

Link State Routing
Internet-Draft
Intended status: Standards Track
Expires: February 15, 2021

Z. Li
Z. Hu
D. Cheng
Huawei Technologies
K. Talaulikar, Ed.
P. Psenak
Cisco Systems
August 14, 2020

OSPFv3 Extensions for SRv6
draft-ietf-lsr-ospfv3-srv6-extensions-01

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 OSPFv3 extensions required to support Segment Routing over an IPv6 data plane (SRv6).

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 February 15, 2021.

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	2
1.1.	Requirements Language	3
2.	SRV6 Capabilities TLV	3
3.	Advertisement of Supported Algorithms	5
4.	Advertisement of SRH Operation Limits	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.	Advertisement of SRV6 Locator and End SIDs	6
6.	SRV6 Locator LSA	7
6.1.	SRV6 Locator TLV	9
7.	Advertisement of SRV6 End SIDs	11
8.	Advertisement of SRV6 SIDs Associated with Adjacencies	12
8.1.	SRV6 End.X SID Sub-TLV	13
8.2.	SRV6 LAN End.X SID Sub-TLV	15
9.	SRV6 SID Structure Sub-TLV	16
10.	Advertising Endpoint Behaviors	17
11.	Security Considerations	18
12.	IANA Considerations	19
12.1.	OSPF Router Information TLVs	19
12.2.	OSPFv3 LSA Function Codes	19
12.3.	OSPFv3 Extended-LSA Sub-TLVs	19
12.4.	OSPFv3 Locator LSA TLVs	19
12.5.	OSPFv3 Locator LSA Sub-TLVs	20
13.	Acknowledgements	20
14.	References	21
14.1.	Normative References	21
14.2.	Informative References	22
	Authors' Addresses	23

[1.](#) Introduction

Segment Routing (SR) architecture [[RFC8402](#)] specifies how a node can steer a packet through an ordered list of instructions, called segments. These segments are identified through Segment Identifiers (SIDs).

Segment Routing can be 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 dataplane. The network programming paradigm for SRv6 is specified in [\[I-D.ietf-spring-srv6-network-programming\]](#) which describes several well-known behaviors that can be bound to SRv6 SIDs.

This document specifies extensions to OSPFv3 in order to support SRv6 as defined in [\[I-D.ietf-spring-srv6-network-programming\]](#) by signaling the SRv6 capabilities of the node and certain SRv6 SIDs with their endpoint behaviors (e.g., End, End.X, etc.) that are instantiated on the SRv6 capable router.

At a high level, the extensions to OSPFv3 are comprised of the following:

1. SRv6 Capabilities TLV to advertise the SRv6 features and SRH operations supported by the router
2. SRv6 Locator TLV to advertise the SRv6 Locator - a form of summary address for the algorithm specific SIDs instantiated on the router
3. TLVs and Sub-TLVs to advertise the SRv6 SIDs instantiated on the router along with their endpoint behaviors

[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 TLV](#)


The SRv6 Capabilities TLV is used by an OSPFv3 router to advertise its SRv6 support along with its related capabilities for SRv6 functionality. This is an optional top level TLV of the OSPFv3 Router Information LSA [\[RFC7770\]](#) which MUST be advertised by an SRv6 enabled router.

This TLV SHOULD be advertised only once in the OSPFv3 Router Information LSA. When multiple SRv6 Capabilities TLVs are received from a given router, the receiver MUST use the first occurrence of the TLV in the OSPFv3 Router Information Opaque LSA. If the SRv6 Capabilities TLV appears in multiple OSPFv3 Router Information Opaque LSAs that have different flooding scopes, the TLV in the OSPFv3 Router Information Opaque LSA with the area-scoped flooding scope MUST be used. If the SRv6 Capabilities TLV appears in multiple

OSPFv3 Router Information Opaque LSAs that have the same flooding scope, the TLV in the OSPFv3 Router Information Opaque LSA with the numerically smallest Instance ID MUST be used and subsequent instances of the TLV MUST be ignored.

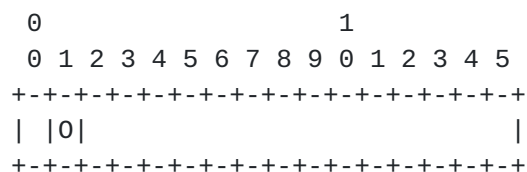
The OSPFv3 Router Information Opaque LSA can be advertised at any of the defined opaque flooding scopes (link, area, or Autonomous System (AS)). For the purpose of SRv6 Capabilities TLV advertisement, area-scoped flooding is REQUIRED.

