

Network Working Group
Internet Draft
Intended status: Standards Track
Expires: October 8, 2022

W. Cheng
W. Jiang
China Mobile
C. Lin
M. Chen
New H3C Technologies
L. Gong
China Mobile

Y.Liu

ZTE Corporation

April 8, 2022

**IS-IS Extension to Advertise SRv6 SIDs using SID Block
draft-cheng-lsr-isis-srv6-sid-block-00**

Abstract

This document proposes a simplified method to advertise SRv6 SIDs in IS-IS. The SRv6 SID Block is composed of a number of continuous SIDs within the address range of a Locator. When a SID is assigned from the SID Block, it is described by an index based on the SID Block, instead of the whole 128-bit IPv6 address.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on October 8, 2022.

Copyright Notice

Copyright (c) 2022 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Requirements Language	3
2.	SRv6 SID Block sub-TLV	4
3.	SRv6 End.X SID Index sub-TLV	5
4.	SRv6 LAN End.X SID Index sub-TLV	6
5.	SRv6 End SID Index sub-TLV	7
6.	Computations of SID Index	8
7.	Rule-based SID Block for Large-Scale End.X SID Advertisement.	9
8.	Security Considerations	12
9.	IANA Considerations	13
10.	References	13
10.1.	Normative References	13
10.2.	Informative References	13
11.	Acknowledgments	14
	Authors' Addresses	15

[1.](#) Introduction

The Segment Routing (SR) allows for a flexible definition of end-to-end paths by encoding paths as sequences of topological sub-paths, called "segments". As defined in [\[RFC8402\]](#) and [\[RFC8986\]](#), an SRv6 Segment Identifier (SID) is an IPv6 address explicitly associated with the segment and consists of Locator, Function and Argument parts.

[I-D.ietf-lsr-isis-srv6-extensions] defines the SRv6 End SID sub-TLV, the SRv6 End.X SID sub-TLV, and the SRv6 LAN End.X SID sub-TLV in IS-IS.

The SRv6 End SID sub-TLV is used to advertise an SRv6 SID with Endpoint behaviors which do not require a particular neighbor. The SRv6 End SID sub-TLV is used to advertise an SRv6 SID associated with a point to point adjacency. The SRv6 LAN End.X SID sub-TLV sub-TLV is used to advertise an SRv6 SID associated with a LAN adjacency. Each of these sub-TLVs contains a complete 128-bit SID and the sub-TLV length is quite long.

Multiple SRv6 End.X SIDs can be associated with the same point to point adjacency or the same physical LAN neighbor. Each SID is advertised in a single SRv6 End.X SID sub-TLV or SRv6 LAN End.X SID sub-TLV. These SIDs are possibly associated to the same Locator, therefore the main differences among the sub-TLVs may be a few bits in the Function part of SID and the Endpoint Behavior value indicating different flavors.

The number of End.X SIDs has a positive correlation with the number of neighbors. Assume that, each neighbor is assigned with End.X SIDs, and each End.X behavior has several different flavors, such as PSP, USP, USD, no PSP/USP/USD, etc. Then, the number of End.X SIDs will be at least the number of neighbors multiplied by the number of flavors.

If Flexible-Algorithm is applied on SRv6 forwarding plane as defined in [[I-D.ietf-lsr-flex-algo](#)], a node generally advertises a Flex-Algorithm specific locator for each Flex-Algorithm it participates in and also advertises associated SRv6 END.X SIDs for every link that has not been pruned from the Flex-Algorithm computation.

There may be other scenarios in which the number of End.X SIDs is quite large. Under these circumstances, advertising SRv6 SIDs needs considerable amounts of IS-IS LSPs and may be a burden of IS-IS nodes.

This document proposes a simplified method to advertise SRv6 SIDs in IS-IS. A number of continuous SIDs are assigned within the address range of a Locator, which is defined as the SRv6 SID Block. When a SID is assigned from the block, it is described by an index based on the block, instead of the whole 128-bit IPv6 address.

