

Workgroup: Network Working Group
Internet-Draft:
draft-ietf-idr-bgp-sr-segtypes-ext-02
Published: 4 February 2024

Intended Status: Experimental
Expires: 7 August 2024

Authors: K. Talaulikar, Ed. C. Filsfils
Cisco Systems Cisco Systems
S. Previdi P. Mattes D. Jain
Huawei Technologies Microsoft Google

Segment Routing Segment Types Extensions for BGP SR Policy

Abstract

This document specifies the signaling of additional Segment Routing Segment Types for BGP SR Policy SAFI.

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 7 August 2024.

Copyright Notice

Copyright (c) 2024 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. Segment Type Sub-TLVs](#)
 - [2.1. Segment Type C](#)
 - [2.2. Segment Type D](#)
 - [2.3. Segment Type E](#)
 - [2.4. Segment Type F](#)
 - [2.5. Segment Type G](#)
 - [2.6. Segment Type H](#)
 - [2.7. Segment Type I](#)
 - [2.8. Segment Type J](#)
 - [2.9. Segment Type K](#)
 - [2.10. Segment Flags](#)
 - [3. IANA Considerations](#)
 - [3.1. SR Policy Segment List Sub-TLVs](#)
 - [3.2. SR Policy Segment Flags](#)
 - [4. Security Considerations](#)
 - [5. References](#)
 - [5.1. Normative References](#)
 - [5.2. Informational References](#)
- [Authors' Addresses](#)

1. Introduction

BGP SR Policy SAFI was introduced by [[I-D.ietf-idr-sr-policy-safi](#)] for the advertisement of Segment Routing (SR) Policy [[RFC8402](#)]. [[I-D.ietf-idr-sr-policy-safi](#)] introduced the base SR Segment Types A and B as specified by the SR Policy Architecture [[RFC9256](#)].

This document specifies the extensions for the advertisement of the remaining SR Segment Types defined in [[RFC9256](#)] in the SR Policy SAFI for both SR-MPLS [[RFC8660](#)] and SRv6 [[RFC8754](#)] [[RFC8986](#)].

The extensions in this document do not impact the SR Policy operations or fault management as specified in [[I-D.ietf-idr-sr-policy-safi](#)].

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. Segment Type Sub-TLVs

The Segment List sub-TLV [[I-D.ietf-idr-sr-policy-safi](#)] encodes a single explicit path towards the endpoint as described in section 5.1 of [[RFC9256](#)]. The Segment List sub-TLV includes the elements of the paths (i.e., segments).

A Segment sub-TLV describes a single segment in a segment list (i.e., a single element of the explicit path).

Section 4 of [[RFC9256](#)] defines several Segment Types:

Type A: SR-MPLS Label
Type B: SRv6 SID
Type C: IPv4 Prefix with optional SR Algorithm
Type D: IPv6 Global Prefix with optional SR Algorithm for SR-MPLS
Type E: IPv4 Prefix with Local Interface ID
Type F: IPv4 Addresses for link endpoints as Local, Remote pair
Type G: IPv6 Prefix and Interface ID for link endpoints as Local, Remote pair for SR-MPLS
Type H: IPv6 Addresses for link endpoints as Local, Remote pair for SR-MPLS
Type I: IPv6 Global Prefix with optional SR Algorithm for SRv6
Type J: IPv6 Prefix and Interface ID for link endpoints as Local, Remote pair for SRv6
Type K: IPv6 Addresses for link endpoints as Local, Remote pair for SRv6

Figure 1: SR Segment Types

[[I-D.ietf-idr-sr-policy-safi](#)] specifies Segment Type Sub-TLVs for the segment types A and B. The following sub-sections specify the sub-TLVs used for encoding each of the other Segment Types above.

As specified in section 2.4.4.2 of [[I-D.ietf-idr-sr-policy-safi](#)], the Segment Type Sub-TLVs specified in this document are also used only by the SRPM [[I-D.ietf-idr-sr-policy-safi](#)] as described in section 4 of [[RFC9256](#)] on the same lines as segment types A and B. Their validation is, therefore, beyond the scope of BGP.