The format of OSPFv3 SRv6 Capabilities TLV is shown below



Where:

- o Type: 16 bit field. Value is TBD.
- o Length: 16 bit field. Length of Capability TLV + length of Sub-TLVs
- o Reserved : 16 bit field. SHOULD be set to 0 and MUST be ignored on receipt.
- o Flags: 16 bit field. The following flags are defined and others SHOULD be set to 0 and MUST be ignored on receipt:



where:

- * 0-flag: If set, then the router is capable of supporting the SRH 0-bit, as specified in [[I-D.ietf-6man-spring-srv6-oam](#)].

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

3. Advertisement of Supported Algorithms

SRv6 enabled OSPFv3 router advertises its algorithm support using the SR Algorithm TLV defined in [[RFC8665](#)] as described in [[RFC8666](#)].

4. Advertisement of SRH Operation Limits

An SRv6 enabled router may have different capabilities and limits when it comes to SRH processing and this needs to be advertised to other routers in the SRv6 domain.

[RFC8476] defines the means to advertise node/link specific values for Maximum SID Depth (MSD) types. Node MSDs are advertised using the Node MSD TLV in the OSPFv3 Router Information LSA [[RFC7770](#)] while Link MSDs are advertised using the Link MSD Sub-TLV of the E-Router-LSA TLV [[RFC8362](#)]. The format of the MSD types for OSPFv3 is defined in [[RFC8476](#)].

The MSD types for SRv6 that are defined in section 4 of [[I-D.ietf-lsr-isis-srv6-extensions](#)] for IS-IS are also used by OSPFv3. These MSD Types are allocated under the IGP MSD Types registry maintained by IANA that are shared by IS-IS and OSPF. They are described below:

4.1. Maximum Segments Left MSD Type

The Maximum Segments Left MSD Type specifies the maximum value of the "SL" field [[RFC8754](#)] in the SRH of a received packet before applying the Endpoint behavior associated with a SID. 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 SRH for which the router can apply Penultimate Segment Pop (PSP) or Ultimate Segment Pop (USP) as defined in [[I-D.ietf-spring-srv6-network-programming](#)] flavors. If the advertised value is zero or no value is advertised, then it is assumed that the router cannot apply PSP or USP.

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](#)]. 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 [\[I-D.ietf-spring-srv6-network-programming\]](#). 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. Advertisement of SRv6 Locator and End SIDs

An SRv6 Segment Identifier (SID) is 128 bits and comprises of Locator, Function and Argument parts as described in [\[I-D.ietf-spring-srv6-network-programming\]](#).

A node is provisioned with algorithm specific locators for each algorithm supported by that node. Each locator is a prefix subsuming all SIDs provisioned on that node which have the matching algorithm.

Locators MUST be advertised in the SRv6 Locator LSA (see [Section 6](#)). Forwarding entries for the locators advertised in the SRv6 Locator LSA MUST be installed in the forwarding plane of receiving SRv6 capable routers when the associated algorithm is supported by the receiving node. Locators can be of different route types similar to existing OSPF LSA route types - Intra-Area, Inter-Area, External, and NSSA. The computation of locator reachability and their advertisement are similar to how normal OSPF prefix reachability LSAs are processed as part of the route computation.

Locators are routable and MAY also be advertised via Prefix LSAs of different types - Inter-Area Prefix LSA, AS-External LSA, NSSA LSA, or Intra-Area Prefix LSA (or their equivalent extended LSAs [\[RFC8362\]](#)). Locators associated with Flexible Algorithms [\[I-D.ietf-lsr-flex-algo\]](#) SHOULD NOT be advertised via Prefix LSAs. Locators associated with algorithm 0 (for all supported topologies) SHOULD be advertised in Prefix LSAs 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 LSA and an SRv6 Locator LSA, the Prefix LSA 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 of E-Router-Link TLV.

SRv6 SIDs are not directly routable. SRv6 SIDs learnt by via advertisements from remote routers 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 algorithm will be forwarded correctly, while SRv6 SIDs associated with an unsupported algorithm will be dropped. NOTE: The drop behavior depends on the absence of a default/summary route covering a given locator.

