

LSR Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 25, 2018

U. Chunduri
Y. Qu
Huawei USA
J. Tantsura
Nuage Networks
March 24, 2018

Usage of Non Shortest Path Forwarding (NSPF) IDs in OSPF
draft-ct-ospf-nspf-id-for-sr-paths-00

Abstract

This document specifies the advertisement of Non Shortest Path Forwarding Identifier (NSPF ID) TLV and the computation procedures for the same in OSPFv2 and OSPFv3 protocols. NSPF ID allows to simplify the data plane path description of data traffic in SR deployments. This helps to mitigate the MTU issues that are caused by additional SR overhead of the packet and allows traffic statistics.

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 [RFC2119](#) [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 25, 2018.

Internet-Draft Usage of Non Shortest Path Forwarding (NSPF) IDs 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
1.1.	Acronyms	3
2.	OSPF NSPF ID TLV	4
2.1.	Flags	5
2.2.	NSPF-ID Fields	5
2.3.	NSP sub-TLVs	6
2.4.	Non-NSP sub-TLVs	7
3.	OSPFv3 NSPF ID TLV	7
3.1.	OSPFv3 NSPF-ID Fields	9
3.2.	OSPFv3 NSP sub-TLVs	10
3.3.	OSPFv3 Non-NSP sub-TLVs	11
4.	Other Considerations	11
5.	Acknowledgements	11
6.	IANA Considerations	11
7.	Security Considerations	12
8.	References	12
8.1.	Normative References	12
8.2.	Informative References	12
	Authors' Addresses	14

[1.](#) Introduction

In a network implementing source routing, packets may be transported through the use of segment identifiers (SIDs), where a SID uniquely identifies a segment as defined in [[I-D.ietf-spring-segment-routing](#)]. In SR-MPLS, a segment is encoded as a label and an ordered list of

segments is encoded as a stack of labels. In SRv6, a segment is encoded as an IPv6 address, with a new type of IPv6 routing header called SRH. An ordered list of segments is encoded as an ordered list of IPv6 addresses in SRH [[I-D.ietf-6man-segment-routing-header](#)].

Internet-Draft Usage of Non Shortest Path Forwarding (NSPF) IDs March 2018

A segment may include one or more nodes, unidirectional adjecencies between two nodes or service instruction by a particular node in the network. A Non Shortest Path (NSP) could be a Traffic Engineered (TE) path or an explicitly provisioned FRR path or a service chained path. NSP can be described using list of segments in SR. However, this creates a problem of having a relatively large stack imposed on the data packet. A path that is encoded with SIDs can be a loose or strict path. In a strict path all the nodes/links on the path are encoded as SIDs, with the expense of number of total SIDs in the stack.

The issues caused by the large SID depth, and existing methods for mitigation are introduced in [[I-D.ct-isis-nspf-id-for-sr-paths](#)] [section 1.1](#) and 1.2. To mitigate the these issues, and also to facilitate forwarding plane a mechanism to identify the SR path with a corresponding data plane identifier for accounting of traffic for SR paths, this draft proposes a new OSPFv2 TLV ([Section 2](#)), OSPFv3 TLV ([Section 3](#)) to advertise the NSPs with Non Shortest Path Forwarding IDentifier (NSPF ID).

With corresponding data plane, [Section 3](#) mechanism as in [[I-D.ct-isis-nspf-id-for-sr-paths](#)], reduces the SID stack in the data plane with a single NSPF ID.

[1.1](#). Acronyms

EL	- Entropy Label
ELI	- Entropy Label Indicator
MPLS	- Multi Protocol Label Switching
MSD	- Maximum SID Depth
MTU	- Maximum Transferrable Unit

- NSP - Non Shortest Path
- SID - Segment Identifier
- SPF - Shortest Path First
- SR - Segment Routing
- SRH - Segment Routing Header
- SR-MPLS - Segment Routing with MPLS data plane

