at a transit node (as defined in [RFC8754]) for forwarding packets. As an example, this information could be used for the following:

- * validation of SRv6 SIDs being instantiated in the network and advertised via IS-IS. These can be learned by controllers via Border Gateway Protocol - Link State (BGP-LS) and then be monitored for conformance to the SRv6 SID allocation scheme chosen by the operator, as described in Section 3.2 of [RFC8986].
- * verification and automation for securing the SRv6 domain by provisioning filtering rules at SR domain boundaries, as described in Section 5 of [RFC8754].

The details of these potential applications are outside the scope of this document.

10. Advertising Endpoint Behaviors

Endpoint behaviors are defined in [RFC8986]. The codepoints for the Endpoint behaviors are defined in the "SRv6 Endpoint Behaviors" registry defined in [RFC8986]. If a behavior is advertised, it MUST only be advertised in the TLV(s) marked with "Y" in the table below and MUST NOT be advertised in the TLV(s) marked with "N" in the table below.

Endpoint Behavior	Endpoint Behavior Codepoint	End SID	End.X SID	Lan End.X SID
End (PSP, USP, USD)	1-4, 28-31	Y	N	N
End.X (PSP, USP, USD)	5-8, 32-35	N	Y	Y
End.DX6	16	N	Y	Y
End.DX4	17	N	Y	Y
End.DT6	18	Y	N	N
End.DT4	19	Y	N	N
End.DT46	20	Y	N	N

Table 1: Endpoint Behaviors

11. IANA Considerations

This document requests allocation for the following TLVs, sub-TLVs, and sub-sub-TLVs by updating the existing registries and defining new registries under the "IS-IS TLV Codepoints" grouping.

11.1. SRv6 Locator TLV

The SRv6 Locator TLV shares sub-TLV space with TLVs advertising prefix reachability. IANA has updated the "IS-IS Sub-TLVs for TLVs Advertising Prefix Reachability" registry initially defined in [RFC7370] by adding this document as a reference and updating the description of that registry to include the SRv6 Locator TLV (27).

This document makes the following registration in the "IS-IS Top-Level TLV Codepoints" registry:

Value	Name	IIH	LSP	SNP	Purge
27	SRv6 Locator	n	y	n	n

Table 2: IS-IS Top-Level TLV Codepoints Registry

11.1.1. SRv6 End SID Sub-TLV

This document makes the following registration:

Type	Description	27	135	235	236	237	Reference
5	SRv6 End SID	y	n	n	n	n	RFC 9352, Section 7.2

Table 3: IS-IS Sub-TLVs for TLVs Advertising Prefix Reachability Registry

11.1.2. IS-IS Sub-TLVs for TLVs Advertising Prefix Reachability Registry

IANA has updated the "IS-IS Sub-TLVs for TLVs Advertising Prefix

Reachability" registry to include a column for the SRv6 Locator TLV (27) as shown below:

Type	Description	27	135	235	236	237
1	32-bit Administrative Tag Sub-TLV	y	y	y	y	y
2	64-bit Administrative Tag Sub-TLV	y	y	y	y	y
3	Prefix Segment Identifier	n	y	y	y	y
4	Prefix Attribute Flags	y	y	y	y	y
6	Flexible Algorithm Prefix Metric (FAPM)	n	y	y	y	y
11	IPv4 Source Router ID	y	y	y	y	y
12	IPv6 Source Router ID	y	y	y	y	y
32	BIER Info	n	y	y	y	y

Table 4: IS-IS Sub-TLVs for TLVs Advertising Prefix Reachability Registry

11.2. SRv6 Capabilities Sub-TLV

This document makes the following registration in the "IS-IS Sub-TLVs for IS-IS Router CAPABILITY TLV" registry:

Value	Description	Reference
25	SRv6 Capabilities	RFC 9352, Section 2

Table 5: IS-IS Sub-TLVs for IS-IS Router CAPABILITY TLV Registry

11.3. IS-IS Sub-Sub-TLVs for the SRv6 Capabilities Sub-TLV Registry

IANA has created the "IS-IS Sub-Sub-TLVs for SRv6 Capabilities Sub-TLV" registry under the "IS-IS TLV Codepoints" grouping for the assignment of sub-TLV types for the SRv6 Capabilities sub-TLV specified in this document (Section 2). This registry defines sub-sub-TLVs for the SRv6 Capabilities sub-TLV (25) advertised in the IS-IS Router CAPABILITY TLV (242).

