

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 6, 2018

Z. Li
Z. Hu
D. Cheng
Huawei Technologies
K. Talaulikar
P. Psenak
Cisco Systems
March 5, 2018

OSPFv3 Extensions for SRv6
draft-li-ospf-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 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", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <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 September 6, 2018.

Internet-Draft

OSPFv3 Extensions for SRv6

March 2018

Copyright Notice

Copyright (c) 2018 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
2.	OSPFv3 Extensions for SRv6	3
2.1.	SRv6-Capabilities TLV	3
2.1.1.	Maximum SL Sub-TLV	5
2.1.2.	Maximum End Pop SRH Sub-TLV	5
2.1.3.	Maximum T.Insert SRH Sub-TLV	6
2.1.4.	Maximum T.Encap SRH Sub-TLV	6
2.1.5.	Maximum End D SRH Sub-TLV	7
2.2.	SRv6 Node SID TLV	8
2.3.	SRv6 SIDs Associated with Adjacencies	10
2.3.1.	SRv6 SID Link Attribute Sub-TLV	11
2.3.2.	SRv6 SID LAN Link Attribute Sub-TLV	12
3.	Security Considerations	14
4.	IANA Considerations	14
4.1.	OSPF Parameters	14
4.2.	OSPFv3 Parameters	14
5.	Acknowledgements	15
6.	References	15
6.1.	Normative References	15
6.2.	Informative References	16
	Authors' Addresses	16

[1.](#) Introduction

Segment Routing (SR) architecture [[I-D.ietf-spring-segment-routing](#)] 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

[[I-D.ietf-6man-segment-routing-header](#)]. SRv6 refers to this SR instantiation on the IPv6 dataplane. The network programming paradigm for SRv6 is specified in [[I-D.filsfils-spring-srv6-network-programming](#)] which describes several well-known functions that can be bound to SRv6 SIDs.

This document proposes extensions to OSPFv3 in order to support SRv6 as defined in [[I-D.filsfils-spring-srv6-network-programming](#)] by signaling the SRv6 capabilities of the node and certain functions (e.g. END, END.X, etc.) that are instantiated on the SRv6 capable router.

At a high level, the extensions to OSPFv3 comprise of a SRv6 Capabilities TLV to advertise the support for SRv6 features supported by the router. Also included are extensions in the form of TLVs and sub-TLVs for advertising the SRv6 SIDs corresponding to functions related to the node (e.g. END) and those related to the adjacencies (e.g. END.X) for the SRv6 enabled OSPFv3 router.

[2.](#) OSPFv3 Extensions for SRv6

[2.1.](#) SRv6-Capabilities TLV

When Segment Routing (SR) is instantiated using the IPv6 data plane (SRv6), the list of segments is expressed using the segment routing header (SRH) as defined in [[I-D.ietf-6man-segment-routing-header](#)].

A router that supports SRv6 MUST be able to process the SRH as described in [[I-D.ietf-6man-segment-routing-header](#)], as well as apply function behaviors and flavors as described in [[I-D.filsfils-spring-srv6-network-programming](#)]. A 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.

The SRv6 Capabilities TLV is designed for an OSPFv3 router to

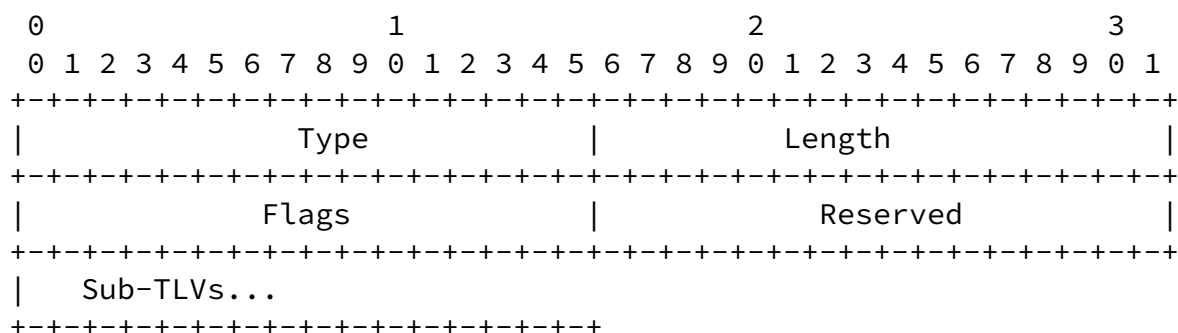
advertise its SRv6 support along with its related capabilities for SRv6 functionality. This is a new optional top level TLV of OSPFv3 Router Information LSA TLV LSA ([[RFC7770](#)]) which MUST be advertised by a SRv6 enabled router.