6. SRv6 Locator LSA

The SRv6 Locator LSA has a function code of TBD while the S1/S2 bits are dependent on the desired flooding scope for the LSA. The flooding scope of the SRv6 Locator LSA depends on the scope of the advertised SRv6 Locator and is under the control of the advertising router. The U bit will be set indicating that the LSA should be flooded even if it is not understood.

Multiple SRv6 Locator LSAs can be advertised by an OSPFv3 router and they are distinguished by their Link State IDs (which are chosen arbitrarily by the originating router).

The format of SRv6 Locator LSA is shown below:

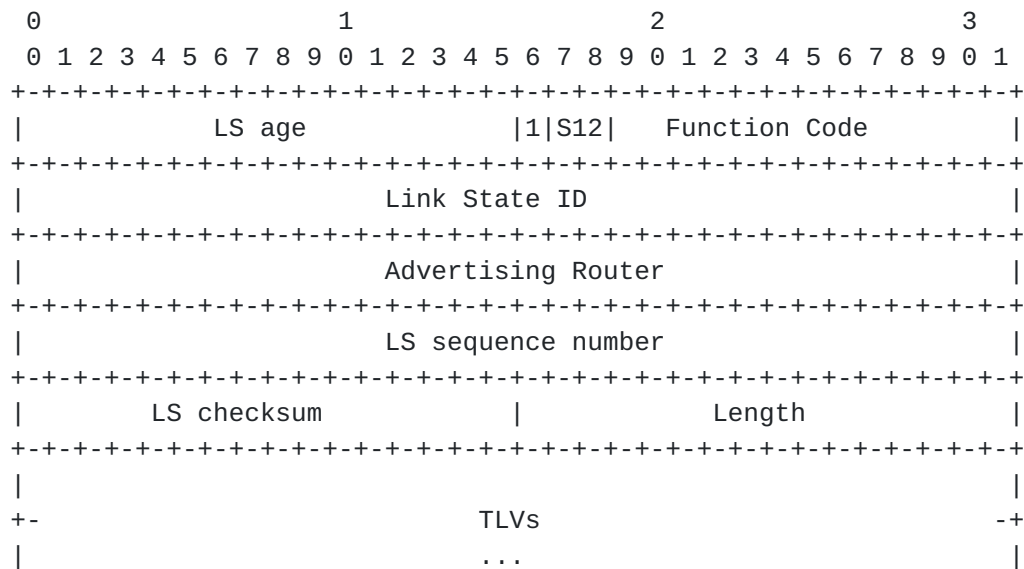


Figure 1: SRv6 Locator LSA

The format of the TLVs within the body of the SRv6 Locator LSA is the same as the format used by [\[RFC3630\]](#). The variable TLV section consists of one or more nested TLV tuples. Nested TLVs are also referred to as Sub-TLVs. The format of each TLV is:




Figure 2: SRv6 Locator LSA TLV Format

The Length field defines the length of the value portion in octets (thus, a TLV with no value portion would have a length of 0). The TLV is padded to 4-octet alignment; padding is not included in the Length field (so a 3-octet value would have a length of 3, but the total size of the TLV would be 8 octets). Nested TLVs are also 32-bit aligned. For example, a 1-byte value would have the Length field set to 1, and 3 octets of padding would be added to the end of the value portion of the TLV. The padding is composed of zeros.

6.1. SRv6 Locator TLV

The SRv6 Locator TLV is a top-level TLV of the SRv6 Locator LSA that is used to advertise an SRv6 Locator, its attributes, and SIDs associated with it. Multiple SRv6 Locator TLVs MAY be advertised in each SRv6 Locator LSA. However, since the S12 bits define the flooding scope, the LSA flooding scope MUST satisfy the application-specific requirements for all the locators included in a single SRv6 Locator LSA.

When multiple SRv6 Locator TLVs are received from a given router in an SRv6 Locator LSA for the same Locator, the receiver MUST use the first occurrence of the TLV in the LSA. If the SRv6 Locator TLV for the same Locator appears in multiple SRv6 Locator LSAs that have different flooding scopes, the TLV in the SRv6 Locator LSA with the area-scoped flooding scope MUST be used. If the SRv6 Locator TLV for the same Locator appears in multiple SRv6 Locator LSAs that have the same flooding scope, the TLV in the SRv6 Locator LSA with the numerically smallest Link-State ID MUST be used and subsequent instances of the TLV MUST be ignored.