The registration procedure is "Expert Review", as defined in [RFC8126]. Guidance for the designated experts is provided in [RFC7370]. No sub-sub-TLVs are defined by this document, except for the reserved type 0.

Value	Description	Reference
0	Reserved	RFC 9532
1-255	Unassigned	

Table 6: IS-IS Sub-Sub-TLVs for
SRv6 Capabilities Sub-TLV
Registry

11.4. SRv6 End.X SID and SRv6 LAN End.X SID Sub-TLVs

This document makes the following registrations in the "IS-IS Sub-TLVs for TLVs Advertising Neighbor Information" registry:

Type	Description	22	23	25	141	222	223	Reference
43	SRv6 End.X SID	y	y	y	y	y	y	RFC 9352, Section 8.1
44	SRv6 LAN End.X SID	y	y	y	y	y	y	RFC 9352, Section 8.2

Table 7: IS-IS Sub-TLVs for TLVs Advertising Neighbor Information
Registry

11.5. MSD Types

This document makes the following registrations in the "IGP MSD-Types" registry:

Value	Name	Reference
41	SRH Max SL	RFC 9352
42	SRH Max End Pop	RFC 9352
44	SRH Max H.encaps	RFC 9352
45	SRH Max End D	RFC 9352

Table 8: IGP MSD-Types

11.6. IS-IS Sub-Sub-TLVs for SRv6 SID Sub-TLVs Registry

IANA has created the "IS-IS Sub-Sub-TLVs for SRv6 SID Sub-TLVs" registry under the "IS-IS TLV Codepoints" grouping to assign sub-TLV types for the SID sub-TLVs specified in this document (Sections 7.2, 8.1, and 8.2).

This registry defines sub-sub-TLVs for SRv6 SID sub-TLVs. This includes the following sub-TLVs:

- * SRv6 End SID (5) (Advertised in SRv6 Locator TLV (27))
- * SRv6 End.X SID (43) (Advertised in TLVs advertising neighbor information)
- * SRv6 LAN End.X SID (44) (Advertised in TLVs advertising neighbor information)

The registration procedure is "Expert Review", as defined in [RFC8126]. Guidance for the designated experts is provided in [RFC7370]. The following assignments are made by this document:

Type	Description	5	43	44	Reference
0	Reserved				RFC 9352
1	SRv6 SID Structure	y	y	y	RFC 9352
2-255	Unassigned				

Table 9: IS-IS Sub-Sub-TLVs for SRv6 SID Sub-TLVs Registry

11.7. Prefix Attribute Flags Sub-TLV

This document adds a new bit in the "IS-IS Bit Values for Prefix Attribute Flags Sub-TLV" registry:

Bit #	Name	Reference
4	Anycast Flag (A-flag)	RFC 9352, Section 6

Table 10: IS-IS Bit Values for Prefix Attribute Flags Sub-TLV Registry

11.8. IS-IS SRv6 Capabilities Sub-TLV Flags Registry

IANA has created the "IS-IS SRv6 Capabilities Sub-TLV Flags" registry under the "IS-IS TLV Codepoints" grouping to assign bits 0 to 15 in the Flags field of the IS-IS SRv6 Capabilities sub-TLV specified in this document (Section 2). This registry defines bit values advertised in the Flags field of the SRv6 Capabilities sub-TLV (25). This sub-TLV is advertised in the IS-IS Router CAPABILITY TLV (242).

