

Internet Engineering Task Force (IETF)
Request for Comments: 9513
Category: Standards Track
ISSN: 2070-1721

Z. Li
Z. Hu
Huawei Technologies
K. Talaulikar, Ed.
P. Psenak
Cisco Systems
December 2023

OSPFv3 Extensions for Segment Routing over IPv6 (SRv6)

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 an MPLS or IPv6 data plane. This document describes the OSPFv3 extensions required to support SR over the IPv6 data plane.

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/rfc9513>.

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 TLV
3. Advertisement of Supported Algorithms
4. Advertisement of 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
 - 5.1. SRv6 Flexible Algorithm
6. Advertisement of Anycast Property

- 7. SRv6 Locator LSA
 - 7.1. SRv6 Locator TLV
 - 7.2. SRv6 Locator Sub-TLVs
- 8. Advertisement of SRv6 End SIDs
- 9. Advertisement of SRv6 SIDs Associated with Adjacencies
 - 9.1. SRv6 End.X SID Sub-TLV
 - 9.2. SRv6 LAN End.X SID Sub-TLV
- 10. SRv6 SID Structure Sub-TLV
- 11. Advertising Endpoint Behaviors
- 12. Security Considerations
- 13. IANA Considerations
 - 13.1. OSPFv3 Router Information TLVs
 - 13.2. OSPFv3 LSA Function Codes
 - 13.3. OSPFv3 Prefix Options
 - 13.4. OSPFv3 SRv6 Capabilities TLV Flags
 - 13.5. OSPFv3 SRv6 End SID Sub-TLV Flags
 - 13.6. OSPFv3 SRv6 Adjacency SID Sub-TLV Flags
 - 13.7. OSPFv3 Extended-LSA Sub-TLVs
 - 13.8. OSPFv3 SRv6 Locator LSA TLVs
 - 13.9. OSPFv3 SRv6 Locator LSA Sub-TLVs
 - 13.10. OSPFv3 Extended-LSA Sub-TLVs
- 14. References
 - 14.1. Normative References
 - 14.2. Informative References
- Acknowledgements
- Authors' Addresses

1. Introduction

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

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

The network programming paradigm for SRv6 is specified in [RFC8986]. It describes how any behavior can be bound to a SID and how any network program can be expressed as a combination of SIDs. It also describes several well-known behaviors that can be bound to SRv6 SIDs.

This document specifies OSPFv3 extensions to support SRv6 capabilities as defined in [RFC8986], [RFC8754], and [RFC9259]. The extensions include advertisement of an OSPFv3 router's SRv6 capabilities, SRv6 Locators, and required SRv6 SIDs along with their supported Endpoint behaviors. Familiarity with [RFC8986] is necessary to understand the extensions specified in this document.

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

1. An SRv6 Capabilities TLV to advertise the SRv6 features and SRH operations supported by an OSPFv3 router.
2. Several sub-TLVs to advertise various SRv6 Maximum SID Depths.
3. An SRv6 Locator TLV using an SRv6 Locator Link State Advertisement (LSA) to advertise the SRv6 Locator -- a form of summary address for the IGP algorithm-specific SIDs instantiated on an OSPFv3 router.
4. TLVs and sub-TLVs to advertise the SRv6 SIDs instantiated on an OSPFv3 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 support for the SR Segment Endpoint Node [RFC8754] functionality along with its SRv6-related capabilities. This is an optional top-level TLV of the OSPFv3 Router Information LSA [RFC7770] that MUST be advertised by an SRv6-enabled router.

This TLV MUST 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 LSA. If the SRv6 Capabilities TLV appears in multiple OSPFv3 Router Information LSAs that have different flooding scopes, the TLV in the OSPFv3 Router Information LSA with the area-scoped flooding scope MUST be used. If the SRv6 Capabilities TLV appears in multiple OSPFv3 Router Information LSAs that have the same flooding scope, the TLV in the OSPFv3 Router Information LSA with the numerically smallest Link State ID MUST be used, and subsequent instances of the TLV MUST be ignored.

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

The format of the OSPFv3 SRv6 Capabilities TLV is shown below:

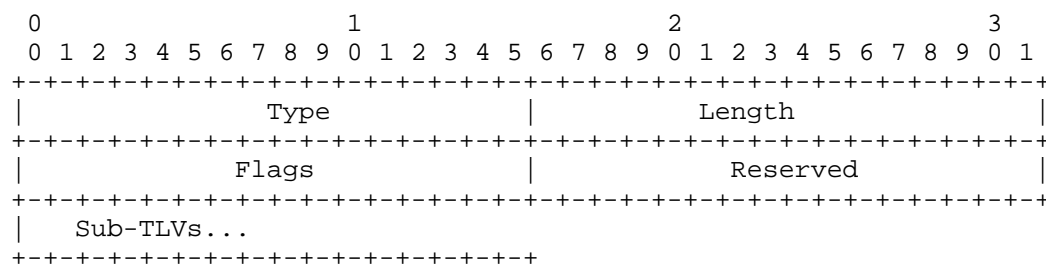


Figure 1: SRv6 Capabilities TLV


where:

Type: 2-octet field. The value for this type is 20.