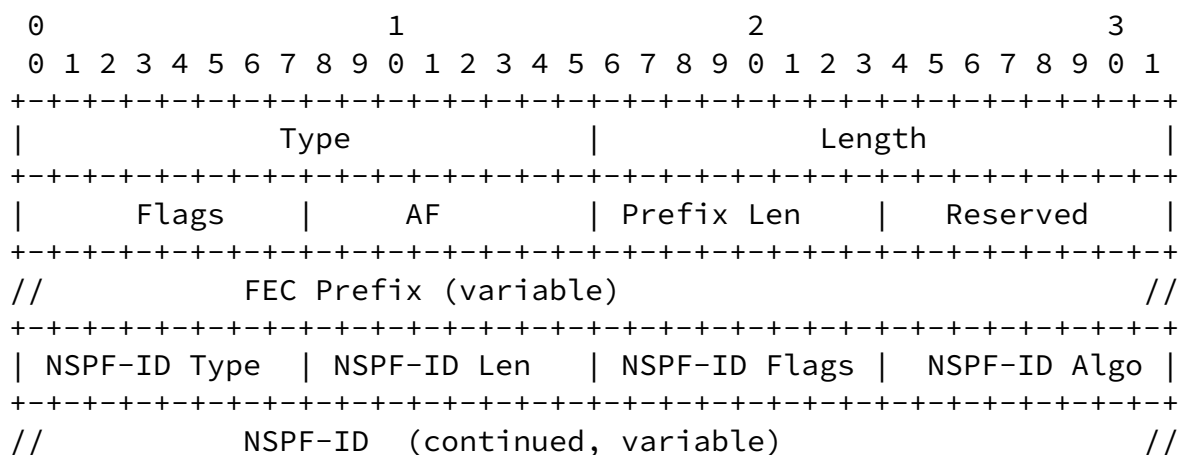
Internet-Draft Usage of Non Shortest Path Forwarding (NSPF) IDs March 2018

- SRv6 - Segment Routing with Ipv6 data plane with SRH
- SRH - IPv6 Segment Routing Header
- TE - Traffic Engineering

2. OSPF NSPF ID TLV

Extended Prefix Opaque LSAs defined in [\[RFC7684\]](#) are used for advertisements of NSPF ID TLV. Multiple OSPF NSPF ID TLVs MAY be advertised in each OSPF Extended Prefix Opaque LSA, but all TLVs included in a single OSPF Extended Prefix Opaque LSA MUST have the same flooding scope.

The NSPF-ID TLV has Type TBD (suggested value xxx), and has the following format:



```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|No.of NSP-STs |   NSP sub-TLVs (Variable)                                     //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|No.of Other-STs| Non-NSP sub-TLVs(variable)                                //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 1: NSPF ID TLV Format

Type - TBD (IANA) from OSPF Extended Prefix Opaque LSA registry.

Length - Total length of the value field in bytes (variable).

Reserved - 1 Octet reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.

Flags - Flags for this TLV are described in [Section 2.1](#).

AF - Address family for the prefix. Currently, the only supported value is 0 for IPv4 unicast. The inclusion of address family in this TLV allows for future extension.

Prefix Len - contains the length of the prefix in bits.

FEC Prefix - represents the Forwarding Equivalence Class at the tail-end of the advertised NSP. The "FEC Prefix" corresponds to a routable prefix of the originating node. Value of this field MUST be 4 octets for IPv4 "FEC Prefix".

[2.1](#). Flags

Flags: 1 octet field of NSPD ID TLV has following flags defined:

```

NSPF ID Flags Format
 0  1  2  3  4  5  6  7
+---+---+---+---+---+---+---+---+
|IA|  Rsrvd                    |
+---+---+---+---+---+---+---+---+

```

w=Where:

IA-Flag: Inter-Area flag. If set, advertisement is of inter-area type. An ABR that is advertising the OSPF NSPF ID TLV between areas MUST set this bit.

Rsrvd - reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.

2.2. NSPF-ID Fields

This represents the actual data plane identifier in the packet and could be of any data plane as defined in type field. Both "FEC Prefix" and NSPF-ID MUST belong to a same node in the network.

1. NSPF-ID Type: This is a new registry (TBD IANA) for this TLV and the defined types are as follows. Type: 1 - MPLS SID/Label Type: 2 Native IPv4 Address
2. NSPF-ID Len: Length of the NSPF Identifier field in octets and this depends on the NSPF-ID type. See NSPF-ID below for the length of this field and other considerations.
3. NSPF-ID Flags: 1 Octet field for NSPF-ID flags. Some of the bits could be NSPF-ID type specific and each new type MUST define the flags applicable to the NSPF-ID type. For NSPF-ID Type 1, the flags are same as definition in

[[I-D.ietf-ospf-segment-routing-extensions](#)]. Undefined flags for each NSPF-ID type MUST be considered as reserved. Reserved flag bits in each NSPF-ID type specific flags MUST be reset on transmission and ignored on receive.

4. NSPF-ID Algo: 1 octet value represents the SPF algorithm. Algorithm registry is as defined in [[I-D.ietf-ospf-segment-routing-extensions](#)].
5. NSPF-ID: This is the NSP forwarding identifier that would be on the data packet. The value of this field is variable and it depends on the NSPF-ID Type. For Type 1, this is and MPLS SID/Label. For Type 2 this is a 4-byte IPv4 address. For NSPF-ID Type 2, if the NSPF-ID Len is set to 0, then FEC Prefix would also become the NSPF-ID. In the case when NSPF-ID Len is 0,

NSPF-ID Type is 2, then FEC Prefix length MUST be a 4-byte IPv4 address.

6. No.of NSP-STs: Total number of the NSP sub-TLVs are defined with this 1-octet field. The value MUST NOT be zero.

[2.3.](#) NSP sub-TLVs

A new sub-TLV registry is created (TBD IANA) called NSP sub-TLVs. These are used to describe the path in the form of set of contiguous and ordered sub-TLVs, with first sub-TLV representing the top of the stack or first segment. These set of ordered TLVs can have both topological SIDs and non-topological SIDs (e.g., service segments).

Type 1: SID/Label sub-TLV as defined in [[I-D.ietf-ospf-segment-routing-extensions](#)]. Only Type is defined and Length/Value fields are per section 2.1 of the referenced document.

Type 2: Prefix SID sub-TLV as defined in [[I-D.ietf-ospf-segment-routing-extensions](#)]. Only Type is defined and Length/Value fields are per [section 5](#) of the referenced document.

Type 3: Adjacency SID sub-TLV as defined in [[I-D.ietf-ospf-segment-routing-extensions](#)]. Only Type is defined and Length/Value fields are per [section 6](#) of the referenced document.

Type 4: Length 4 bytes, value is 4 bytes IPv4 address encoded similar to IPv4 FEC Prefix described above.

[2.4.](#) Non-NSP sub-TLVs

NSPF ID TLV also defines a new sub-TLV registry (TBD IANA) for defining extensible set of sub-TLVs other than describing the path sub-TLVs. Total number of the path sub-TLVs to describe the path are defined in 1-octet field "No.of Other-STs" just before the Non-NSP sub-TLVs. This field serves as a demarcation for set of ordered NSP sub-TLVs and Non-NSP sub-TLVs.

Type 1: Length 0 No value field. Specifies a counter to count number of packets forwarded on this NSPF-ID.