2.1. Segment Type C

The Type C Segment Sub-TLV encodes an IPv4 node address, SR Algorithm, and an optional SR-MPLS SID. The format is as follows:

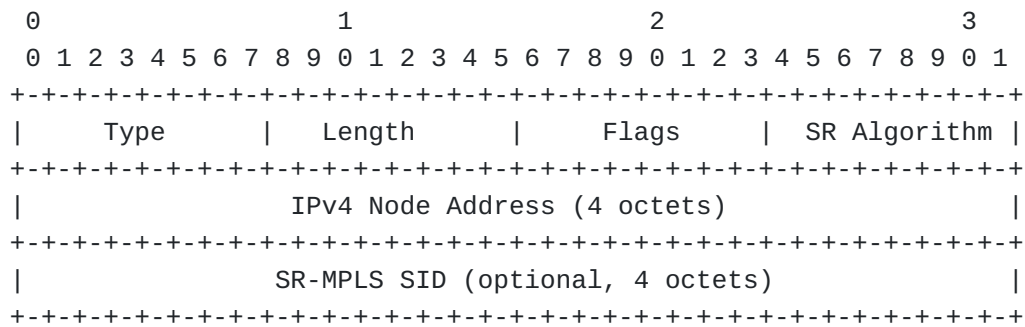


Figure 2: Type C Segment sub-TLV

where:

*Type: 3.

*Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 10 when the SR-MPLS SID is present else it MUST be 6.

*Flags: 1 octet of flags as defined in [Section 2.10](#).

*SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [\[RFC8402\]](#) when A-Flag as defined in [Section 2.10](#) is present. SR Algorithm is used by SRPM [\[I-D.ietf-idr-sr-policy-safi\]](#) as described in section 4 in [\[RFC9256\]](#). When A-Flag is not encoded, this field MUST be set to zero on transmission and MUST be ignored on receipt.

*IPv4 Node Address: a 4-octet IPv4 address representing a node.

*SR-MPLS SID: optional, 4-octet field containing label, TC, S and TTL as defined for Segment Type A [\[I-D.ietf-idr-sr-policy-safi\]](#).

2.2. Segment Type D

The Type D Segment Sub-TLV encodes an IPv6 node address, SR Algorithm, and an optional SR-MPLS SID. The format is as follows:

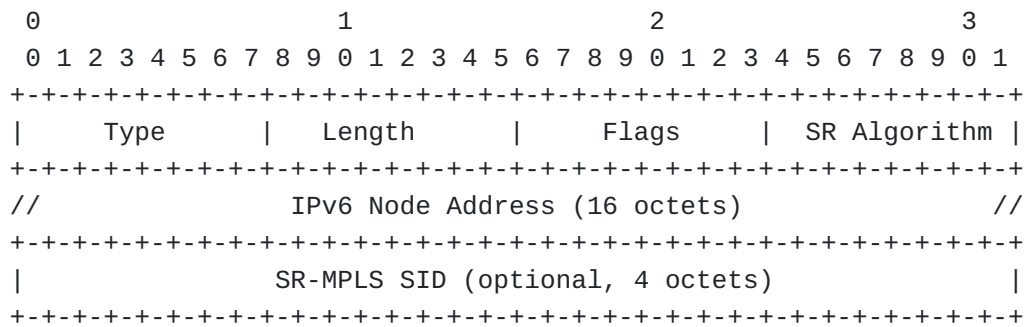


Figure 3: Type D Segment sub-TLV

where:

*Type: 4

*Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 22 when the SR-MPLS SID is present else it MUST be 18.

*Flags: 1 octet of flags as defined in [Section 2.10](#).

*SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [\[RFC8402\]](#) when A-Flag as defined in [Section 2.10](#) is present. SR Algorithm is used by SRPM [\[I-D.ietf-idr-sr-policy-safi\]](#) as described in section 4 in [\[RFC9256\]](#). When A-Flag is not encoded, this field MUST be set to zero on transmission and MUST be ignored on receipt.

*IPv6 Node Address: a 16-octet IPv6 address representing a node.

*SR-MPLS SID: optional, 4-octet field containing label, TC, S and TTL as defined for Segment Type A [\[I-D.ietf-idr-sr-policy-safi\]](#).

