

Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 18, 2020

P. Psenak, Ed.
C. Filsfils
Cisco Systems
A. Bashandy
Arrcus
B. Decraene
Orange
Z. Hu
Huawei Technologies
January 15, 2020

IS-IS Extension to Support Segment Routing over IPv6 Dataplane
draft-ietf-lsr-isis-srv6-extensions-04.txt

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". Segment routing architecture can be implemented over an MPLS data plane as well as an IPv6 data plane. This draft describes the IS-IS extensions required to support Segment Routing over an IPv6 data plane.

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 July 18, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	SRv6 Capabilities sub-TLV	4
3.	Advertising Supported Algorithms	4
4.	Advertising Maximum SRv6 SID Depths	5
4.1.	Maximum Segments Left MSD Type	5
4.2.	Maximum End Pop MSD Type	5
4.3.	Maximum H.Encaps MSD Type	5
4.4.	Maximum End D MSD Type	6
5.	SRv6 SIDs and Reachability	6
6.	Advertising Anycast Property	7
7.	Advertising Locators and End SIDs	8
7.1.	SRv6 Locator TLV Format	8
7.2.	SRv6 End SID sub-TLV	10
8.	Advertising SRv6 Adjacency SIDs	12
8.1.	SRv6 End.X SID sub-TLV	12
8.2.	SRv6 LAN End.X SID sub-TLV	14
9.	SRv6 SID Structure Sub-Sub-TLV	15
10.	Advertising Endpoint Behaviors	16
11.	Implementation Status	17
11.1.	Cisco	17
11.2.	Huawei	17
11.3.	Juniper	18
11.4.	Interoperability Testing.	18
12.	IANA Considerations	18
12.1.	SRv6 Locator TLV	18
12.1.1.	SRv6 End SID sub-TLV	18
12.1.2.	Revised sub-TLV table	19
12.2.	SRv6 Capabilities sub-TLV	19
12.3.	SRv6 End.X SID and SRv6 LAN End.X SID sub-TLVs	19
12.4.	MSD Types	20
12.5.	Sub-Sub-TLVs for SID Sub-TLVs	20

12.6.	Prefix Attribute Flags Sub-TLV	20
13.	Security Considerations	21
14.	Contributors	21
15.	References	22
15.1.	Normative References	22
15.2.	Informative References	24
15.3.	URIs	25
Authors'	Addresses	25

1. Introduction

With Segment Routing (SR) [[I-D.ietf-spring-segment-routing](#)], a node steers a packet through an ordered list of instructions, called segments.

Segments are identified through Segment Identifiers (SIDs).

Segment Routing can be directly instantiated on the IPv6 data plane through the use of the Segment Routing Header defined in [[I-D.ietf-6man-segment-routing-header](#)]. SRv6 refers to this SR instantiation on the IPv6 dataplane.

The network programming paradigm [[I-D.ietf-spring-srv6-network-programming](#)] 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 SID's.

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 [[I-D.ietf-spring-srv6-network-programming](#)] is necessary to understand the extensions specified in this document.

This document defines one new top level IS-IS TLV and several new IS-IS sub-TLVs.

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

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

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 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.

2. SRv6 Capabilities sub-TLV

A node indicates that it has support for SRv6 by advertising a new SRv6 Capabilities sub-TLV of the router capabilities 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: Suggested value 25, to be assigned by IANA

Length: 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| |0|                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

0-flag: If set, the router supports use of the 0-bit in the Segment Routing Header(SRH) as defined in [[I-D.ietf-6man-spring-srv6-oam](#)].

3. Advertising Supported Algorithms

SRv6 capable router indicates supported algorithm(s) by advertising the SR Algorithm TLV as defined in [[RFC8667](#)].

[4.](#) Advertising Maximum SRv6 SID Depths

[RFC8491] defines the means to advertise node/link specific values for Maximum SID Depths (MSD) of various types. Node MSDs are advertised in a sub-TLV of the Router Capabilities TLV [[RFC7981](#)]. Link MSDs are advertised in a sub-TLV of TLVs 22, 23, 141, 222, and 223.