Length: 2-octet field. The total length (in octets) of the value portion of the TLV, including nested sub-TLVs.

Reserved: 2-octet field. It MUST be set to 0 on transmission and MUST be ignored on receipt.

Flags: 2-octet field. The flags are defined as follows:



where:

O-flag: If set, then the router is capable of supporting the O-flag in the SRH flags, as specified in [RFC9259].

Other flags are not defined and are reserved for future use. They MUST be set to 0 on transmission and MUST be ignored on receipt.

The SRv6 Capabilities TLV may contain optional sub-TLVs. No sub-TLVs are defined in this specification.

3. Advertisement of Supported Algorithms

An SRv6-enabled OSPFv3 router advertises its algorithm support using the SR-Algorithm TLV defined in [RFC8665] and as described in [RFC8666].

4. Advertisement of Maximum SRv6 SID Depths

An SRv6-enabled router may have different capabilities and limits related to SRH processing. These need to be advertised to other OSPFv3 routers in the SRv6 domain.

[RFC8476] defines the means to advertise node- and 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 Router-Link 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 [RFC9352] for IS-IS are also used by OSPFv3. These MSD types are allocated in the "IGP MSD-Types" registry maintained by IANA and are shared by IS-IS and OSPF. They are described in the subsections below.

4.1. Maximum Segments Left MSD Type

The Maximum Segments Left MSD Type signals the maximum value of the Segments Left field of 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 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", which are flavors defined in [RFC8986]. If the advertised value is zero or no value is advertised, then the router cannot apply the 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 as part of the H.Encaps behavior as defined in [RFC8986]. 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. A non-zero SRH Max H.Encaps MSD indicates that the headend can insert an SRH with SIDs up to the advertised value.

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. These include, but are not limited to, End.DX6, End.DT4, End.DT46, End with USD, and End.X with USD as defined in [RFC8986]. 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 when the outer IPv6 header contains an SRH.

5. SRv6 SIDs and Reachability

An SRv6 SID is 128 bits and consists of locator, function, and argument parts as described in [RFC8986].

An OSPFv3 router is provisioned with algorithm-specific locators for each algorithm supported by that router. Each locator is a covering prefix for all SIDs provisioned on that router that have the matching algorithm.

Locators MUST be advertised within an SRv6 Locator TLV (see Section 7.1) using an SRv6 Locator LSA (see Section 7). The SRv6 Locator LSA is introduced instead of reusing the respective Extended Prefix LSAs [RFC8362] for a clear distinction between the two different types of reachability advertisements (viz., the base OSPFv3 prefix reachability advertisements and the SRv6 Locator reachability advertisements).

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 algorithm is supported by the receiving OSPFv3 router. Locators can be of different route types that map to existing OSPFv3 LSA types: Intra-Area, Inter-Area, External, and Not-So-Stubby Area (NSSA). The advertisement and propagation of the SRv6 Locator LSAs also follow the OSPFv3 [RFC5340] specifications for the respective LSA types. 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 OSPFv3 [RFC5340] specifications for the respective LSA types.

Locators associated with algorithms 0 and 1 (refer to Section 3.1.1 of [RFC8402]) SHOULD also be advertised using Extended LSA types with extended TLVs [RFC8362] so that routers that do not support SRv6 will install a forwarding entry for SRv6 traffic matching those locators. When operating in Extended LSA sparse-mode [RFC8362], these locators SHOULD also be advertised using Legacy LSAs [RFC5340].

When SRv6 Locators are also advertised as Intra-Area-Prefix-LSAs and/or E-Intra-Area-Prefix-LSAs, the SRv6 Locator MUST be considered as a prefix associated with the router, and the referenced LSA type MUST point to the Router LSA of the advertising router as specified in Section 4.4.3.9 of [RFC5340].

In cases where a locator advertisement is received both in a prefix reachability advertisement (i.e., via Legacy LSAs and/or Extended Prefix TLVs using Extended LSAs) and an SRv6 Locator TLV, the prefix reachability advertisement in the Legacy LSA or Extended LSA MUST be preferred over the advertisement in the SRv6 Locator TLV when installing entries in the forwarding plane. This prevents inconsistent forwarding entries between SRv6-capable and SRv6-incapable OSPFv3 routers. Such preference for prefix reachability advertisement does not have any impact on the rest of the data advertised in the SRv6 Locator TLV.

SRv6 SIDs are advertised as sub-TLVs in the SRv6 Locator TLV except for SRv6 End.X SIDs and LAN End.X SIDs, which are associated with a specific neighbor/link and are therefore advertised as sub-TLVs of the E-Router-Link TLV.

SRv6 SIDs received from other OSPFv3 routers 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 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 matching the locator prefix.

If the locator associated with SRv6 SID advertisements is the longest prefix match installed in the forwarding plane for those SIDs, this will 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 algorithm is the longest prefix match.
- * Another prefix advertised via Legacy or Extended LSA advertisement is the longest prefix match.