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. 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 by receiver.
- o Flags: 16 bit field. The following flags are defined:

```

      0                               1
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|E|O|                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

- * E-flag: If set, then router is able to apply "T.Encap" operation as specified in [\[I-D.filsfils-spring-srv6-network-programming\]](#)

- * O-flag: If set, then router is capable of supporting SRH O-bit Flags, as specified in [\[I-D.ietf-6man-segment-routing-header\]](#).

The following sections define the supported sub-TLVs.

[2.1.1.](#) Maximum SL Sub-TLV

The Maximum Segments Left Sub-TLV of the SRv6 Capabilities TLV specifies the maximum value of the "SL" field (refer to [\[I-D.ietf-6man-segment-routing-header\]](#)) in the SRH of a received packet before applying the function associated with a SID.

```

      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                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Max SL   |
+---+---+---+---+---+---+

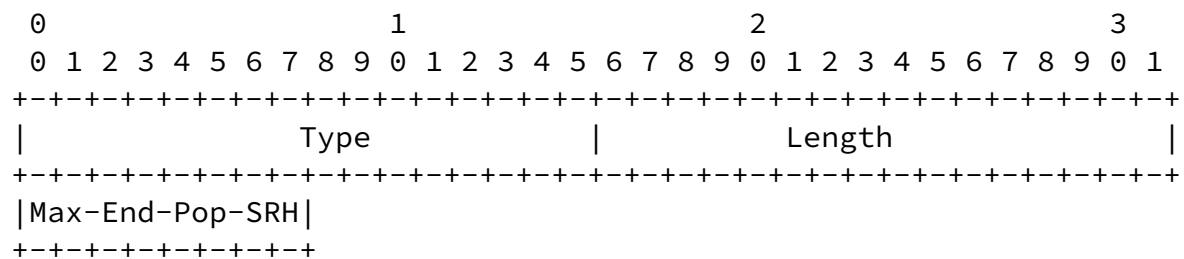
```

- o Type: 1

- If the sub-TLV is not advertised by an SRv6 capable router, then the value MUST be considered to be 0.

The Maximum End Pop SRH Sub-Sub-TLV specifies the maximum number of SIDs in the top SRH in an SRH stack to which the router can apply "PSP" or USP" flavors

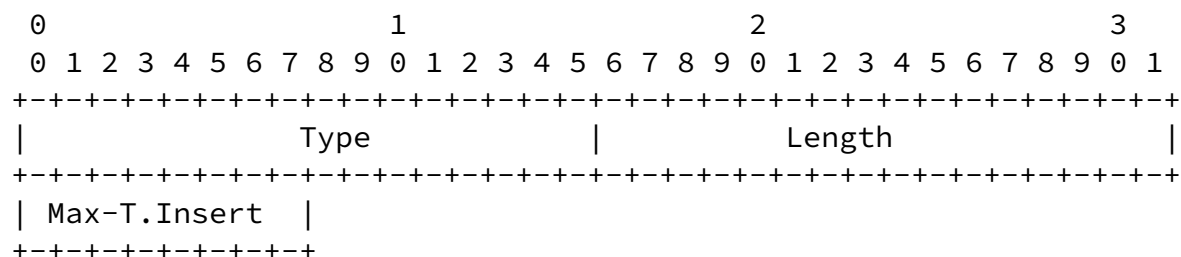
([\[I-D.filsfils-spring-srv6-network-programming\]](#)).



- If the value is 0 or the sub-TLV is not advertised by an SRv6 capable router, then it MUST be considered that the router cannot apply PSP or USP flavors.

2.1.3. Maximum T.Insert SRH Sub-TLV

The Maximum T.Insert SRH Sub-Sub-TLV specifies the maximum number of SIDs that can be inserted as part of the "T.insert" behavior([\[I-D.filsfils-spring-srv6-network-programming\]](#)).

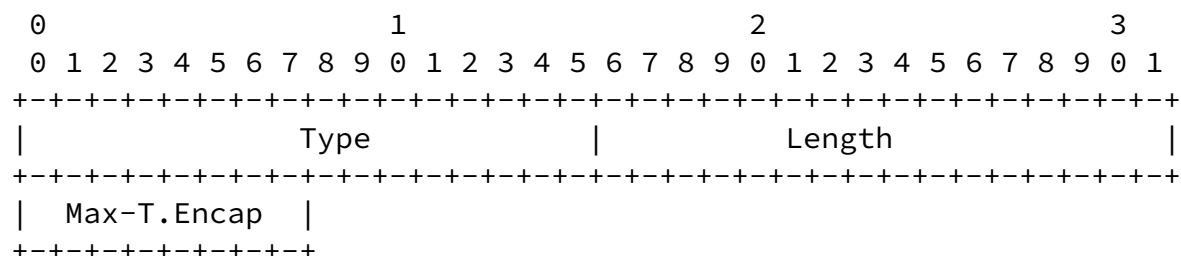


- o Type: 3
- o Length: 4 (including padding to 32 bit aligned boundary for OSPF TLVs)