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 SID Block sub-TLV

An SRv6 SID Block is composed of a number of continuous SIDs within the address range of a Locator. When any SID assigned in an SRv6 SID Block is advertised in IS-IS, the new sub-TLVs defined in this document SHOULD be used, instead of the SRv6 End SID sub-TLV, the SRv6 End.X SID sub-TLV, and the SRv6 LAN End.X SID sub-TLV defined in [[I-D.ietf-lsr-isis-srv6-extensions](#)].

The SRv6 SID Block sub-TLV is defined in this document to advertise an SRv6 SID Block. This new sub-TLV is carried in the SRv6 Locator TLV defined in [[I-D.ietf-lsr-isis-srv6-extensions](#)].

The SRv6 SID Block sub-TLV has the following format:

```

      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      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Block-ID   |      Offset     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Start SID (128 bits) . . . |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Start SID (cont . . .) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Start SID (cont . . .) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Start SID (cont . . .) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     End SID (128 bits) . . . |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     End SID (cont . . .) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     End SID (cont . . .) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     End SID (cont . . .) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Sub-sub-tlv-len|      Sub-sub-TLVs (variable) . . . |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

o Type: TBD.

o Length: variable.

o Block-ID: 1 octet. Identity of the advertised SRv6 SID Block. It MUST be unique in a node.

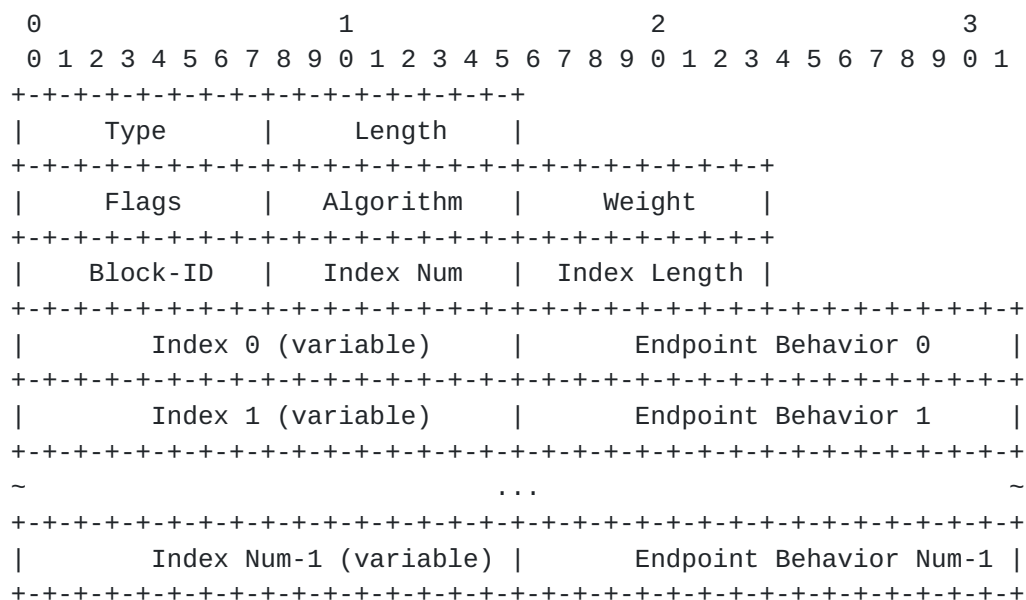
- o Offset: 1 octet. Number of shifting bits for the index of SIDs in this block.
- o Start SID: 16 octets. The smallest SID in this block.
- o End SID: 16 octets. The largest SID in this block.
- o Sub-sub-tlv-len: 1 octet. Number of octets used by sub-sub-TLVs.
- o Optional Sub-sub-TLVs.

The address range of an SRv6 SID Block, indicated by the Start SID and the End SID, must be covered by the associated Locator. Multiple SRv6 SID Blocks can be associated to the same Locator, and their address ranges should not overlap.

3. SRv6 End.X SID Index sub-TLV

[I-D.ietf-lsr-isis-srv6-extensions] defines SRv6 End.X SID sub-TLV to advertise an SRv6 SID associated with a point to point adjacency. If such SID is assigned in an SRv6 SID Block, it must be advertised by SRv6 End.X SID Index sub-TLV instead.

The SRv6 End.X SID Index sub-TLV has the following format:



- o Type: TBD.
- o Length: variable.

- o Flags: 1 octet. Same as the Flags field in SRv6 End.X SID sub-TLV.
- o Algorithm: 1 octet. Same as the Algorithm field in SRv6 End.X SID sub-TLV.
- o Weight: 1 octet. Same as the Weight field in SRv6 End.X SID sub-TLV.
- o Block-ID: 1 octet. Indicating the associated SRv6 SID Block in which the SIDs are assigned.
- o Index Num: 1 octet. Number of the following indexes.
- o Index Length: 1 octet. Number of octets used by a single index.
- o Index N: variable. Index of a SID in the SRv6 SID Block.
- o Endpoint Behavior N: 2 octets. Behavior of the SID indicated by the preceding Index N field. Same as the Endpoint Behavior field in SRv6 End.X SID sub-TLV.

The SRv6 End.X SID Index sub-TLV can carry multiple SIDs which are assigned in an SRv6 SID Block identified by the Block-ID. These SIDs are indicated by the indexes in the SRv6 SID Block, instead of the whole 128-bit IPv6 addresses. They share the same Flags, Algorithm and Weight, but their Endpoint Behavior can be different.

The conversion from 128-bit SID into index, or vice versa, is described in [Section 6](#).

4. SRv6 LAN End.X SID Index sub-TLV

The SRv6 LAN End.X SID Index sub-TLV is used to advertise SRv6 SIDs associated with a LAN adjacency, as substitution of the SRv6 LAN End.X SID sub-TLV for SID Block. The main difference from the SRv6 End.X SID Index sub-TLV is that the SRv6 LAN End.X SID Index sub-TLV has to include the System ID of the physical neighbor on the LAN with which the SRv6 SIDs are associated.

The SRv6 LAN End.X SID Index sub-TLV has the following format:

```

      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   |                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Neighbor System-ID (ID length octets)
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Flags   | Algorithm |   Weight   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Block-ID | Index Num | Index Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Index 0 (variable) | Endpoint Behavior 0 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Index 1 (variable) | Endpoint Behavior 1 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                                     ...                                     ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Index Num-1 (variable) | Endpoint Behavior Num-1 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

o Type: TBD.

o Length: variable.

o Neighbor System-ID: IS-IS System-ID of length "ID Length" as defined in [ISO10589]. Same as the Neighbor System-ID field in SRv6 LAN End.X SID sub-TLV.

o Flags, Algorithm, Weight, Block-ID, Index Num, Index Length, Index N, Endpoint Behavior N: Same with the corresponding fields in SRv6 End.X SID Index sub-TLV.

5. SRv6 End SID Index sub-TLV

[I-D.ietf-lsr-isis-srv6-extensions] defines SRv6 End SID sub-TLV to advertise SRv6 End SID. If such SID is assigned in an SRv6 SID Block, it must be advertised by SRv6 End SID Index sub-TLV instead.

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								
+--+																																							

o Type: TBD.

o Length: variable.

o Flags: 1 octet. Same as the Flags field in SRv6 End SID sub-TLV.

o Block-ID: 1 octet. Indicating the associated SRv6 SID Block in which the SIDs are assigned.

o Index Num: 1 octet. Number of the following indexes.

o Index Length: 1 octet. Number of octets used by a single index.

o Index N: variable. Index of a SID in the SRv6 SID Block.

o Endpoint Behavior N: 2 octets. Behavior of the SID indicating by the preceding Index N field. Used as the same as the Endpoint Behavior field in SRv6 End SID sub-TLV.