Type 2: Length 0 No value field. Specifies a counter to count number of bytes forwarded on this NSPF-ID specified in the network header (e.g. IPv4, IPv6).

Type 3: Length 4 bytes, and Value is metric of this path represented through the NSPF-ID. Different nodes can advertise the same NSPF-ID for the same FEC-Prefix with a different set of NSP sub-TLVs and the receiving node MUST consider the lowest metric value (TBD more, what happens when metric is same for two different set of NSP sub-TLVs).

3. OSPFv3 NSPF ID TLV

The OSPFv3 NSPF ID TLV s a top level TLV of the following LSAs defined in [[I-D.ietf-ospf-ospfv3-lsa-extend](#)].

E-Intra-Area-Prefix-LSA

E-Inter-Area-Prefix-LSA

E-AS-External-LSA

E-Type-7-LSA

Multiple OSPFv3 NSPF ID TLVs MAY be advertised in each LSA mentioned above. The OSPFv3 NSPF ID TLV has the following format:

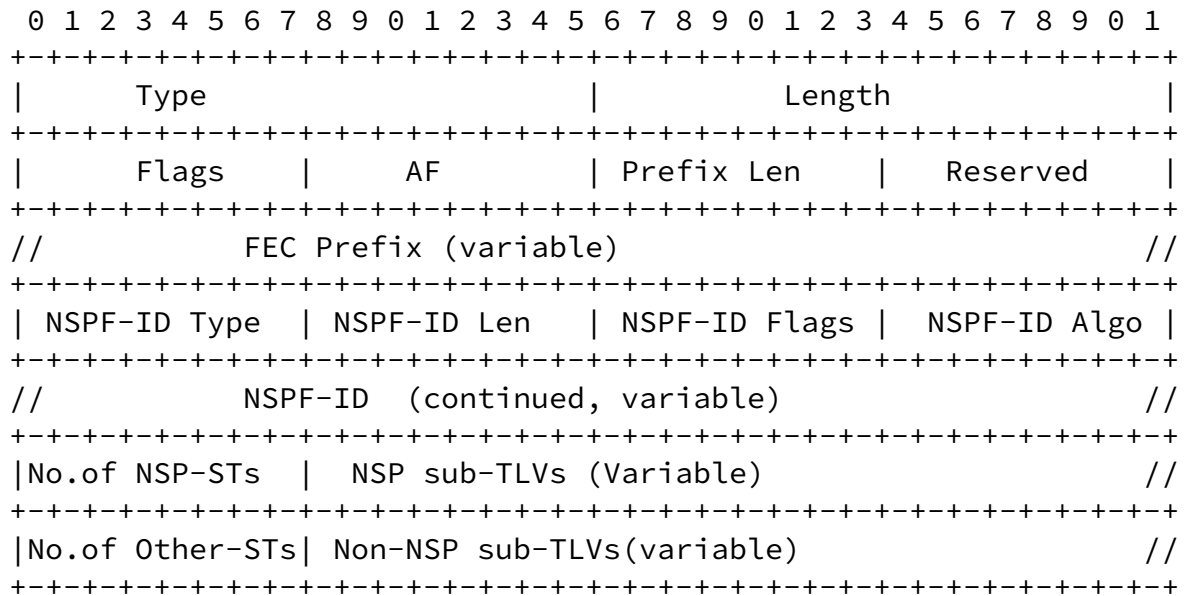


Figure 2: OSPFv3 NSPF ID TLV Format

where:

Type: TBD

Length: Variable, in octets, depends on Sub-TLVs.

prefix length: Length of prefix in bytes.


AF: Address family for the prefix.

AF: 0 - IPv4 unicast

AF: 1 - IPv6 unicast

Flags: Single octet field. The following flags are defined:

NSPF ID Flags Format



IA-Flag: Inter-Area flag. If set, advertisement is of inter-area type. An ABR that is advertising the OSPF NSPF ID TLV between areas MUST set this bit.

[[I-D.ietf-ospf-ospfv3-segment-routing-extensions](#)]

Rsrvd - reserved bits for future use. Reserved bits MUST be reset on transmission and ignored on receive.

FEC Prefix - represents the Forwarding Equivalence Class at the tail-end of the advertised NSP. The "FEC Prefix" corresponds to a routable prefix of the originating node. Value of this field MUST be 4 octets for IPv4 "FEC Prefix". Value of this field MUST be 16 octets for IPv6 "FEC Prefix".

[3.1.](#) OSPFv3 NSPF-ID Fields

This represents the actual data plane identifier in the packet and could be of any data plane as defined in type field. Both "FEC Prefix" and NSPF-ID MUST belong to a same node in the network.

1. NSPF-ID Type: This is a new registry (TBD IANA) for this TLV and the defined types are as follows. Type: 1 - MPLS SID/Label Type: 2 Native IPv4 Address Type: 3 Native IPv6 Address Type 4: IPv6 SID in SRv6 with SRH
2. NSPF-ID Len: Length of the NSPF Identifier field in octets and this depends on the NSPF-ID type. See NSPF-ID below for the length of this field and other considerations.
3. NSPF-ID Flags: 1 Octet field for NSPF-ID flags. Some of the bits could be NSPF-ID type specific and each new type MUST define the flags applicable to the NSPF-ID type. For NSPF-ID Type 1, the flags are same as [Section 2.1](#) definition in [[I-D.ietf-ospf-segment-routing-extensions](#)]. For NSPF-ID Type 2, 3 and NSPF-ID Type 4 only 'R' flag is applicable. Undefined flags for each NSPF-ID type MUST be considered as reserved. Reserved flag bits in each NSPF-ID type specific flags MUST be reset on transmission and ignored on receive.
4. NSPF-ID Algo: 1 octet value represents the SPF algorithm. Algorithm registry is as defined in [[I-D.ietf-ospf-segment-routing-extensions](#)].
5. NSPF-ID: This is the NSP forwarding identifier that would be on the data packet. The value of this field is variable and it depends on the NSPF-ID Type. For Type 1, this is and MPLS SID/

Label. For Type 2 this is a 4 byte IPv4 address. For Type 3 and Type 4, it is a 16 byte IPv6 address. For NSPF-ID Type 2, 3 or

4, if the NSPF-ID Len is set to 0, then FEC Prefix would also become the NSPF-ID. In the case when NSPF-ID Len is 0, NSPF-ID Type is 2, then FEC Prefix length MUST be a 4 byte IPv4 address. Similarly, if NSPF-ID Type is 3 or 4 with NSPF-ID Len is set to 0, then FEC Prefix MUST be of a 16 byte IPv6 Address.