- o Max-T.Insert Value: 1 octet
- o An 8 bit unsigned integer.

If the value is 0 or the sub-TLV is not advertised by an SRv6 capable router, then it MUST be considered that the router does not support any variation of the "T.insert" behavior.

[2.1.4.](#) Maximum T.Encap SRH Sub-TLV

The Maximum T.Encap SRH Sub-Sub-TLV specifies the maximum number of SIDs that can be included as part of the "T.Encap" behavior ([\[I-D.filsfils-spring-srv6-network-programming\]](#)).

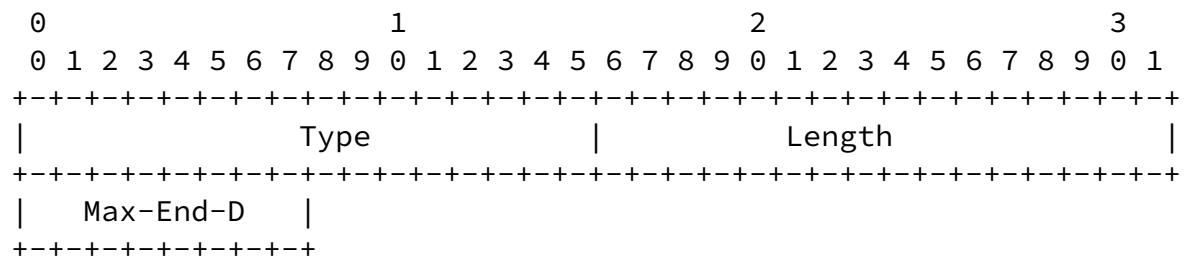


- o Type: 4
- o Length: 4 (including padding to 32 bit aligned boundary for OSPF TLVs)
- o Max-T.Encap Value: 1 octet
- o An 8 bit unsigned integer.

If this value is 0 or the sub-TLV is not advertised by an SRv6 capable router and the "E" flag is set in the associated SRv6 Capabilities sub-TLV, then it MUST be considered that the router can apply T.Encap by encapsulating the incoming packet in another IPv6 header without SRH the same way as IP-in-IP encapsulation is performed. If the "E" flag is clear, then this sub-TLV SHOULD NOT be advertised and MUST be ignored on receipt.

[2.1.5.](#) Maximum End D SRH Sub-TLV

The Maximum End D SRH sub-sub-TLV specifies the maximum number of SIDs in an SRH when applying "End.DX6" and "End.DT6" functions.



- o Type: 5
- o Length: 4 (including padding to 32 bit aligned boundary for OSPF TLVs)
- o Max End D Value: 1 octet
- o An 8 bit unsigned integer.

If this value is zero or the sub-TLV is not advertised by an SRv6

capable router, then it MUST be considered that the router cannot apply "End.DX6" or "End.DT6" functions if the extension header right underneath the outer IPv6 header is an SRH.

2.2. SRv6 Node SID TLV

An OSPFv3 SRv6 enabled router may have multiple SRv6 SIDs instantiated in its "My Local SID Table" (refer [\[I-D.filsfils-spring-srv6-network-programming\]](#)). OSPFv3 needs to advertise the SRv6 SIDs associated with the node and its functions (e.g. END, END.T, etc.) as specified in [\[I-D.filsfils-spring-srv6-network-programming\]](#) so that other routers in the area learn the SRv6 SIDs that can be used to program SRv6 paths through this node.