The registration procedure is "Expert Review", as defined in [RFC8126]. Guidance for the designated experts is provided in [RFC7370]. The following assignments are made by this document:

Type	Description	Reference
0	Unassigned	
1	O-flag	RFC 9352, Section 2
2-15	Unassigned	

Table 11: IS-IS SRv6 Capabilities Sub-TLV Flags Registry

11.9. IS-IS SRv6 Locator TLV Flags Registry

IANA has created the "IS-IS SRv6 Locator TLV Flags" registry under the "IS-IS TLV Codepoints" grouping to assign bits 0 to 7 in the Flags field of the SRv6 Locator TLV specified in this document (Section 7.1). This registry defines bit values advertised in the Flags field of the SRv6 Locator TLV (27).

The registration procedure is "Expert Review", as defined in [RFC8126]. Guidance for the designated experts is provided in [RFC7370]. The following assignments are made by this document:

=====

Value	Description	Reference
0	D-flag	RFC 9352, Section 7.1
1-7	Unassigned	

Table 12: IS-IS SRv6 Locator TLV Flags Registry

11.10. IS-IS SRv6 End SID Sub-TLV Flags Registry

IANA has created the "IS-IS SRv6 End SID Sub-TLV Flags" registry under the "IS-IS TLV Codepoints" grouping to assign bits 0 to 7 in the Flags field of the IS-IS SRv6 End SID sub-TLV specified in this document (Section 7.2). This registry defines bit values advertised in the Flags field of the SRv6 End SID sub-TLV (5), which is advertised in the SRv6 Locator TLV (27).

The registration procedure is "Expert Review", as defined in [RFC8126]. Guidance for the designated experts is provided in [RFC7370]. No assignments are made by this document.

Value	Description
0-7	Unassigned

Table 13: IS-IS SRv6
End SID Sub-TLV Flags
Registry

11.11. IS-IS SRv6 Adjacency SID Sub-TLVs Flags Registry

IANA has created the "IS-IS SRv6 Adjacency SID Sub-TLVs Flags" registry under the "IS-IS TLV Codepoints" grouping to assign bits 0 to 7 in the Flags field of the IS-IS SRv6 End.X SID and LAN End.X SID sub-TLVs (Sections 8.1 and 8.2).

This registry defines bit values advertised in the Flags field of SRv6 SID sub-TLVs associated with adjacencies. These sub-TLVs are advertised in TLVs advertising neighbor information. The list of sub-TLVs includes:

- * SRv6 End.X SID (43)
- * SRv6 LAN End.X SID (44)

The registration procedure is "Expert Review", as defined in [RFC8126]. Guidance for the designated experts is provided in [RFC7370]. The following assignments are made by this document:

Value	Description	Reference
0	B-flag	RFC 9352, Section 8.1
1	S-flag	RFC 9352, Section 8.1
2	P-flag	RFC 9352, Section 8.1
3-7	Unassigned	

Table 14: IS-IS SRv6 Adjacency SID Sub-TLVs
Flags Registry

12. Security Considerations

Security concerns for IS-IS are addressed in [ISO10589], [RFC5304], and [RFC5310]. While IS-IS is deployed under a single administrative domain, there can be deployments where potential attackers have access to one or more networks in the IS-IS routing domain. In these deployments, the stronger authentication mechanisms defined in the aforementioned documents SHOULD be used.

This document describes the IS-IS extensions required to support SR over an IPv6 data plane. The security considerations for SR are discussed in [RFC8402]. [RFC8986] defines the SRv6 Network Programming concept and specifies the main SR behaviors to enable the creation of interoperable overlays; the security considerations from that document apply too.

The advertisement for an incorrect MSD value may have negative consequences; see [RFC8491] for additional considerations.

Security concerns associated with the setting of the O-flag are described in [RFC9259].

Security concerns associated with the usage of Flexible Algorithms are described in [RFC9350]).