The format of SRv6 Locator TLV is shown below:

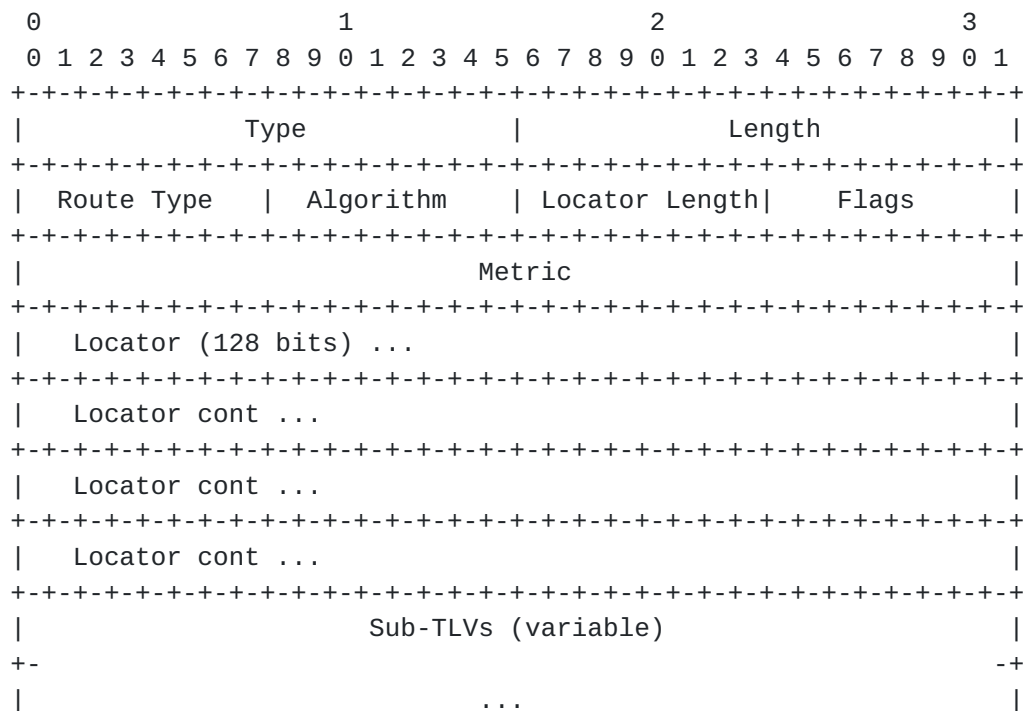


Figure 3: SRv6 Locator TLV

Where:

Type: 16 bit field. The value is 1 for this type.

Length: 16 bit field. The total length of the value portion of the TLV including Sub-TLVs.

Route Type : 8 bit field. The type of the locator route. Supported types are the ones listed below and other other types MUST be ignored on receipt.

- 1 - Intra-Area
- 2 - Inter-Area
- 3 - AS External
- 4 - NSSA External

Figure 4

Algorithm: 8 bit field. Associated algorithm. Algorithm values are defined in the IGP Algorithm Type registry.

Locator Length: 8 bit field. Carries the length of the Locator prefix as the number of locator bits (1-128).