SRv6 Node SID TLV is a new optional top-level TLV of OSPFv3 Router Information LSA ([\[RFC7770\]](#)) and is used to advertise the SRv6 SIDs belonging to the node along with their associated functions. Every SRv6 enabled OSPFv3 router SHOULD advertise at least one SRv6 SID associated with END function for its node as specified in [\[I-D.filsfils-spring-srv6-network-programming\]](#).

The router MAY advertise multiple instances of the SRv6 Node SID TLV in its OSPFv3 Router Information LSA - one for each of the SRv6 SIDs to be advertised. It is RECOMMENDED that the TLVs are ordered by increasing values of the SRv6 SIDs within a single instance of the OSPFv3 Router LSA. When multiple instances of the OSPFv3 Router Information LSA are necessary to accomodate a large number of SRv6 SIDs, it is RECOMMENDED that the SRv6 Node SID TLVs are ordered by increasing values of the SRv6 SIDs across increasing instance numbers of the OSPFv3 Router Information LSA.

When multiple SRv6 Node SID TLVs are received from a given router for the same SRv6 SID value, the receiver MUST use the first occurrence of the TLV in the OSPFV3 Router Information Opaque LSA. If the SRv6 Node SID TLV for the same SRv6 SID value 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 Node SID TLV for the same SRv6 SID value 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 Node SID TLV advertisement, area-scoped flooding is REQUIRED.

The format of OSPFv3 SRv6 Node SID TLV is shown below

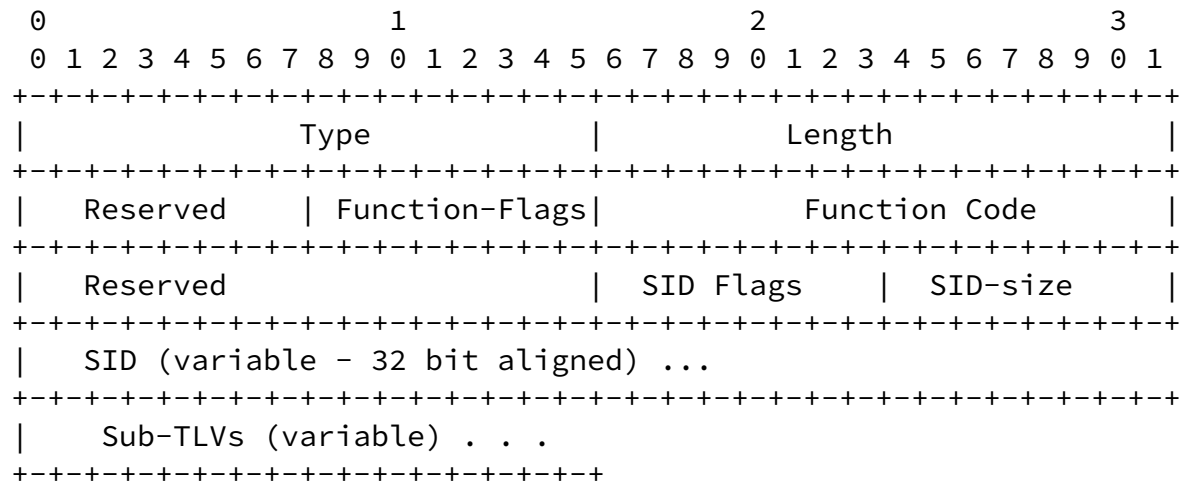


Figure 1: SRv6 SID Node TLV

Where:

Type: 16 bit field. TBD

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

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

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

Function Code: 16 bit field. The function code point for this SRv6 SID as defined in [\[I-D.filsfils-spring-srv6-network-programming\]](#).

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

SID Flags: 8 bit field which define the flags associated with the SID