6. No. of NSP-STs: Total number of the NSP sub-TLVs are defined with this 1-octet field. The value MUST NOT be zero.

3.2. OSPFv3 NSP sub-TLVs

A new sub-TLV registry is created (TBD IANA) called NSP sub-TLVs. These are used to describe the path in the form of set of contiguous and ordered sub-TLVs, with first sub-TLV representing the top of the stack or first segment. These set of ordered TLVs can have both topological SIDs and non-topological SIDs (e.g., service segments).

Type 1: SID/Label sub-TLV as defined in [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#). Only Type is defined and Length/Value fields are per [section 2.1](#) of the referenced document.

Type 2: Prefix SID sub-TLV as defined in [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#). Only Type is defined and Length/Value fields are per [section 5](#) of the referenced document.

Type 3: Adjacency SID sub-TLV as defined in [\[I-D.ietf-ospf-ospfv3-segment-routing-extensions\]](#). Only Type is defined and Length/Value fields are per [section 6](#) of the referenced document.

Type 4: Length 4 bytes, value is 4 bytes IPv4 address encoded similar to IPv4 FEC Prefix described above.

Type 5: Length 16 bytes; value is 16 bytes IPv6 address encoded similar to IPv6 FEC Prefix described above.

Type 6: SRv6 Node SID TLV as defined in

[[I-D.li-ospf-ospfv3-srv6-extensions](#)]. Only Type is defined and Length/Value fields are in the referenced document.

Type 7: SRv6 Adjacency-SID sub-TLV as defined in [[I-D.li-ospf-ospfv3-srv6-extensions](#)]. Only Type is defined and Length/Value fields are in the referenced document.

Type 8: SRv6 LAN Adjacency-SID sub-TLV as defined in [[I-D.li-ospf-ospfv3-srv6-extensions](#)]. Only Type is defined and Length/Value fields are in the referenced document.

[3.3.](#) OSPFv3 Non-NSP sub-TLVs

NSPF ID TLV also defines a new sub-TLV registry (TBD IANA) for defining extensible set of sub-TLVs other than describing the path sub-TLVs. Total number of the path sub-TLVs to describe the path are defined in 1-octet field "No.of Other-STs" just before the Non-NSP sub-TLVs. This field serves as a demarcation for set of ordered NSP sub-TLVs and Non-NSP sub-TLVs.

Type 1: Length 0 No value field. Specifies a counter to count number of packets forwarded on this NSPF-ID.

Type 2: Length 0 No value field. Specifies a counter to count number of bytes forwarded on this NSPF-ID specified in the network header (e.g. IPv4, IPv6).

Type 3: Length 4 bytes, and Value is metric of this path represented through the NSPF-ID. Different nodes can advertise the same NSPF-ID for the same FEC-Prefix with a different set of NSP sub-TLVs and the receiving node MUST consider the lowest metric value (TBD more, what happens when metric is same for two different set of NSP sub-TLVs).

[4.](#) Other Considerations

Please refer to [[I-D.ct-isis-nspf-id-for-sr-paths](#)] [section 3](#), 4 and 5.

[5.](#) Acknowledgements

Thanks to Richard Li, Alex Clemm, Kiran Makhijani and Lin Han for initial discussions on this topic.

Earlier versions of [draft-ietf-ospf-segment-routing-extensions](#) have a mechanism to advertise EROs through Binding SID.

6. IANA Considerations

This document requests the following new TLV in IANA OSPFv2 and OSPFv3 TLV code-point registry.

TLV #	Name
-----	-----
TBD	NSPF ID TLV

Chunduri, et al.

Expires September 25, 2018

[Page 11]

Internet-Draft Usage of Non Shortest Path Forwarding (NSPF) IDs March 2018

This document also requests IANA to create new registries for NSPF ID TLV Flags field, NSPF-ID Type, NSPF-ID Flags, NSP sub-TLVs and Non-NSP sub-TLVs in NSPF ID TLV as described in [Section 2](#) and [Section 3](#).