This document defines the relevant SRv6 MSDs and requests MSD type assignments in the MSD Types registry created by [[RFC8491](#)].

[4.1.](#) Maximum Segments Left MSD Type

The Maximum Segments Left MSD Type specifies the maximum value of the "SL" field [[I-D.ietf-6man-segment-routing-header](#)] in the SRH of a received packet before applying the Endpoint behavior associated with a SID.

SRH Max SL Type: 41 (Suggested value - to be assigned by IANA)

If no value is advertised the supported value is assumed to be 0.

[4.2.](#) Maximum End Pop MSD Type

The Maximum End Pop MSD Type specifies the maximum number of SIDs in the top SRH in an SRH stack to which the router can apply "PSP" or USP" as defined in [[I-D.ietf-spring-srv6-network-programming](#)] flavors.

SRH Max End Pop Type: 42 (Suggested value - to be assigned by IANA)

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

[4.3.](#) Maximum H.Encaps MSD Type

The Maximum H.Encaps MSD Type specifies the maximum number of SIDs that can be included as part of the "H.Encaps" behavior as defined in [[I-D.ietf-spring-srv6-network-programming](#)] .

SRH Max H.encaps Type: 44 (Suggested value - to be assigned by IANA)

If the advertised value is zero then the router can apply H.Encaps only by encapsulating the incoming packet in another IPv6 header without SRH the same way IPinIP encapsulation is performed.

If the advertised value is non-zero then the router supports both IPinIP and SRH encapsulation subject to the SID limitation specified by the advertised value.

4.4. Maximum End D MSD Type

The Maximum End D MSD Type specifies the maximum number of SIDs in an SRH when performing decapsulation associated with "End.Dx" behaviors (e.g., "End.DX6" and "End.DT6") as defined in [\[I-D.ietf-spring-srv6-network-programming\]](#).

SRH Max End D Type: 45 (Suggested value - to be assigned by IANA)

If the advertised value is zero or no value is advertised then it is assumed that the router cannot apply "End.DX6" or "End.DT6" behaviors if the extension header right underneath the outer IPv6 header is an SRH.

5. SRv6 SIDs and Reachability

As discussed in [\[I-D.ietf-spring-srv6-network-programming\]](#), 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 which have the matching topology/algorithm.

Locators MUST be advertised in the SRv6 Locator TLV (see [Section 6.1](#)). 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.

Locators are routable and MAY also be advertised in Prefix Reachability TLVs (236 or 237).

Locators associated with Flexible Algorithms [\[I-D.ietf-lsr-flex-algo\]](#) SHOULD NOT be advertised in Prefix Reachability TLVs (236 or 237).

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

In cases where a locator advertisement is received in both in 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 on SRv6 capable/SRv6 incapable routers.

SRv6 SIDs are advertised as sub-TLVs in the SRv6 Locator TLV except for SRv6 End.X SIDs/LAN End.X SIDs which are associated with a specific Neighbor/Link and are therefore advertised as sub-TLVs in TLVs 22, 23, 222, 223, and 141.

SRv6 SIDs 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 assure 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. There are a number of ways in which this requirement could be compromised

- o Another locator associated with a different topology/algorithm is the longest match
- o A 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 and 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 "Bit Values for Prefix Attribute Flags Sub-TLV" registry [[RFC7794](#)] is defined to advertise the anycast property:

Bit #: 4 (Suggested - to be assigned by IANA)
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 leaked between levels.

The A-flag and the N-flag MUST NOT both be set.

If both N-flag and 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.