```

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

```

Figure 2

- * D bit (0x1) : When the SID is leaked from OSPFv3 backbone area to other areas, the D bit MUST be set. Otherwise, this bit MUST be clear. SIDs with the D bit set MUST NOT be leaked to OSPFv3 backbone area from others. This is to prevent looping.
- * Other flags are not defined and SHOULD be set to 0 and MUST be ignored on receipt.

SID Size : 8 bit field. Number of bits in the SID field.

SID : 1-16 octets. This field encodes the advertised SRv6 SID. The "SID-size" field can have the values 1-128 and indicates the number of bits in the SID. The SRv6 SID is encoded in the minimal number of 32-bit aligned space for the given number of bits.

Sub-TLVs : currently none defined. Used to advertise sub-TLVs that provide additional attributes for the given SRv6 SID.

2.3. SRv6 SIDs Associated with Adjacencies

The SRv6 functions are defined in [\[I-D.filsfils-spring-srv6-network-programming\]](#) include certain functions which are specific to links or adjacencies. The most basic of this which is critical for link state routing protocols like OSPFv3 is the END.X function that is an instruction to forward to a specific neighbor on a specific link. These END.X SRv6 SIDs are instantiated by SRv6 enabled OSPFv3 router for all its adjacencies and installed in the local node's "My Local SID Table". These SRv6 SIDs along with others that are defined in [\[I-D.filsfils-spring-srv6-network-programming\]](#) which are specific to links or adjacencies need to be advertised by OSPFv3 so that this information is available with all routers in the domain to influence

the packet path via these specific functions over the specified adjacencies.


The advertising of SRv6 SIDs and their functions that are specific to a particular neighbor are done via two different optional sub-TLVs of the Router-Link TLV as defined in [[I-D.ietf-ospf-ospfv3-lsa-extend](#)] as follows:

- o SRv6 SID Link Attribute Sub-TLV: for OSPFv3 adjacency 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 SID LAN Link Attribute Sub-TLV: for OSPFv3 adjacency 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 multiple instances of this TLV for each neighbor to be explicitly advertised under the Router-Link TLV for the same link.

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

2.3.1. SRv6 SID Link Attribute Sub-TLV

The format of the SRv6 SID Link Attribute Sub-TLV is shown below



Reserved	SID Flags	SID-size
SID (variable - 32 bit aligned) ...		
Sub-TLVs (variable) . . .		

Where:

Type is TBD

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

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

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

Function Code: 16 bit field. The function code point for this SRv6 SID as defined in [\[I-D.filsfils-spring-srv6-network-programming\]](#).

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

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.

SID-size: Number of bits in the SID field.

SID: 1-16 octets. This field encodes the advertised SRv6 SID. The "SID-size" field can have the values 1-128 and indicates the number of bits in the SID. The SRv6 SID is encoded in the minimal number of 32-bit aligned space for the given number of bits.

- o Type: TBD
- o Length: 16 bit value. Variable
- o Reserved : 8 bit field. Should be set to 0 and MUST be ignored on receipt.
- o Function Flags: 8 bit field which define the flags associated with the function. No flags are currently defined and SHOULD be set to 0 and MUST be ignored on receipt.