13. References

13.1. Normative References

- [ISO10589] ISO, "Information technology - Telecommunications and information exchange between systems - Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", Second Edition, ISO/IEC 10589:2002, November 2002.
- [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>>.
- [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>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <<https://www.rfc-editor.org/info/rfc5308>>.
- [RFC7370] Ginsberg, L., "Updates to the IS-IS TLV Codepoints Registry", RFC 7370, DOI 10.17487/RFC7370, September 2014, <<https://www.rfc-editor.org/info/rfc7370>>.
- [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>>.
- [RFC7981] Ginsberg, L., Previdi, S., and M. Chen, "IS-IS Extensions

- for Advertising Router Information", RFC 7981,
DOI 10.17487/RFC7981, October 2016,
<<https://www.rfc-editor.org/info/rfc7981>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017,
<<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.
- [RFC8402] Filts, 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>>.
- [RFC8491] Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg, "Signaling Maximum SID Depth (MSD) Using IS-IS", RFC 8491, DOI 10.17487/RFC8491, November 2018,
<<https://www.rfc-editor.org/info/rfc8491>>.
- [RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filts, 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>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filts, 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>>.
- [RFC8754] Filts, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020,
<<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC8986] Filts, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021,
<<https://www.rfc-editor.org/info/rfc8986>>.
- [RFC9259] Ali, Z., Filts, C., Matsushima, S., Voyer, D., and M. Chen, "Operations, Administration, and Maintenance (OAM) in Segment Routing over IPv6 (SRv6)", RFC 9259, DOI 10.17487/RFC9259, June 2022,
<<https://www.rfc-editor.org/info/rfc9259>>.
- [RFC9350] Psenak, P., Ed., Hegde, S., Filts, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023,
<<https://www.rfc-editor.org/rfc/rfc9350>>.

13.2. Informative References

- [RFC5286] Atlas, A., Ed. and A. Zinin, Ed., "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, DOI 10.17487/RFC5286, September 2008,
<<https://www.rfc-editor.org/info/rfc5286>>.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October

2008, <<https://www.rfc-editor.org/info/rfc5304>>.

[RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <<https://www.rfc-editor.org/info/rfc5310>>.

[RFC8355] Filsfils, C., Ed., Previdi, S., Ed., Decraene, B., and R. Shakir, "Resiliency Use Cases in Source Packet Routing in Networking (SPRING) Networks", RFC 8355, DOI 10.17487/RFC8355, March 2018, <<https://www.rfc-editor.org/info/rfc8355>>.

Acknowledgements

Thanks to Christian Hopps for his review comments and shepherd work.

Thanks to Alvaro Retana and John Scudder for AD review and comments.

Contributors

The following people gave a substantial contribution to the content of this document and should be considered coauthors:

Stefano Previdi
Huawei Technologies
Email: stefano@previdi.net

Paul Wells
Cisco Systems
Saint Paul, Minnesota
United States of America
Email: pauwells@cisco.com

Daniel Voyer
Email: daniel.voyer@bell.ca

Satoru Matsushima
Email: satoru.matsushima@g.softbank.co.jp

Bart Peirens
Email: bart.peirens@proximus.com

Hani Elmalky
Email: hani.elmalky@ericsson.com

Prem Jonnalagadda
Email: prem@barefootnetworks.com

Milad Sharif
Email: msharif@barefootnetworks.com

Robert Hanzl
Cisco Systems
Millenium Plaza Building, V Celnici 10, Prague 1
Prague
Czech Republic
Email: rhanzl@cisco.com

Ketan Talaulikar
Cisco Systems, Inc.
Email: ketant@cisco.com

Authors' Addresses

Peter Psenak (editor)
Cisco Systems
Pribinova Street 10
81109 Bratislava
Slovakia
Email: ppsenak@cisco.com

Clarence Filsfils
Cisco Systems
Brussels
Belgium
Email: cfilsfil@cisco.com

Ahmed Bashandy
Cisco Systems
Milpitas,
United States of America
Email: bashandy@cisco.com

Bruno Decraene
Orange
Chatillon
France
Email: bruno.decraene@orange.com

Zhibo Hu
Huawei Technologies
Email: huzhibo@huawei.com