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																			
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-

Followed by one or more locator entries of the form:

[illegible]


```
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

Type: 27 (Suggested value to be assigned by IANA)

Length: variable.

MTID: Multitopology Identifier as defined in [[RFC5120](#)].

Note that the value 0 is legal.

Locator entry:

Metric: 4 octets. As described in [[RFC5305](#)].

Flags: 1 octet. The following flags are defined

```

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

```

where:

D bit: When the Locator is leaked from level-2 to level-1, the D bit MUST be set. Otherwise, this bit MUST be clear. Locators with the D bit set MUST NOT be leaked from level-1 to level-2. This is to prevent looping.

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

Algorithm: 1 octet. Associated algorithm. Algorithm values are defined in the IGP Algorithm Type registry.

Loc-Size: 1 octet. Number of bits in the Locator field.
(1 - 128)

Locator: 1-16 octets. This field encodes the advertised SRv6 Locator. The Locator is encoded in the minimal number of octets for the given number of bits.

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

Optional sub-TLVs.

7.2. SRv6 End SID sub-TLV

The SRv6 End SID sub-TLV is introduced to advertise SRv6 Segment Identifiers (SID) with Endpoint behaviors which do not require a particular neighbor in order to be correctly applied [[I-D.ietf-spring-srv6-network-programming](#)]. SRv6 SIDs associated with a neighbor are advertised using the sub-TLVs defined in [Section 6](#).

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:


```

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

```

Type: 5 (Suggested value to be assigned by IANA)

Length: variable.

Flags: 1 octet. No flags are currently defined.

Endpoint Behavior: 2 octets, as defined in [[I-D.ietf-spring-srv6-network-programming](#)].

Legal behavior values for this sub-TLV are defined in [Section 9](#) of this document.

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

The SRv6 End SID MUST be a subnet of the associated Locator. SRv6 End SIDs which are NOT a subnet of 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

[[I-D.ietf-spring-srv6-network-programming](#)] must be associated with a particular neighbor, and in case of multiple links to the same neighbor, with a particular link in order to be correctly applied.

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

IS-IS Neighbor advertisements are topology specific - but not algorithm specific. End.X SIDs therefore inherit the topology from the associated neighbor advertisement, but the algorithm is specified in the individual SID.

All End.X SIDs MUST be a subnet of a Locator with matching topology and algorithm which is advertised by the same node in an SRv6 Locator TLV. End.X SIDs which do not meet this requirement MUST be ignored.

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-tlv-len|      Sub-sub-TLVs (variable) . . .      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type: 43 (Suggested value to be assigned by IANA)

Length: variable.

Flags: 1 octet.

```

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

```

where:

B-Flag: Backup flag. If set, the End.X SID is eligible for protection (e.g., using IPFRR) as described in [\[RFC8355\]](#).

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

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

Algorithm: 1 octet. Associated algorithm. Algorithm values are defined in the IGP Algorithm Type registry.

Weight: 1 octet. The value represents the weight of the End.X SID for the purpose of load balancing. The use of the weight is defined in [\[I-D.ietf-spring-segment-routing\]](#).

Endpoint Behavior: 2 octets. As defined in [\[I-D.ietf-spring-srv6-network-programming\]](#) Legal behavior values for this sub-TLV are defined in [Section 9](#).

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

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

Note that multiple TLVs for the same neighbor may be required in order to advertise all of the SRv6 End.X 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 a large number of neighbors may exist on a given LAN a large number of 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 of the SRv6 End.X 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      |   System ID (6 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 (Suggested value to be assigned by IANA)

Length: variable.

System-ID: 6 octets of 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 |
+---+---+---+---+---+

```


where B,S, and P flags are as described in [Section 8.1](#).
Other bits: MUST be zero when originated and ignored when received.

Algorithm: 1 octet. Associated algorithm. Algorithm values are defined in the IGP Algorithm Type registry.

Weight: 1 octet. The value represents the weight of the End.X SID for the purpose of load balancing. The use of the weight is defined in [[I-D.ietf-spring-segment-routing](#)].

Endpoint Behavior: 2 octets. As defined in [[I-D.ietf-spring-srv6-network-programming](#)]
Legal behavior values for this sub-TLV are defined in [Section 9](#).

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

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

9. SRv6 SID Structure Sub-Sub-TLV

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](#))

SRv6 SID Structure Sub-Sub-TLV is used to advertise the length of each individual part of the SRv6 SID as defined in [[I-D.ietf-spring-srv6-network-programming](#)]. 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type      |      Length      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  LB Length  |  LN Length  | Fun. Length  | Arg. Length  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

Type: 1

Length: 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.