5.1. SRv6 Flexible Algorithm

[RFC9350] specifies IGP Flexible Algorithm mechanisms for OSPFv3. Section 14.2 of [RFC9350] explains SRv6 forwarding for Flexible Algorithms, and analogous procedures apply for supporting SRv6 Flexible Algorithms using OSPFv3. When the algorithm value that is advertised in the SRv6 Locator TLV (refer to Section 7.1) represents a Flexible Algorithm, the procedures described in Section 14.2 of [RFC9350] are followed for the programming of those specific SRv6 Locators.

Locators associated with Flexible Algorithms SHOULD NOT be advertised in the base OSPFv3 prefix reachability advertisements. Advertising the Flexible Algorithm locator in a regular prefix reachability advertisement would make it available for non-Flexible Algorithm forwarding (i.e., algorithm 0).

The procedures for OSPFv3 Flexible Algorithm for SR-MPLS, as specified in [RFC9350], also apply for SRv6; these procedures include a) ASBR reachability, b) inter-area, external, and NSSA prefix advertisements, and c) the use of those prefix advertisements in Flexible Algorithm route computation.

6. Advertisement of 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.

The AC-bit (value 0x80) in the OSPFv3 PrefixOptions field [RFC5340] is defined to advertise the anycast property:

```

      0  1  2  3  4  5  6  7
+---+---+---+---+---+---+---+
|AC|EL| N|DN| P| x|LA|NU|
+---+---+---+---+---+---+---+

```

Figure 2: OSPFv3 Prefix Options Field

When the prefix/SRv6 Locator is configured as anycast, the AC-bit

MUST be set. Otherwise, this flag MUST be clear.

The AC-bit MUST be preserved when re-advertising the prefix/SRv6 Locator across areas.

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

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

A prefix/SRv6 Locator that is advertised by a single node and without an AC-bit is considered node-specific.

All the nodes advertising the same anycast SRv6 Locator MUST instantiate the exact same set of SIDs under that anycast SRv6 Locator. Failure to do so may result in traffic being dropped or misrouted.

The PrefixOptions field is common to the prefix reachability advertisements (i.e., the base OSPFv3 prefix LSA types defined in [RFC5340], the OSPFv3 Extended Prefix TLV types defined in [RFC8362]), and the SRv6 Locator TLV advertisements specified in Section 7.1 of this document. When a router originates both the prefix reachability advertisement and the SRv6 Locator advertisement for a given prefix, the router SHOULD advertise the same PrefixOptions bits in both advertisements. In the case of any inconsistency between the PrefixOptions advertised in the SRv6 Locator and in the prefix reachability advertisements, the ones advertised in the prefix reachability advertisement MUST be preferred.

7. SRv6 Locator LSA

The SRv6 Locator LSA has a function code of 42. 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 the SRv6 Locator LSA is shown below:

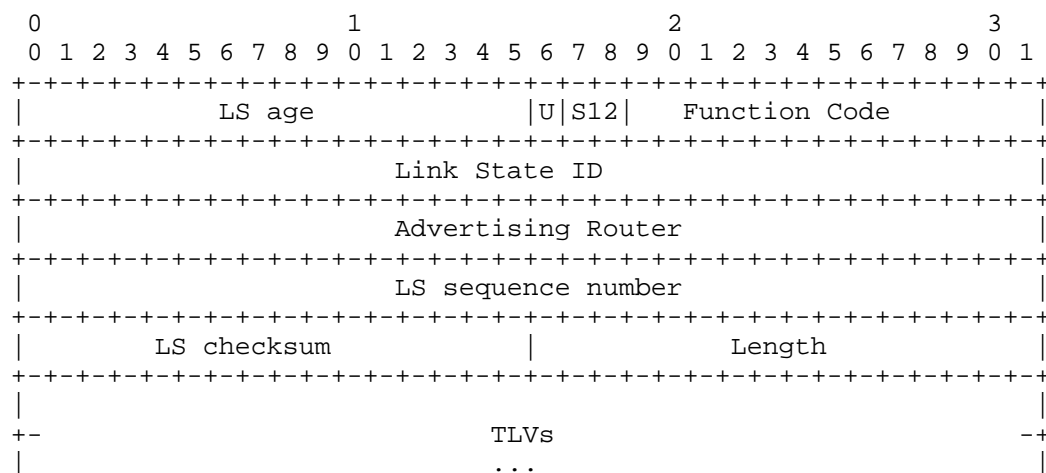


Figure 3: SRv6 Locator LSA

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:

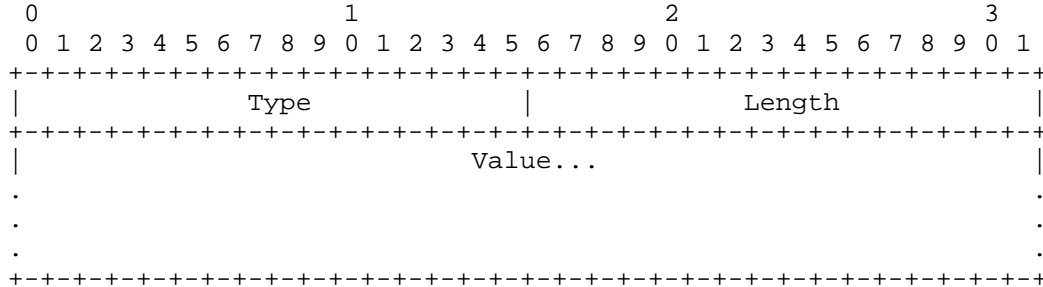


Figure 4: SRv6 Locator LSA TLV Format

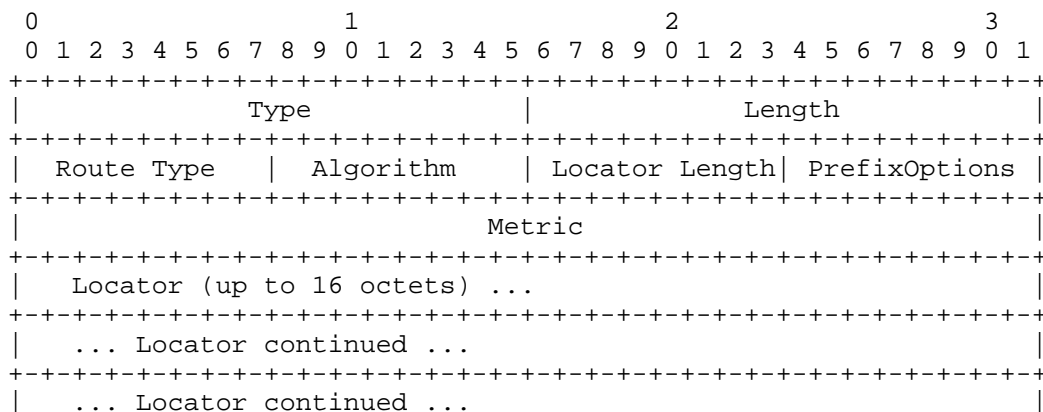
The Length field defines the length of the value portion in octets (thus, a TLV with no value portion would have a length of 0). The TLV is padded to 4-octet alignment; padding is not included in the Length field (so a 3-octet value would have a length of 3, but the total size of the TLV would be 8 octets). 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.

7.1. 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 has to 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 the SRv6 Locator TLV is shown below:



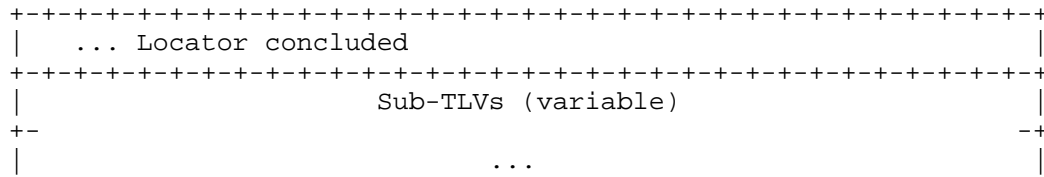


Figure 5: SRv6 Locator TLV

where:

Type: 2-octet field. The value for this type is 1.

Length: 2-octet field. The total length (in octets) of the value portion of the TLV, including nested sub-TLVs.

Route Type: 1-octet field. The type of the locator route. The only supported types are the ones listed below, and the SRv6 Locator TLV MUST be ignored on receipt of any other type.

- 1: Intra-Area
- 2: Inter-Area
- 3: AS External Type 1
- 4: AS External Type 2
- 5: NSSA External Type 1
- 6: NSSA External Type 2

Algorithm: 1-octet field. The algorithm associated with the SRv6 Locator. Algorithm values are defined in the "IGP Algorithm Types" registry [RFC8665].

Locator Length: 1-octet field. Specifies the length of the locator prefix as the number of locator bits from the range (1-128).

PrefixOptions: 1-octet field. Specifies the prefix options bits/flags as specified in [RFC5340] and further extended by [RFC8362] and Section 6 of this document.

Metric: 4-octet field. The metric value associated with the SRv6 Locator. The metric value of 0xFFFFFFFF MUST be considered as unreachable.

Locator: 1-16 octets. This field encodes the advertised SRv6 Locator as an IPv6 Prefix as specified in Appendix A.4.1 of [RFC5340].

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

7.2. SRv6 Locator Sub-TLVs

The following OSPFv3 Extended-LSA sub-TLVs corresponding to the Extended Prefix LSAs are also applicable for use as sub-TLVs of the SRv6 Locator TLV using code points as specified in Section 13.9:

- * IPv6-Forwarding-Address sub-TLV [RFC8362]
- * Route-Tag sub-TLV [RFC8362]
- * Prefix Source OSPF Router-ID sub-TLV [RFC9084]
- * Prefix Source Router Address sub-TLV [RFC9084]

8. 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 7). It is used to advertise the SRv6 SIDs belonging to the router along with their associated Endpoint behaviors. SIDs associated with adjacencies are advertised as described in Section 9. Every SRv6-enabled OSPFv3 router SHOULD advertise at least one SRv6 SID associated with an End behavior for itself as specified in [RFC8986], although it MAY omit doing so if that node is not going to be used as a Segment Endpoint (e.g., for TE or Topology Independent Loop-Free Alternate (TI-LFA)) by any SR Source Node.

SRv6 End SIDs inherit the algorithm from the parent locator. 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.