2.3. Segment Type E

The Type E Segment Sub-TLV encodes an IPv4 node address, a local interface Identifier (Local Interface ID), and an optional SR-MPLS SID. The format is as follows:

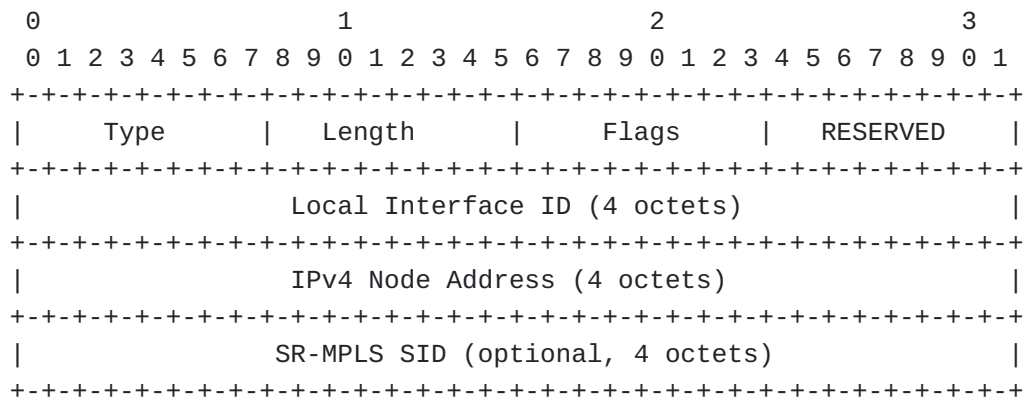


Figure 4: Type E Segment sub-TLV

where:

*Type: 5.

*Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 14 when the SR-MPLS SID is present else it MUST be 10.

*Flags: 1 octet of flags as defined in [Section 2.10](#).

*RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

*Local Interface ID: 4 octets of interface index of local interface (refer TLV 258 of [[RFC9552](#)]).

*IPv4 Node Address: a 4-octet IPv4 address representing a node.

*SR-MPLS SID: optional, 4-octet field containing label, TC, S and TTL as defined for Segment Type A [[I-D.ietf-idr-sr-policy-safi](#)].

2.4. Segment Type F

The Type F Segment Sub-TLV encodes an adjacency local address, an adjacency remote address, and an optional SR-MPLS SID. The format is as follows:

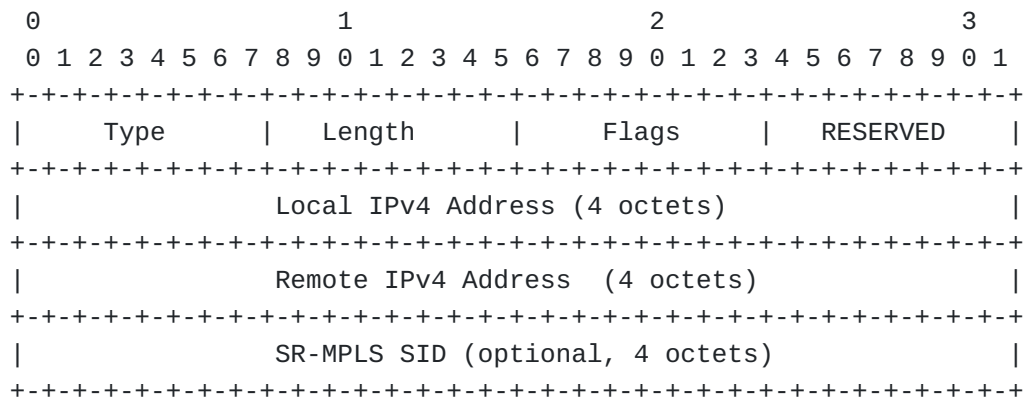


Figure 5: Type F Segment sub-TLV

where:

*Type: 6.

*Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 14 when the SR-MPLS SID is present else it MUST be 10.

*Flags: 1 octet of flags as defined in [Section 2.10](#).

*RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

*Local IPv4 Address: a 4-octet IPv4 address representing the local link address of the node.

*Remote IPv4 Address: a 4-octet IPv4 address representing the link address of the neighbor node.

*SR-MPLS SID: optional, 4-octet field containing label, TC, S and TTL as defined for Segment Type A [[I-D.ietf-idr-sr-policy-safi](#)].