7. Security Considerations

Existing security extensions as described in [\[RFC2328\]](#) and [\[RFC7684\]](#) apply to these segment routing extensions. While OSPF is under a single administrative domain, there can be deployments where potential attackers have access to one or more networks in the OSPF routing domain. In these deployments, stronger authentication mechanisms such as those specified in [\[RFC7474\]](#) SHOULD be used.

Advertisement of the additional information defined in this document introduces no new security concerns in OSPF protocol. However as this extension is related to SR-MPLS and SRH data planes as defined in [\[I-D.ietf-spring-segment-routing\]](#), those particular data plane security considerations does apply here.

8. References

8.1. Normative References

[I-D.ct-isis-nspfid-for-sr-paths]

Chunduri, U., Tantsura, J., and Y. Qu, "Usage of Non

Shortest Path Forwarding (NSPF) IDs in IS-IS", [draft-ct-isis-nspfid-for-sr-paths-01](#) (work in progress), March 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>>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.

8.2. Informative References

[I-D.hegde-spring-traffic-accounting-for-sr-paths]
Hegde, S., "Traffic Accounting for MPLS Segment Routing Paths", [draft-hegde-spring-traffic-accounting-for-sr-paths-01](#) (work in progress), October 2017.

[I-D.ietf-6man-segment-routing-header]
Previdi, S., Filsfils, C., Leddy, J., Matsushima, S., and d. daniel.voyer@bell.ca, "IPv6 Segment Routing Header (SRH)", [draft-ietf-6man-segment-routing-header-10](#) (work in progress), March 2018.

[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-ospf-ospfv3-segment-routing-extensions]
Psenak, P., Filsfils, C., Previdi, S., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPFv3 Extensions for Segment Routing", [draft-ietf-ospf-ospfv3-segment-routing-extensions-11](#) (work in progress), January 2018.

- [I-D.ietf-ospf-segment-routing-extensions]
Psenak, P., Previdi, S., Filsfils, C., Gredler, H.,
Shakir, R., Henderickx, W., and J. Tantsura, "OSPF
Extensions for Segment Routing", [draft-ietf-ospf-segment-
routing-extensions-24](#) (work in progress), December 2017.
- [I-D.ietf-ospf-segment-routing-msd]
Tantsura, J., Chunduri, U., Aldrin, S., and P. Psenak,
"Signaling MSD (Maximum SID Depth) using OSPF", [draft-
ietf-ospf-segment-routing-msd-09](#) (work in progress),
February 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.
- [I-D.li-ospf-ospfv3-srv6-extensions]
Li, Z., Hu, Z., Cheng, D., Talaulikar, K., and P. Psenak,
"OSPFv3 Extensions for SRv6", [draft-li-ospf-
ospfv3-srv6-extensions-01](#) (work in progress), March 2018.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P.
Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF",
[RFC 4915](#), DOI 10.17487/RFC4915, June 2007,
<https://www.rfc-editor.org/info/rfc4915>.

- [RFC7474] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed.,
"Security Extension for OSPFv2 When Using Manual Key
Management", [RFC 7474](#), DOI 10.17487/RFC7474, April 2015,
<https://www.rfc-editor.org/info/rfc7474>.
- [RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W.,
Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute
Advertisement", [RFC 7684](#), DOI 10.17487/RFC7684, November
2015, <https://www.rfc-editor.org/info/rfc7684>.

Uma Chunduri
Huawei USA
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: uma.chunduri@huawei.com

Yingzhen Qu
Huawei USA
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: yingzhen.qu@huawei.com

Jeff Tantsura
Nuage Networks
755 Ravendale Drive
Mountain View, CA 94043
USA

Email: jefftant.ietf@gmail.com