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 the SRv6 End SID sub-TLV is shown below

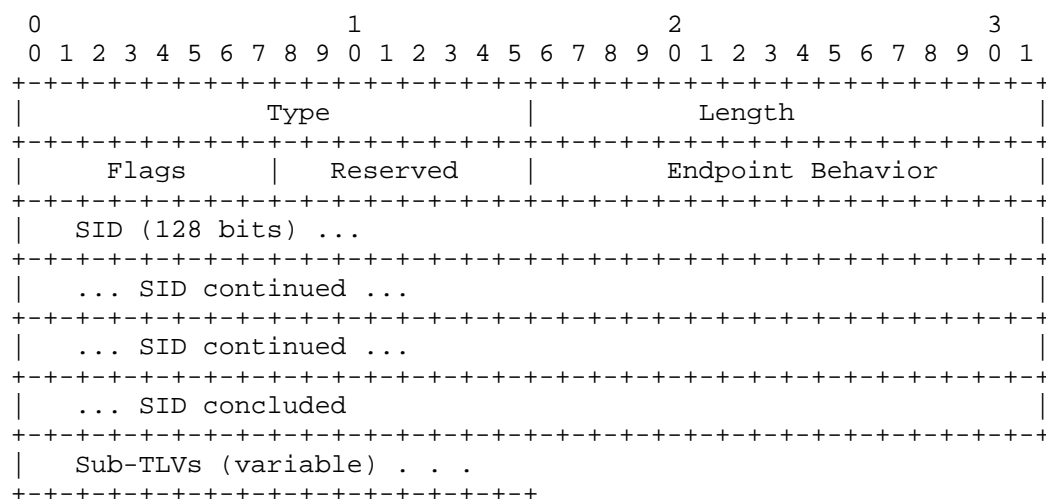


Figure 6: SRv6 End SID Sub-TLV

where:

Type: 2-octet field. The value for this type is 1.

Length: 2-octet field. The total length (in octets) of the value portion of the sub-TLV, including its nested sub-TLVs.

Flags: 1-octet field. Specifies the flags associated with the SID. No flags are currently defined, and this field MUST be set to 0 on transmission and MUST be ignored on receipt.

Reserved: 1-octet field. It MUST be set to 0 on transmission and MUST be ignored on receipt.

Endpoint Behavior: 2-octet field. The Endpoint behavior code point for this SRv6 SID as defined in [RFC8986]. Supported behavior values for this sub-TLV are defined in Section 11 of this document. Unsupported or unrecognized behavior values are ignored by the receiver.

SID: 16-octet field. This field encodes the advertised SRv6 SID.

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

9. Advertisement of SRv6 SIDs Associated with Adjacencies

The SRv6 Endpoint behaviors defined in [RFC8986] include certain behaviors that 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, which is an instruction to forward to a specific neighbor on a specific link. These SRv6 SIDs and others that are defined in [RFC8986], which are specific to links or adjacencies, need to be advertised to OSPFv3 routers within an area to steer SRv6 traffic over a specific link or adjacency.

Therefore, SRv6 SIDs that are specific to a particular neighbor, such as End.X, are not advertised as a sub-TLVs of the SRv6 Locator TLV. Instead, they are advertised via two different optional sub-TLVs of the E-Router-Link TLV defined in [RFC8362]:

SRv6 End.X SID sub-TLV: Used for OSPFv3 adjacencies over point-to-point or point-to-multipoint links and for the adjacency to the Designated Router (DR) over broadcast and Non-Broadcast-Multi-Access (NBMA) links.


SRv6 LAN End.X SID sub-TLV: Used 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 as a sub-TLV of 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 neighbors, although it MAY omit doing so if features like TE or TI-LFA that require End.X SID are not in use. A router MAY instantiate more than one SRv6 End.X SID 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 that is advertised by the same OSPFv3 router in an SRv6 Locator TLV. End.X SIDs that do not meet this requirement MUST be ignored. This ensures that the OSPFv3 router 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.

9.1. SRv6 End.X SID Sub-TLV

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



```

|   ... SID continued ...   |
+++++
|   ... SID concluded   |
+++++
|   Sub-TLVs (variable) . . .   |
+++++

```

Figure 7: SRv6 End.X SID Sub-TLV

where:

Type: 2-octet field. The value for this type is 31.

Length: 2-octet field. The total length (in octets) of the value portion of the sub-TLV, including its nested sub-TLVs.

Endpoint Behavior: 2-octet field. The Endpoint behavior code point for this SRv6 SID as defined in [RFC8986]. Supported behavior values for this sub-TLV are defined in Section 11 of this document. Unsupported or unrecognized behavior values are ignored by the receiver.