2.5. Segment Type G

The Type G Segment Sub-TLV encodes an IPv6 link-local adjacency with IPv6 local node address, a local interface identifier (Local Interface ID), IPv6 remote node address, a remote interface identifier (Remote Interface ID), and an optional SR-MPLS SID. The format is as follows:

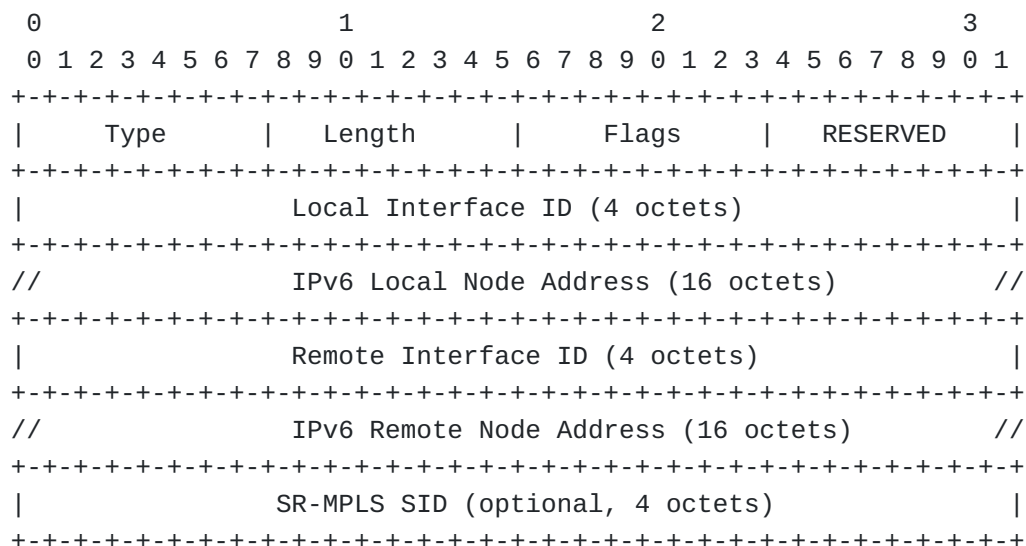


Figure 6: Type G Segment sub-TLV

where:

*Type: 7

*Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 46 when the SR-MPLS SID is present else it MUST be 42.

*Flags: 1 octet of flags as defined in [Section 2.10](#).

*RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

*Local Interface ID: 4 octets of interface index of local interface (refer TLV 258 of [\[RFC9552\]](#)).

*IPv6 Local Node Address: a 16-octet IPv6 address representing the node.

*Remote Interface ID: 4 octets of interface index of remote interface (refer TLV 258 of [\[RFC9552\]](#)). The value MAY be set to zero when the local node address and interface identifiers are sufficient to describe the link.

*IPv6 Remote Node Address: a 16-octet IPv6 address. The value MAY be set to zero when the local node address and interface identifiers are sufficient to describe the link.

*SR-MPLS SID: optional, 4-octet field containing label, TC, S and TTL as defined for Segment Type A [\[I-D.ietf-idr-sr-policy-safi\]](#).

2.6. Segment Type H

The Type H Segment Sub-TLV encodes an adjacency local address, an adjacency remote address, and an optional SR-MPLS SID. The format is as follows:

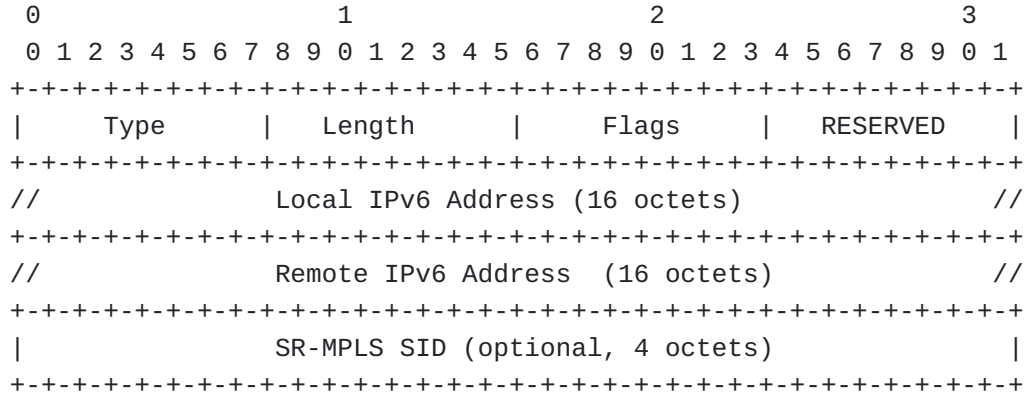


Figure 7: Type H Segment sub-TLV

where:

*Type: 8

*Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 38 when the SR-MPLS SID is present else it MUST be 34.