ISIS 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 TLV, the parent TLV MUST be ignored by the receiver.

10. Advertising Endpoint Behaviors

Endpoint behaviors are defined in [[I-D.ietf-spring-srv6-network-programming](#)] and [[I-D.ietf-6man-spring-srv6-oam](#)]. The codepoints for the Endpoint behaviors are defined in the "SRv6 Endpoint Behaviors" registry defined in [[I-D.ietf-spring-srv6-network-programming](#)]. This section lists the Endpoint behaviors and their codepoints, which MAY be advertised by IS-IS and the SID sub-TLVs in which each type MAY appear.

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.T (PSP, USP, USD)	9-12, 36-39	Y	N	N
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.DT64	20	Y	N	N
End.OP	40	Y	N	N
End.OTP	41	Y	N	N

[11.](#) Implementation Status

This section describes the implementation status of the ISIS SRv6 extensions.

[11.1.](#) Cisco

Cisco's ISIS SRv6 implementation supports following functionalities:

Types of SID supported: End, End.X, LAN End.X, END.OP

Intra/Inter area/level support: Yes

Anycast SID support: Yes, including A-flag ([Section 6](#))

SID Structure Sub-Sub-TLV: Yes

[11.2.](#) Huawei

Huawei ISIS SRv6 implementation supports following functionalities:

Types of SID supported: End, End.X

Intra/Inter area/level support: Yes

Anycast SID support: Yes, no A-flag support ([Section 6](#))

SID Structure Sub-Sub-TLV: No

[11.3.](#) Juniper

Juniper's implementation is in beta stage and supports most of the sections of the this draft.

[11.4.](#) Interoperability Testing.

EANTC interoperability testing has been performed in April 2019. Results can be found at EANTC site [[1](#)].

[12.](#) IANA Considerations

This document requests allocation for the following TLVs, sub-TLVs, and sub-sub-TLVs as well updating the ISIS TLV registry and defining a new registry.

[12.1.](#) SRv6 Locator TLV

This document adds one new TLV to the IS-IS TLV Codepoints registry.

Value: 27 (suggested - to be assigned by IANA)

Name: SRv6 Locator

This TLV shares sub-TLV space with existing "Sub-TLVs for TLVs 135, 235, 236 and 237 registry". The name of this registry needs to be changed to "Sub-TLVs for TLVs 27, 135, 235, 236 and 237 registry".

[12.1.1.](#) SRv6 End SID sub-TLV

This document adds the following new sub-TLV to the (renamed) "Sub-TLVs for TLVs 27, 135, 235, 236 and 237 registry".

Value: 5 (suggested - to be assigned by IANA)

Name: SRv6 End SID

This document requests the creation of a new IANA managed registry for sub-sub-TLVs of the SRv6 End SID sub-TLV. The registration procedure is "Expert Review" as defined in [[RFC7370](#)]. Suggested registry name is "sub-sub-TLVs for SRv6 End SID sub-TLV". No sub-sub-TLVs are defined by this document except for the reserved value.

0: Reserved

1-255: Unassigned

12.1.2. Revised sub-TLV table

The revised table of sub-TLVs for the (renamed) "Sub-TLVs for TLVs 27, 135, 235, 236 and 237 registry" is shown below:

Type	27	135	235	236	237
1	n	y	y	y	y
2	n	y	y	y	y
3	n	y	y	y	y
4	y	y	y	y	y
5	y	n	n	n	n
11	y	y	y	y	y
12	y	y	y	y	y

12.2. SRv6 Capabilities sub-TLV

This document adds the definition of a new sub-TLV in the "Sub- TLVs for TLV 242 registry".

Type: 25 (Suggested - to be assigned by IANA)

Description: SRv6 Capabilities

This document requests the creation of a new IANA managed registry for sub-sub-TLVs of the SRv6 Capability sub-TLV. The registration procedure is "Expert Review" as defined in [[RFC7370](#)]. Suggested registry name is "sub-sub-TLVs for SRv6 Capability sub-TLV". No sub-sub-TLVs are defined by this document except for the reserved value.