Flags: 1-octet field. The flags are defined as follows:

```

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.

Other flags are not defined and are reserved for future use. They MUST be set to 0 on transmission and MUST be ignored on receipt.

Reserved1: 1-octet field. It MUST be set to 0 on transmission and MUST be ignored on receipt.

Algorithm: 1-octet field. The algorithm associated with the SRv6 Locator from which the SID is allocated. Algorithm values are defined in the "IGP Algorithm Types" registry [RFC8665].

Weight: 1-octet field. Its value represents the weight of the End.X SID for load-balancing. The use of the weight is defined in [RFC8402].

Reserved2: 2-octet field. It MUST be set to 0 on transmission and MUST be ignored on receipt.

SID: 16-octet 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.

9.2. SRv6 LAN End.X SID Sub-TLV

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

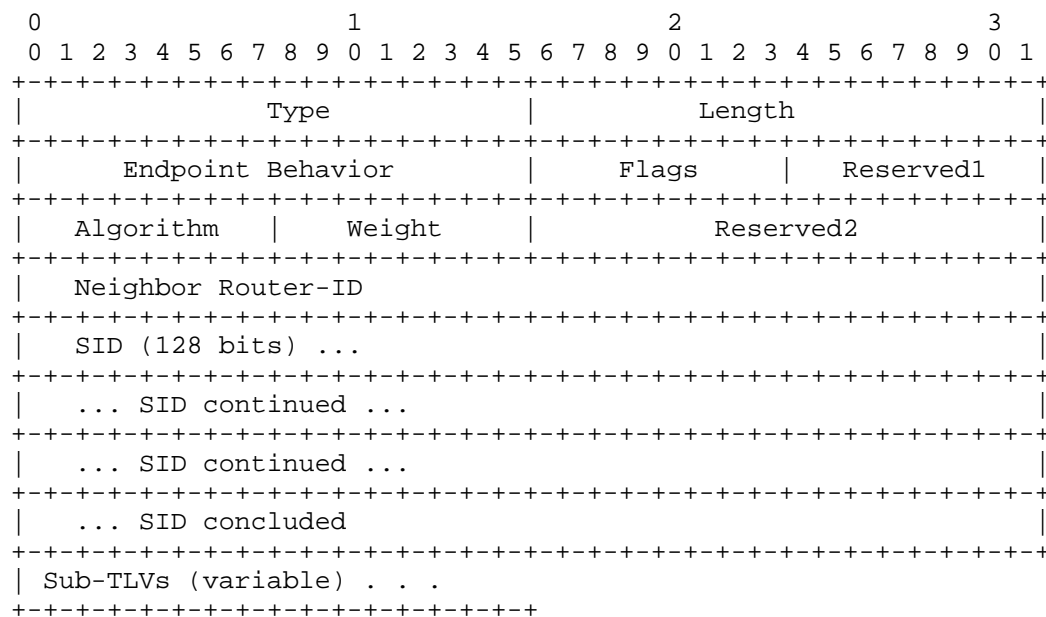


Figure 8: SRv6 LAN End.X SID Sub-TLV

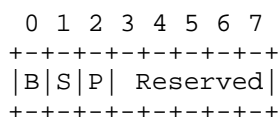
where:

Type: 2-octet field. The value for this type is 32.

Length: 2-octet field. The total length (in octets) of the value portion of the sub-TLV, including its nested sub-TLVs.

Endpoint Behavior: 2-octet field. The code point for the Endpoint behavior for this SRv6 SID as defined in Section 9.2 of [RFC8986].

Flags: 1-octet field. The flags are defined as follows:



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.

Other flags are not defined and are reserved for future use. They MUST be set to 0 on transmission and MUST be ignored on receipt.

Reserved1: 1-octet field. It MUST be set to 0 on transmission and MUST be ignored on receipt.

Algorithm: 1-octet field. The algorithm associated with the SRv6 Locator from which the SID is allocated. Algorithm values are defined in the "IGP Algorithm Types" registry [RFC8665].

Weight: 1-octet field. Its value represents the weight of the End.X SID for load balancing. The use of the weight is defined in

[RFC8402].

Reserved2: 2-octet field. It MUST be set to 0 on transmission and MUST be ignored on receipt.

Neighbor Router-ID: 4-octet field. It specifies the OSPFv3 Router-ID of the neighbor.

SID: 16-octet field. This field encodes the advertised SRv6 SID.

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

10. SRv6 SID Structure Sub-TLV

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

- * SRv6 End SID sub-TLV (refer to Section 8)
- * SRv6 End.X SID sub-TLV (refer to Section 9.1)
- * SRv6 LAN End.X SID sub-TLV (refer to Section 9.2)

The sub-TLV has the following format:

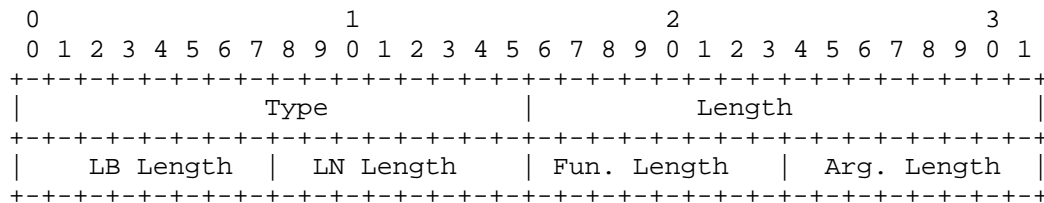


Figure 9: SRv6 SID Structure Sub-TLV

where:

Type: 2-octet field. The value for this type is 30.

Length: 2-octet field. The value MUST be 4.

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

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

Function Length: 1-octet field. SRv6 SID Function length in bits.

Argument Length: 1-octet field. SRv6 SID Argument length in bits.

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

The SRv6 SID Structure 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 OSPFv3. These can be learned by controllers via BGP-LS [RFC9514] and then monitored for conformance to the SRv6 SID allocation scheme chosen by the operator as described in Section 3.2 of [RFC8986].

- * Verification and the 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.

11. Advertising Endpoint Behaviors

Endpoint behaviors are defined in [RFC8986]. The code points for the Endpoint behaviors are defined in the "SRv6 Endpoint Behaviors" registry of [RFC8986]. This section lists the Endpoint behaviors and their code points, which MAY be advertised by OSPFv3 and the sub-TLVs in which each type MAY appear.

Endpoint Behavior	Endpoint Behavior Code Point	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.DT64	20	Y	N	N

Table 1: SRv6 Endpoint Behaviors in OSPFv3

12. Security Considerations

This document introduces extensions to the OSPFv3 protocol and, as such, does not affect existing security considerations for OSPFv3 as documented in [RFC5340]. [RFC7166] describes an alternative and improved authentication mechanism to IPsec for OSPFv3. The use of authentication is RECOMMENDED for OSPFv3 deployment.