- o Function Code: 16 bit field. The function code point for this SRv6 SID as defined in [\[I-D.filsfils-spring-srv6-network-programming\]](#).
- o Reserved : 16 bit field. Should be set to 0 and MUST be ignored on receipt.
- 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 SID Size : 8 bit field. Number of bits in the SID field.
- o SID : 1-16 octets. This field encodes the advertised SRv6 SID. The "SID-size" field can have the values 1-128 and indicates the number of bits in the SID. The SRv6 SID is encoded in the minimal number of 32-bit aligned space for the given number of bits.

- o Neighbor ID : 4 octets of OSPFv3 Router-id of the neighbor
- o Sub-TLVs : currently none defined. Used to advertise sub-TLVs that provide additional attributes for the given SRv6 SID.

3. Security Considerations

Existing security extensions as described in [\[RFC5340\]](#) and [\[I-D.ietf-ospf-ospfv3-lsa-extend\]](#) 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 [\[RFC4552\]](#) 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.

[4.](#) IANA Considerations

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

[4.1.](#) OSPF Parameters

This document proposes the following new code points in the OSPF Router Information (RI) TLVs registry for OSPFv3 Extensions in order to support SRv6:

1. Type TBD: SRv6-Capabilities TLV: Refer to [Section 2.1](#).
2. Type TBD: SRv6 Node SID TLV : Refer to [Section 2.2](#).

[4.2.](#) OSPFv3 Parameters

This document proposes the following new code points in the OSPFv3 Extend-LSA Sub-TLV registry for OSPFv3 Extensions in order to support SRv6:

1. Type TBD: SRv6 SID Link Attribute Sub-TLV : Refer to [Section 2.3.1](#).

2. Type TBD: SRv6 SID LAN Link Attribute Sub-TLV : Refer to [Section 2.3.2](#).

[5.](#) Acknowledgements

6. References

6.1. Normative References

- [I-D.filsfils-spring-srv6-network-programming]
Filsfils, C., Leddy, J., daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d., Steinberg, D., Raszuk, R., Matsushima, S., Lebrun, D., Decraene, B., Peirens, B., Salsano, S., Naik, G., Elmalky, H., Jonnalagadda, P., Sharif, M., Ayyangar, A., Mynam, S., Henderickx, W., Bashandy, A., Raza, K., Dukes, D., Clad, F., and P. Camarillo, "SRv6 Network Programming", [draft-filsfils-spring-srv6-network-programming-03](#) (work in progress), December 2017.
- [I-D.ietf-ospf-ospfv3-lsa-extend]
Lindem, A., Roy, A., Goethals, D., Vallem, V., and F. Baker, "OSPFv3 LSA Extendibility", [draft-ietf-ospf-ospfv3-lsa-extend-23](#) (work in progress), January 2018.
- [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.
- [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>>.
- [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>>.
- [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>>.

[6.2.](#) Informative References

- [I-D.ietf-6man-segment-routing-header]
Previdi, S., Filsfils, C., Raza, K., Dukes, D., Leddy, J., Field, B., daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d., Matsushima, S., Leung, I., Linkova, J., Aries, E., Kosugi, T., Vyncke, E., Lebrun, D., Steinberg, D., and R. Raszuk, "IPv6 Segment Routing Header (SRH)", [draft-ietf-6man-segment-routing-header-08](#) (work in progress), January 2018.

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
Cisco Systems
India

Email: ketant@cisco.com

Peter Psenak
Cisco Systems
Slovakia

Email: ppsenak@cisco.com

Li, et al.

Expires September 6, 2018

[Page 16]