0: Reserved

1-255: Unassigned

12.3. SRv6 End.X SID and SRv6 LAN End.X SID sub-TLVs

This document adds the definition of two new sub-TLVs in the "sub-TLVs for TLV 22, 23, 25, 141, 222 and 223 registry".

Type: 43 (suggested - to be assigned by IANA)

Description: SRv6 End.X SID

Type: 44 (suggested - to be assigned by IANA)

Description: SRv6 LAN End.X SID

Type	22	23	25	141	222	223
43	y	y	y	y	y	y
44	y	y	y	y	y	y

12.4. MSD Types

This document defines the following new MSD types. These types are to be defined in the IGP MSD Types registry defined in [[RFC8491](#)] .

All values are suggested values to be assigned by IANA.

Type	Description

41	SRH Max SL
42	SRH Max End Pop
44	SRH Max H.encaps
45	SRH Max End D

12.5. Sub-Sub-TLVs for SID Sub-TLVs

This document creates the following Sub-Sub-TLV Registry within the "IS-IS TLV Codepoints" registry:

Registry: Sub-Sub-TLVs for SID Sub-TLVs

Registration Procedure: Expert review

Reference: This document ([Section 7.2](#), [Section 8.1](#), [Section 8.2](#)).

This document defines the following Sub-Sub-TLV in the "Sub-Sub-TLVs for SID Sub-TLVs" registry:

Type: 1

Description: SRv6 SID Structure Sub-Sub-TLV.

Reference: This document ([Section 9](#)).

12.6. Prefix Attribute Flags Sub-TLV

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

Bit #: 4

Description: A bit

Reference: This document ([Section 6](#)).

[13.](#) Security Considerations

Security concerns for IS-IS are addressed in [[ISO10589](#)], [[RFC5304](#)], and [[RFC5310](#)].

[14.](#) Contributors

The following people gave a substantial contribution to the content of this document and should be considered as co-authors:

Stefano Previdi
Huawei Technologies
Email: stefano@previdi.net

Paul Wells
Cisco Systems
Saint Paul,
Minnesota
United States
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

15. References

15.1. Normative References

[I-D.ietf-6man-segment-routing-header]

Filsfils, C., Dukes, D., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", [draft-ietf-6man-segment-routing-header-26](#) (work in progress), October 2019.

[I-D.ietf-6man-spring-srv6-oam]

Ali, Z., Filsfils, C., Matsushima, S., Voyer, D., and M. Chen, "Operations, Administration, and Maintenance (OAM) in Segment Routing Networks with IPv6 Data plane (SRv6)", [draft-ietf-6man-spring-srv6-oam-03](#) (work in progress), December 2019.

[I-D.ietf-spring-srv6-network-programming]

Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 Network Programming", [draft-ietf-spring-srv6-network-programming-08](#) (work in progress), January 2020.

[ISO10589]

Standardization", I. "O. F., "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), ISO/IEC 10589:2002, Second Edition.", Nov 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>>.

[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>>.

[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>>.

- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", [RFC 8667](#), DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.

[15.2. Informative References](#)

- [I-D.ietf-lsr-flex-algo]
Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", [draft-ietf-lsr-flex-algo-05](#) (work in progress), November 2019.
- [I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-15](#) (work in progress), January 2018.

[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>>.

15.3. URIs

[1] http://www.eantc.de/fileadmin/eantc/downloads/events/2017-2020/MPLS2019/Segment_Routing_IPv6_Srv6_.pdf

Authors' Addresses

Peter Psenak (editor)
Cisco Systems
Pribinova Street 10
Bratislava 81109
Slovakia

Email: ppsenak@cisco.com

Clarence Filsfils
Cisco Systems
Brussels
Belgium

Email: cfilsfil@cisco.com

Ahmed Bashandy
Arcus

Email: abashandy.ietf@gmail.com

Bruno Decraene
Orange
Issy-les-Moulineaux
France

Email: bruno.decraene@orange.com

Zhibo Hu
Huawei Technologies

Email: huzhibo@huawei.com