Flags: 8 bit field. The following flags are defined

```

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

Figure 5

- * N bit : When the locator uniquely identifies a node in the network (i.e., it is provisioned on one and only one node), the N bit MUST be set. Otherwise, this bit MUST be clear.
- * A bit : When the Locator is configured as anycast, the A bit SHOULD be set. Otherwise, this bit MUST be clear. If both the N and A bits are set, then the receiving routers MUST ignore the N bit (i.e., consider it as not set).
- * Other flags are not defined and SHOULD be set to 0 and MUST be ignored on receipt.

Metric : 32 bit field. The metric value associated with the locator.

Locator : 128 bit field. This field encodes the advertised SRv6 Locator.

Sub-TLVs : Used to advertise Sub-TLVs that provide additional attributes for the given SRv6 Locator and SRv6 SIDs associated with it.

7. Advertisement of SRv6 End SIDs

The SRv6 End SID Sub-TLV is a Sub-TLV of the SRv6 Locator TLV in the SRv6 Locator LSA (defined in [Section 6](#)). It is used to advertise the SRv6 SIDs belonging to the node along with their associated endpoint behaviors. SIDs associated with adjacencies are advertised as described in [Section 8](#). Every SRv6 enabled OSPFv3 router SHOULD advertise at least one SRv6 SID associated with an END behavior for its node as specified in [[I-D.ietf-spring-srv6-network-programming](#)].

SRv6 End SIDs inherit the algorithm from the parent locator. The SRv6 End SID MUST be contained in the subnet of the associated Locator. SRv6 End SIDs which are NOT in a subnet of the associated locator MUST be ignored.

The router MAY advertise multiple instances of the SRv6 End SID Sub-TLV within the SRv6 Locator TLV - one for each of the SRv6 SIDs to be advertised. When multiple SRv6 End SID Sub-TLVs are received in the SRv6 Locator TLV from a given router for the same SRv6 SID value, the receiver MUST use the first occurrence of the Sub-TLV in the SRv6 Locator TLV.

The format of SRv6 End SID Sub-TLV is shown below

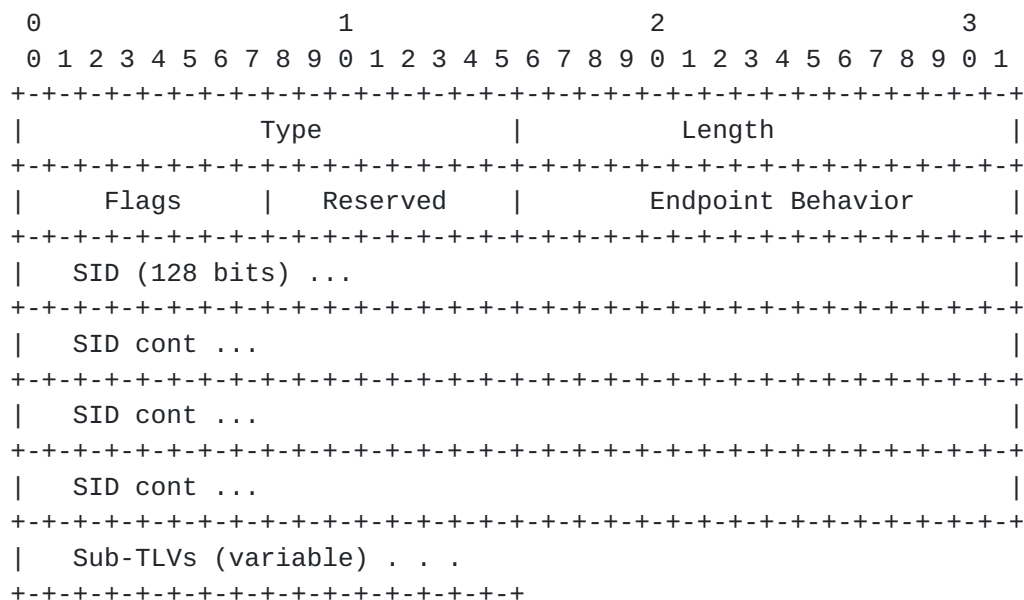


Figure 6: SRv6 End SID Sub-TLV

Where:

Type: 16 bit field. Value is 1 for this type.

Length: 16 bit field. The total length of the value portion of the Sub-TLVs.

Reserved : 8 bit field. Should be set to 0 and MUST be ignored on receipt.

Flags: 8 bit field which define the flags associated with the SID. No flags are currently defined and SHOULD be set to 0 and MUST be ignored on receipt.

Endpoint Behavior: 16 bit field. The endpoint behavior code point for this SRv6 SID as defined in section 9.2 of [\[I-D.ietf-spring-srv6-network-programming\]](#).

SID : 128 bit field. This field encodes the advertised SRv6 SID.

Sub-TLVs : Used to advertise Sub-TLVs that provide additional attributes for the given SRv6 SID.

8. Advertisement of SRv6 SIDs Associated with Adjacencies

The SRv6 endpoint behaviors are defined in [\[I-D.ietf-spring-srv6-network-programming\]](#) include certain behaviors which are specific to links or adjacencies. The most basic of these which is critical for link state routing protocols like OSPFv3 is the End.X behavior that is an instruction to forward to a specific neighbor on a specific link. These SRv6 SIDs along with others that are defined in [\[I-D.ietf-spring-srv6-network-programming\]](#) which are specific to links or adjacencies need to be advertised by OSPFv3 so that this information is available to all routers in the area to influence the packet path via these SRv6 SIDs over the specific adjacencies.

The advertisement of SRv6 SIDs and their behaviors that are specific to a particular neighbor is done via two different optional Sub-TLVs of the E-Router-Link TLV defined in [\[RFC8362\]](#) as follows:

- o SRv6 End.X SID Sub-TLV: For OSPFv3 adjacencies over point-to-point or point-to-multipoint links and the adjacency to the Designated Router (DR) over broadcast and non-broadcast-multi-access (NBMA) links.
- o SRv6 LAN End.X SID Sub-TLV: For OSPFv3 adjacencies on broadcast and NBMA links to the Backup DR and DR-Other neighbors. This Sub-

TLV includes the OSPFv3 router-id of the neighbor and thus allows for an instance of this Sub-TLV for each neighbor to be explicitly advertised under the E-Router-Link TLV for the same link.

Every SRv6 enabled OSPFv3 router SHOULD instantiate at least one unique SRv6 End.X SID corresponding to each of its neighbor. A router MAY instantiate more than one SRv6 End.X SID for for a single neighbor. The same SRv6 End.X SID MAY be advertised for more than one neighbor. Thus multiple instances of the SRv6 End.X SID and SRv6 LAN End.X SID Sub-TLVs MAY be advertised within the E-Router-Link TLV for a single link.

All End.X and LAN End.X SIDs MUST be subsumed by the subnet of a Locator with the matching 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. This ensures that the node advertising the End.X or LAN End.X SID 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

The format of the SRv6 End.X SID Sub-TLV is shown below