Reception of a malformed TLV or sub-TLV SHOULD be counted and/or logged in a rate-limited manner for further analysis.

This document describes the OSPFv3 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 as well.

The advertisement of an incorrect MSD value may have negative consequences. See [RFC8476] 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. IANA Considerations

Per this document, IANA has made allocations in OSPF- and OSPFv3-related registries and created new registries, as detailed in the following subsections.

13.1. OSPF Router Information TLVs

IANA has allocated the following code point in the "OSPF Router Information (RI) TLVs" registry within the "Open Shortest Path First (OSPF) Parameters" registry group:

Value	TLV Name	Reference
20	SRv6 Capabilities	RFC 9513, Section 2

Table 2

13.2. OSPFv3 LSA Function Codes

IANA has allocated the following code point in the "OSPFv3 LSA Function Codes" registry within the "Open Shortest Path First v3 (OSPFv3) Parameters" registry group:

Value	LSA Function Code Name	Reference
42	SRv6 Locator LSA	RFC 9513, Section 7

Table 3

13.3. OSPFv3 Prefix Options

IANA has allocated the following code point in the "OSPFv3 Prefix Options (8 bits)" registry within the "Open Shortest Path First v3 (OSPFv3) Parameters" registry group:

Value	Description	Reference
0x80	AC-bit	RFC 9513, Section 6

Table 4

13.4. OSPFv3 SRv6 Capabilities TLV Flags

IANA has created a new subregistry named "OSPFv3 SRv6 Capabilities TLV Flags" within the "Open Shortest Path First v3 (OSPFv3) Parameters" registry group to control the assignment of bits 0 to 15 in the Flags field of the OSPFv3 SRv6 Capabilities TLV specified in this document. The registration procedure is "Standards Action" as defined in [RFC8126].

The following assignment has been made per this document:

Bit	Description	Reference
-----	-------------	-----------

1	O-flag	RFC 9513, Section 2	
+-----+	+-----+	+-----+	+-----+

Table 5

13.5. OSPFv3 SRv6 End SID Sub-TLV Flags

IANA has created a new subregistry named "OSPFv3 SRv6 End SID Sub-TLV Flags" within the "Open Shortest Path First v3 (OSPFv3) Parameters" registry group to control the assignment of bits 0 to 7 in the Flags field of the OSPFv3 SRv6 End SID sub-TLV specified in this document. The registration procedure is "Standards Action" as defined in [RFC8126].

No assignments are made by this document.

13.6. OSPFv3 SRv6 Adjacency SID Sub-TLV Flags

IANA has created a new subregistry named "OSPFv3 SRv6 Adjacency SID Sub-TLV Flags" within the "Open Shortest Path First v3 (OSPFv3) Parameters" registry group to control the assignment of bits 0 to 7 in the Flags field of the OSPFv3 SRv6 End.X SID and OSPFv3 SRv6 LAN End.X SID sub-TLVs specified in this document. The registration procedure is "Standards Action" as defined in [RFC8126].

The following assignments have been made per this document:

Bit	Description	Reference
0	B-flag	RFC 9513, Sections 9.1 and 9.2
1	S-flag	RFC 9513, Sections 9.1 and 9.2
2	P-flag	RFC 9513, Sections 9.1 and 9.2

Table 6

13.7. OSPFv3 Extended-LSA Sub-TLVs

IANA has allocated the following code points in the "OSPFv3 Extended-LSA Sub-TLVs" registry within the "Open Shortest Path First v3 (OSPFv3) Parameters" registry group:

Value	Description	L2BM	Reference
30	SRv6 SID Structure	Y	RFC 9513, Section 10
31	SRv6 End.X SID	Y	RFC 9513, Section 9.1
32	SRv6 LAN End.X SID	Y	RFC 9513, Section 9.2

Table 7

13.8. OSPFv3 SRv6 Locator LSA TLVs

IANA has created a new subregistry named "OSPFv3 SRv6 Locator LSA TLVs" within the "Open Shortest Path First v3 (OSPFv3) Parameters" registry group to define top-level TLVs for the OSPFv3 SRv6 Locator LSA. The initial assignments are below:

Value	Description	Reference
-------	-------------	-----------

0	Reserved	RFC 9513
1	SRv6 Locator	RFC 9513, Section 7.1

Table 8

Types in the range 0-32767 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 45056-65535 range, there MUST be an IETF specification that specifies IANA Considerations that cover the range being assigned.

13.9. OSPFv3 SRv6 Locator LSA Sub-TLVs

IANA has created a new subregistry named "OSPFv3 SRv6 Locator LSA Sub-TLVs" within the "Open Shortest Path First v3 (OSPFv3) Parameters" registry group to define sub-TLVs at any level of nesting for the SRv6 Locator LSA TLV. The initial assignment are below:

Value	Description	Reference
0	Reserved	RFC 9513
1	SRv6 End SID	RFC 9513, Section 8
2	IPv6-Forwarding-Address	[RFC8362]; RFC 9513, Section 7.2
3	Route-Tag	[RFC8362]; RFC 9513, Section 7.2
4	Prefix Source OSPF Router-ID	[RFC9084]; RFC 9513, Section 7.2
5	Prefix Source Router Address	[RFC9084]; RFC 9513, Section 7.2
10	SRv6 SID Structure	RFC 9513, Section 10

Table 9

Types in the range 0-32767 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 FCFS basis.

Types in the range 45056-65535 are not to be assigned at this time. Before any assignments can be made in the 45056-65535 range, there

MUST be an IETF specification that specifies IANA Considerations that cover the range being assigned.

The following note has been added to this registry to ensure that any document requesting allocations in this registry for sub-TLVs of any of the OSPFv3 SRv6 Locator TLVs checks if allocations are also applicable for the "OSPFv3 Extended-LSA Sub-TLVs" registry.

Note: Allocations made in this registry for sub-TLVs that are associated with OSPFv3 SRv6 Locator TLVs MUST be evaluated for their applicability as OSPFv3 Extended-LSA sub-TLVs, which are required to be allocated in the "OSPFv3 Extended-LSA Sub-TLVs" registry.

13.10. OSPFv3 Extended-LSA Sub-TLVs

IANA has added the following note to the "OSPFv3 Extended-LSA Sub-TLVs" registry within the "Open Shortest Path First v3 (OSPFv3) Parameters" registry group. The purpose of this note is to ensure that any document requesting allocations in this registry for sub-TLVs of any of the OSPFv3 Extended Prefix TLVs checks if allocations are also applicable for the "OSPFv3 SRv6 Locator LSA Sub-TLVs" registry defined in this document.

Note: Allocations made in this registry for sub-TLVs that are associated with OSPFv3 Extended TLVs related to prefix advertisements MUST be evaluated for their applicability as OSPFv3 SRv6 Locator sub-TLVs, which are required to be allocated in the "OSPFv3 SRv6 Locator LSA Sub-TLVs" registry.

14. References

14.1. Normative References

- [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>>.
- [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>>.
- [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>>.
- [RFC8986] Filsfils, 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>>.
- [RFC9084] Wang, A., Lindem, A., Dong, J., Psenak, P., and K. Talaulikar, Ed., "OSPF Prefix Originator Extensions", RFC 9084, DOI 10.17487/RFC9084, August 2021, <<https://www.rfc-editor.org/info/rfc9084>>.
- [RFC9259] Ali, Z., Filsfils, 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., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <<https://www.rfc-editor.org/info/rfc9350>>.
- [RFC9352] Psenak, P., Ed., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane", RFC 9352, DOI 10.17487/RFC9352, February 2023, <<https://www.rfc-editor.org/info/rfc9352>>.

14.2. Informative References

- [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>>.
- [RFC9514] Dawra, G., Filsfils, C., Talaulikar, K., Ed., Chen, M., Bernier, D., and B. Decraene, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing over IPv6 (SRv6)", RFC 9514, DOI 10.17487/RFC9514, December 2023, <<https://www.rfc-editor.org/info/rfc9514>>.

Acknowledgements

The authors would like to acknowledge the contributions of Dean Cheng in the early draft versions of this document. The authors would like to thank Ran Chen and Detao Zhao for their suggestions related to the extension of PrefixOptions for the signaling of the anycast property.

The authors would like to thank Chenzichao, Dirk Goethals, Baalajee S, Yingzhen Qu, Shraddha Hegde, Dhruv Dhody, Martin Vigoureux, and Reese Enghardt for their review and comments on this document. The authors would like to thank Acee Lindem for his detailed shepherd review and feedback for improvement of this document. The authors would like to thank John Scudder for his AD review and suggestions to improve this document.

Authors' Addresses

Zhenbin Li
Huawei Technologies
Email: lizhenbin@huawei.com

Zhibo Hu
Huawei Technologies
Email: huzhibo@huawei.com

Ketan Talaulikar (editor)
Cisco Systems
India
Email: ketant.ietf@gmail.com

Peter Psenak
Cisco Systems
Slovakia
Email: ppsenak@cisco.com