*Flags: 1 octet of flags as defined in [Section 2.10](#).

*RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

*Local IPv6 Address: a 16-octet IPv6 address representing the local link address of the node.

*Remote IPv6 Address: a 16-octet IPv6 address representing the link address of the neighbor node.

*SR-MPLS SID: optional, 4-octet field containing label, TC, S and TTL as defined for Segment Type A [[I-D.ietf-idr-sr-policy-safi](#)].

2.7. Segment Type I

The Type I Segment Sub-TLV encodes an IPv6 node address, SR Algorithm, and an optional SRv6 SID. The format is as follows:

```

0           1           2           3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|   Type   | Length   |   Flags   | SR Algorithm |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
//           IPv6 Node Address (16 octets)           //
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
//           SRv6 SID (optional, 16 octets)           //
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
//   SRv6 Endpoint Behavior and SID Structure (optional)   //
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Figure 8: Type I Segment sub-TLV

where:

*Type: 14

*Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be one of: 42 when both SRv6 SID and SRv6 Endpoint Behavior & SID Structure are present, 34 when only SRv6 SID is present, or 18 when the SRv6 SID is not present.

*Flags: 1 octet of flags as defined in [Section 2.10](#).

*SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [\[RFC8402\]](#) when A-Flag as defined in [Section 2.10](#) is present. SR Algorithm is used by SRPM [\[I-D.ietf-idr-sr-policy-safi\]](#) as described in section 4 in [\[RFC9256\]](#). When A-Flag is not encoded, this field MUST be set to zero on transmission and MUST be ignored on receipt.

*IPv6 Node Address: a 16-octet IPv6 address representing the node.

*SRv6 SID: optional, a 16-octet IPv6 address.

*SRv6 Endpoint Behavior and SID Structure: Optional, as defined in [\[I-D.ietf-idr-sr-policy-safi\]](#) for Segment Type B.

The TLV 10 defined for the advertisement of Segment Type I in the early draft versions of [\[I-D.ietf-idr-sr-policy-safi\]](#) has been deprecated to avoid backward compatibility issues.

2.8. Segment Type J

The Type J Segment Sub-TLV encodes an IPv6 link-local adjacency with local node address, a local interface identifier (Local Interface ID), remote IPv6 node address, a remote interface identifier (Remote Interface ID), and an optional SRv6 SID. The format is as follows:

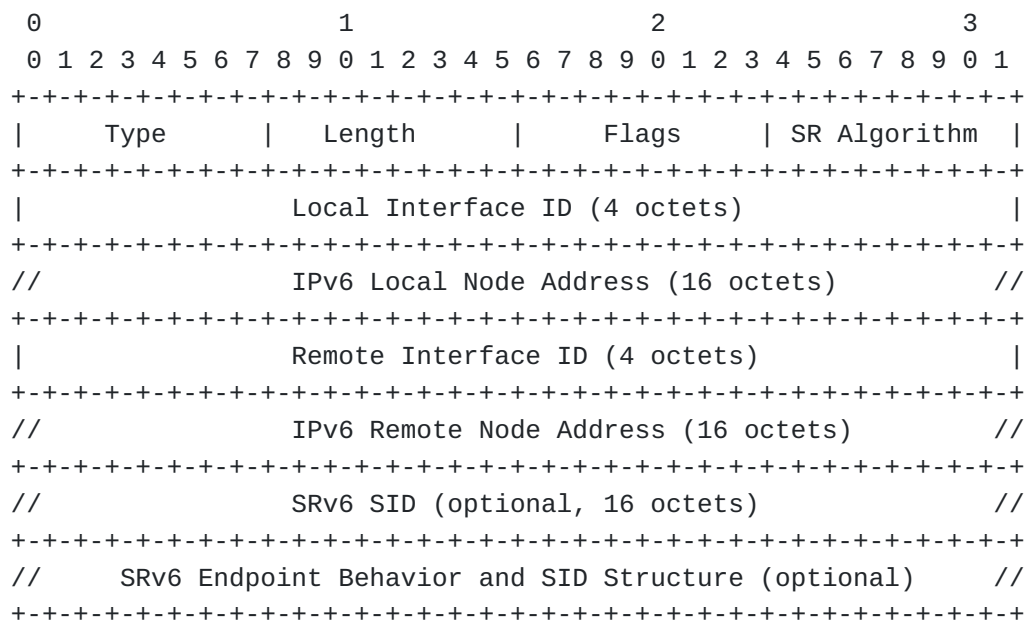


Figure 9: Type J Segment sub-TLV

where:

*Type: 15

*Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be one of: 66 when both SRv6 SID and SRv6 Endpoint Behavior & SID Structure are present, 58 when only SRv6 SID is present, or 42 when the SRv6 SID is not present.

*Flags: 1 octet of flags as defined in [Section 2.10](#).

*SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [\[RFC8402\]](#) when A-Flag as defined in [Section 2.10](#) is present. SR Algorithm is used by SRPM [\[I-D.ietf-idr-sr-policy-safi\]](#) as described in section 4 in [\[RFC9256\]](#). When A-Flag is not encoded, this field MUST be set to zero on transmission and MUST be ignored on receipt.

*Local Interface ID: 4 octets of interface index of local interface (refer TLV 258 of [\[RFC9552\]](#)).

*IPv6 Local Node Address: a 16-octet IPv6 address representing the node.

*Remote Interface ID: 4 octets of interface index of remote interface (refer TLV 258 of [\[RFC9552\]](#)). The value MAY be set to zero when the local node address and interface identifiers are sufficient to describe the link.

*IPv6 Remote Node Address: a 16-octet IPv6 address. The value MAY be set to zero when the local node address and interface identifiers are sufficient to describe the link.

*SRV6 SID: optional, a 16-octet IPv6 address.

*SRV6 Endpoint Behavior and SID Structure: Optional, as defined in [\[I-D.ietf-idr-sr-policy-safi\]](#) for Segment Type B.

The TLV 11 defined for the advertisement of Segment Type J in the early draft versions of [\[I-D.ietf-idr-sr-policy-safi\]](#) has been deprecated to avoid backward compatibility issues.

2.9. Segment Type K

The Type K Segment Sub-TLV encodes an adjacency local address, an adjacency remote address, and an optional SRV6 SID. The format is as follows:

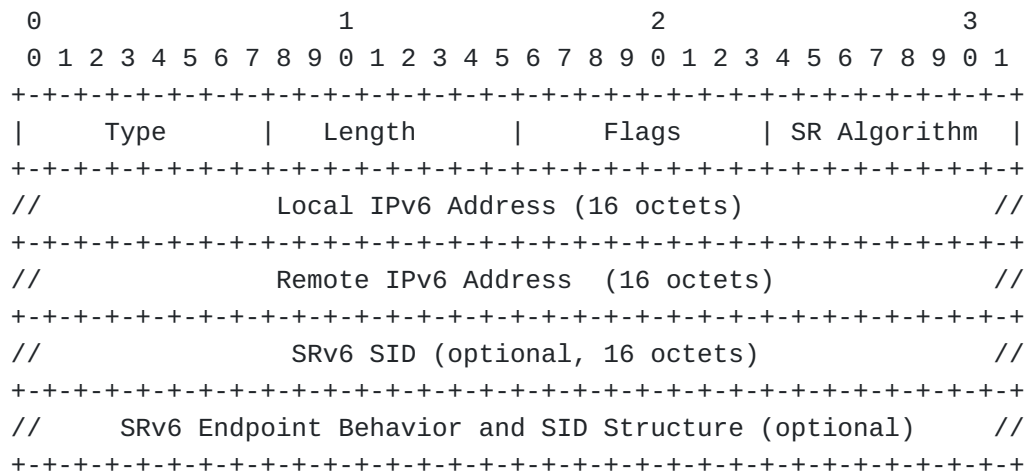


Figure 10: Type K Segment sub-TLV

where:

*Type: 16

*Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be one of: 58 when both SRv6 SID and SRv6 Endpoint Behavior & SID Structure are present, 50 when only SRv6 SID is present, or 34 when the SRv6 SID is not present.

*Flags: 1 octet of flags as defined in [Section 2.10](#).

*SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [\[RFC8402\]](#) when A-Flag as defined in [Section 2.10](#) is present. SR Algorithm is used by SRPM

[[I-D.ietf-idr-sr-policy-safi](#)] as described in section 4 in [[RFC9256](#)]. When A-Flag is not encoded, this field MUST be set to zero on transmission and MUST be ignored on receipt.

*Local IPv6 Address: a 16-octet IPv6 address representing the local link address of the node.

*Remote IPv6 Address: a 16-octet IPv6 address representing the link address of the neighbor node.

*SRV6 SID: optional, a 16-octet IPv6 address.

*SRV6 Endpoint Behavior and SID Structure: Optional, as defined in [[I-D.ietf-idr-sr-policy-safi](#)] for Segment Type B.

The TLV 12 defined for the advertisement of Segment Type K in the early draft versions of [[I-D.ietf-idr-sr-policy-safi](#)] has been deprecated to avoid backward compatibility issues.

2.10. Segment Flags

The Segment Types sub-TLVs described above may contain the following flags in the "Segment Flags" field defined in [[I-D.ietf-idr-sr-policy-safi](#)]. This document introduces additional flags as below:

```
0 1 2 3 4 5 6 7
+--+--+--+--+--+--+
|V|A|S|B|      |
+--+--+--+--+--+--+
```

Figure 11: Segment Flags

where:

V-Flag: existing flag as defined in [[I-D.ietf-idr-sr-policy-safi](#)].

A-Flag: This flag, when set, indicates the presence of SR Algorithm id in the "SR Algorithm" field applicable to various Segment Types. SR Algorithm is used by SRPM [[I-D.ietf-idr-sr-policy-safi](#)] as described in section 4 of [[RFC9256](#)].

S-Flag: This flag, when set, indicates the presence of the SR-MPLS or SRV6 SID depending on the segment type.

B-Flag: existing flag as defined in [[I-D.ietf-idr-sr-policy-safi](#)].

The following applies to the Segment Flags:

*V-Flag applies to all Segment Types including the ones introduced by this document.

*A-Flag applies to Segment Types C, D, I, J, and K. If A-Flag appears with Segment Types A, B, E, F, G, and H, it MUST be ignored.

*S-Flag applies to Segment Types C, D, E, F, G, H, I, J, and K. If S-Flag appears with Segment Types A or B, it MUST be ignored.

*B-Flag applies to Segment Types B, I, J, and K. If B-Flag appears with Segment Types A, C, D, E, F, G, and H, it MUST be ignored.

3. IANA Considerations

This section covers the IANA considerations for this document.

3.1. SR Policy Segment List Sub-TLVs

This document requests the allocation of the following code points from the "SR Policy Segment List Sub-TLVs" registry under the "BGP Tunnel Encapsulation" registry group.

Value	Description	Reference
3	Segment Type C sub-TLV	This document
4	Segment Type D sub-TLV	This document
5	Segment Type E sub-TLV	This document
6	Segment Type F sub-TLV	This document
7	Segment Type G sub-TLV	This document
8	Segment Type H sub-TLV	This document
14	Segment Type I sub-TLV	This document
15	Segment Type J sub-TLV	This document
16	Segment Type K sub-TLV	This document

Table 1: SR Policy Segment List Code Points

3.2. SR Policy Segment Flags

This document requests the allocation of code points from the "SR Policy Segment Flags" registry under the "BGP Tunnel Encapsulation" registry group.

Bit	Description	Reference
1	SR Algorithm Flag (A-Flag)	This document
2	SID Specified Flag (S-Flag)	This document

Table 2: SR Policy Segment Flags

4. Security Considerations

This document does not introduce any security considerations over [I-D.ietf-idr-sr-policy-safi].

5. References

5.1. Normative References

[I-D.ietf-idr-sr-policy-safi]

Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-safi-00, 4 February 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-00>>.

[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>>.

[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>>.

[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>>.

[RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.

[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>>.

[RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.

[RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.

5.2. Informational References

Authors' Addresses

Ketan Talaulikar (editor)
Cisco Systems
India

Email: ketant.ietf@gmail.com

Clarence Filsfils
Cisco Systems
Brussels
Belgium

Email: cfilsfil@cisco.com

Stefano Previdi
Huawei Technologies
Italy

Email: stefano@previdi.net

Paul Mattes
Microsoft
One Microsoft Way
Redmond, WA 98052
United States of America

Email: pamattes@microsoft.com

Dhanendra Jain
Google

Email: dhanendra.ietf@gmail.com