```

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                                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      Endpoint Behavior      |      Flags      |      Reserved1      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      Algorithm      |      Weight      |      Reserved2      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      SID (128 bits) ...      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      SID cont ...      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      SID cont ...      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      SID cont ...      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      Sub-TLVs (variable) . . .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

```

Where:

Type: 16 bit field. Value is TBD.

Length: 16 bit field. The total length of the value portion of the TLV.

Endpoint Behavior: 16 bit field. The code point for the endpoint behavior for this SRv6 SID as defined in section 9.2 of [[I-D.ietf-spring-srv6-network-programming](#)].

Flags: 8 bit field with the following definition:

```

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

```

- * B-Flag: Backup Flag. If set, the SID refers to a path that is eligible for protection.
- * 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: If set, the SID is persistently allocated, i.e., the SID value remains consistent across router restart and session/interface flap.
- * Reserved bits: Reserved for future use and MUST be zero when originated and ignored on receipt.

Reserved1 : 8 bit field. Should be set to 0 and MUST be ignored on receipt.

Algorithm : 8 bit field. Associated algorithm. Algorithm values are defined in the IGP Algorithm Type registry.

Weight: 8 bit field whose value represents the weight of the End.X SID for the purpose of load-balancing. The use of the weight is defined in [[RFC8402](#)].

Reserved2 : 16 bit field. Should be set to 0 and MUST be ignored on receipt.

SID: 128 bit field. This field encodes the advertised SRv6 SID.

Sub-TLVs : Used to advertise Sub-TLVs that provide additional attributes for the given SRv6 End.X SID.

8.2. SRv6 LAN End.X SID Sub-TLV

The format of the SRv6 LAN End.X SID Sub-TLV is as shown below

```

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                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      Endpoint Behavior      |      Flags      |   Reserved1   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Algorithm    |     Weight     |           Reserved2          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| OSPFv3 Router-ID of neighbor                                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| SID (128 bits) ...                                              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| SID cont ...                                                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| SID cont ...                                                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| SID cont ...                                                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Sub-TLVs (variable) . . .                                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

```

Where

- o Type: 16 bit field. Value is TBD.
- o Length: 16 bit field. Variable
- o Endpoint Behavior: 16 bit field. The code point for the endpoint behavior for this SRv6 SID as defined in section 9.2 of [\[I-D.ietf-spring-srv6-network-programming\]](#).
- o SID Flags: 8 bit field which define the flags associated with the SID. No flags are currently defined and SHOULD be set to 0 and MUST be ignored on receipt.
- o Flags: 8 bit field with the following definition:

```

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

```


- * B-Flag: Backup Flag. If set, the SID refers to a path that is eligible for protection.
 - * 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: If set, the SID is persistently allocated, i.e., the SID value remains consistent across router restart and session/interface flap.
 - * Reserved bits: Reserved for future use and MUST be zero when originated and ignored on receipt.
- o Reserved1 : 8 bit field. Should be set to 0 and MUST be ignored on receipt.
 - o Algorithm : 8 bit field. Associated algorithm. Algorithm values are defined in the IGP Algorithm Type registry.
 - o Weight: 8 bit field whose value represents the weight of the End.X SID for the purpose of load balancing. The use of the weight is defined in [[RFC8402](#)].
 - o Reserved2 : 16 bit field. Should be set to 0 and MUST be ignored on receipt.
 - o Neighbor ID : 32 bits of OSPFv3 Router-id of the neighbor
 - o SID: 128 bit field. This field encodes the advertised SRv6 SID.
 - o Sub-TLVs : Used to advertise Sub-TLVs that provide additional attributes for the given SRv6 SID.

9. SRv6 SID Structure Sub-TLV

SRv6 SID Structure 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 is used as an optional Sub-TLV of the following:

- o SRv6 End SID Sub-TLV (refer [Section 7](#))
- o SRv6 End.X SID Sub-TLV (refer [Section 8.1](#))
- o SRv6 LAN End.X SID Sub-TLV (refer [Section 8.2](#))

The Sub-TLV has the following format:

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

Figure 8: SRv6 Endpoint Behaviors in OSPFv3

11. Security Considerations

Existing security extensions as described in [\[RFC5340\]](#) and [\[RFC8362\]](#) apply to these SRv6 extensions. While OSPFv3 is under a single administrative domain, there can be deployments where potential attackers have access to one or more networks in the OSPFv3 routing domain. In these deployments, stronger authentication mechanisms such as those specified in [\[RFC4552\]](#) or [\[RFC7166\]](#) SHOULD be used.

Implementations MUST assure that malformed TLV and Sub-TLV defined in this document are detected and do not provide a vulnerability for attackers to crash the OSPFv3 router or routing process. Reception of malformed TLV or Sub-TLV SHOULD be counted and/or logged for further analysis. Logging of malformed TLVs and Sub-TLVs SHOULD be rate-limited to prevent a Denial of Service (DoS) attack (distributed or otherwise) from overloading the OSPFv3 control plane.

12. IANA Considerations

This document specifies updates to multiple OSPF and OSPFv3 related IANA registries as follows.

12.1. OSPF Router Information TLVs

This document proposes the following new code point in the "OSPF Router Information (RI) TLVs" registry under the "OSPF Parameters" registry for the new TLVs:

Type TBD (suggested 20): SRv6-Capabilities TLV: Refer to [Section 2](#).

12.2. OSPFv3 LSA Function Codes

This document proposes the following new code point in the "OSPFv3 LSA Function Codes" registry under the "OSPFv3 Parameters" registry for the new SRv6 Locator LSA:

- o Type TBD (suggested 42): SRv6 Locator LSA: Refer to [Section 6](#).

12.3. OSPFv3 Extended-LSA Sub-TLVs

This document proposes the following new code points in the "OSPFv3 Extended-LSA Sub-TLVs" registry under the "OSPFv3 Parameters" registry for the new Sub-TLVs:

- o Type TBD (suggested 10): SRv6 SID Structure Sub-TLV : Refer to [Section 9](#).
- o Type TBD (suggested 11): SRv6 End.X SID Sub-TLV : Refer to [Section 8.1](#).
- o Type TBD (suggested 12): SRv6 LAN End.X SID Sub-TLV : Refer to [Section 8.2](#).

12.4. OSPFv3 Locator LSA TLVs

This document proposes setting up of a new "OSPFv3 Locator LSA TLVs" registry that defines top-level TLVs for the OSPFv3 SRv6 Locator LSA to be added under the "OSPFv3 Parameters" registry. The initial code-points assignment is as below:

- o Type 0: Reserved.
- o Type 1: SRv6 Locator TLV : Refer to [Section 6.1](#).

Types in the range 2-32767 are allocated via IETF Review or IESG Approval [[RFC8126](#)].

Types in the range 32768-33023 are Reserved for Experimental Use; these will not be registered with IANA and MUST NOT be mentioned by RFCs.

Types in the range 33024-45055 are to be assigned on a First Come First Served (FCFS) basis.

Types in the range 45056-65535 are not to be assigned at this time. Before any assignments can be made in the 33024-65535 range, there MUST be an IETF specification that specifies IANA Considerations that cover the range being assigned.

[12.5.](#) OSPFv3 Locator LSA Sub-TLVs

This document proposes setting up of a new "OSPFv3 Locator LSA Sub-TLVs" registry that defines Sub-TLVs at any level of nesting for the SRv6 Locator TLVs to be added under the "OSPFv3 Parameters" registry. The initial code-points assignment is as below:

- o Type 0: Reserved.
- o Type 1: SRv6 End SID Sub-TLV : Refer to [Section 7](#).
- o Type 10: SRv6 SID Structure Sub-TLV : Refer to [Section 9](#).

Types in the range 2-9 and 11-32767 are allocated via IETF Review or IESG Approval [[RFC8126](#)].

Types in the range 32768-33023 are Reserved for Experimental Use; these will not be registered with IANA and MUST NOT be mentioned by RFCs.

Types in the range 33024-45055 are to be assigned on a First Come First Served (FCFS) basis.

Types in the range 45056-65535 are not to be assigned at this time. Before any assignments can be made in the 33024-65535 range, there MUST be an IETF specification that specifies IANA Considerations that cover the range being assigned.

[13.](#) Acknowledgements

The authors would like to thank Acee Lindem and Chenzichao for their review and comments on this document.

14. References

14.1. Normative References

- [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-07](#) (work in progress), July 2020.
- [I-D.ietf-lsr-isis-srv6-extensions]
Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extension to Support Segment Routing over IPv6 Dataplane", [draft-ietf-lsr-isis-srv6-extensions-08](#) (work in progress), April 2020.
- [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-17](#) (work in progress), August 2020.
- [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>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", [RFC 5340](#), DOI 10.17487/RFC5340, July 2008, <<https://www.rfc-editor.org/info/rfc5340>>.
- [RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", [RFC 7770](#), DOI 10.17487/RFC7770, February 2016, <<https://www.rfc-editor.org/info/rfc7770>>.
- [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>>.

- [RFC8362] Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", [RFC 8362](#), DOI 10.17487/RFC8362, April 2018, <<https://www.rfc-editor.org/info/rfc8362>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8476] Tantsura, J., Chunduri, U., Aldrin, S., and P. Psenak, "Signaling Maximum SID Depth (MSD) Using OSPF", [RFC 8476](#), DOI 10.17487/RFC8476, December 2018, <<https://www.rfc-editor.org/info/rfc8476>>.
- [RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", [RFC 8665](#), DOI 10.17487/RFC8665, December 2019, <<https://www.rfc-editor.org/info/rfc8665>>.
- [RFC8666] Psenak, P., Ed. and S. Previdi, Ed., "OSPFv3 Extensions for Segment Routing", [RFC 8666](#), DOI 10.17487/RFC8666, December 2019, <<https://www.rfc-editor.org/info/rfc8666>>.
- [RFC8754] Filsfils, 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>>.

14.2. 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-08](#) (work in progress), July 2020.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.
- [RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", [RFC 4552](#), DOI 10.17487/RFC4552, June 2006, <<https://www.rfc-editor.org/info/rfc4552>>.

[RFC7166] Bhatia, M., Manral, V., and A. Lindem, "Supporting Authentication Trailer for OSPFv3", [RFC 7166](#), DOI 10.17487/RFC7166, March 2014, <<https://www.rfc-editor.org/info/rfc7166>>.

Authors' Addresses

Zhenbin Li
Huawei Technologies

Email: lizhenbin@huawei.com

Zhibo Hu
Huawei Technologies

Email: huzhibo@huawei.com

Dean Cheng
Huawei Technologies

Email: dean.cheng@huawei.com

Ketan Talaulikar (editor)
Cisco Systems
India

Email: ketant@cisco.com

Peter Psenak
Cisco Systems
Slovakia

Email: ppsenak@cisco.com