6. Computations of SID Index

When a SID is assigned from a SRv6 SID Block, its index can be computed in the following way:

Index = (SID - Start SID) right shift Offset bits

When a node receives any of the above-mentioned SID Index sub-TLVs, along with the SRv6 SID Block sub-TLV, it can convert the index to SID in the following way:

SID = (Index left shift Offset bits) + Start SID

7. Rule-based SID Block for Large-Scale End.X SID Advertisement

In this section, the rule-based SID Block is proposed to achieve greater compressibility for End.X SID advertisements in massive End.X SIDs scenarios.

Three new sub-TLVs are defined: the Rule-based End.X SID Block sub-TLV, the Rule-based End.X SID Index sub-TLV or the Rule-based LAN End.X SID Index sub-TLV.

The Rule-based End.X SID Block sub-TLV advertises a SID Block, along with the rule for assigning a series of End.X SIDs with different flavors in the SID Block. The Rule-based End.X SID Index sub-TLV or the Rule-based LAN End.X SID Index sub-TLV advertises a start index in the SID Block, indicating that End.X SIDs with different flavors are assigned following the start index: <Index 0 (start index), Endpoint Behavior 0>, <Index 1 (start index + 1), Endpoint Behavior 1>, ..., <Index N (start index + N), Endpoint Behavior N>. N is equal to (Behavior Num - 1).

The Rule-based End.X SID Block sub-TLV has the following format:

- o Flags: 1 octet. Same as the Flags field in SRv6 End.X SID sub-TLV or SRv6 LAN End.X SID sub-TLV. SIDs assigned in the same rule-based SID Block share the same Flags.
- o Algorithm: 1 octet. Same as the Algorithm field in SRv6 End.X SID sub-TLV or SRv6 LAN End.X SID sub-TLV. SIDs assigned in the same rule-based SID Block share the same Algorithm.
- o Weight: 1 octet. Same as the Weight field in SRv6 End.X SID sub-TLV or SRv6 LAN End.X SID sub-TLV. SIDs assigned in the same rule-based SID Block share the same Weight.
- o Behavior Num: 1 octet. Number of End.X SIDs with different flavors for the same neighbor.
- o Endpoint Behavior N: 2 octets. Behavior of the Nth End.X SID. Endpoint Behavior 0 to (Num-1) specify a series of End.X SIDs with different flavors.
- o Sub-sub-tlv-len: 1 octet. Number of octets used by sub-sub-TLVs.
- o Optional Sub-sub-TLVs.

The Rule-based End.X SID Index sub-TLV has the following format:

```

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   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Block-ID | Index Length | Start Index (variable) ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

- o Type: TBD.
- o Length: variable.
- o Block-ID: 1 octet. Indicating the associated Rule-based SID Block in which the SIDs are assigned.
- o Index Length: 1 octet. Number of octets used by the following Start Index field.
- o Start Index: variable. Index of the first SID in the Rule-based SID Block.

The Rule-based LAN End.X SID Index sub-TLV has the following format:

[illegible]

- o Type: TBD.
- o Length: variable.
- o Neighbor System-ID: IS-IS System-ID of length "ID Length" as defined in [[IS010589](#)]. Same as the Neighbor System-ID field in SRv6 LAN End.X SID sub-TLV.
- o Block-ID: 1 octet. Indicating the associated Rule-based SID Block in which the SIDs are assigned.
- o Index Length: 1 octet. Number of octets used by the following Start Index field.
- o Start Index: variable. Index of the first SID in the Rule-based SID Block.

8. Security Considerations

TBD

9. IANA Considerations

Type	Description
----	-----
TBD	SRv6 SID Block sub-TLV
TBD	SRv6 End.X SID Index sub-TLV
TBD	SRv6 LAN End.X SID Index sub-TLV
TBD	SRv6 End SID Index sub-TLV
TBD	Rule-based End.X SID Block sub-TLV
TBD	Rule-based End.X SID Index sub-TLV
TBD	Rule-based LAN End.X SID Index sub-TLV

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), May 2017
- [I-D.ietf-lsr-isis-srv6-extensions] Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over IPv6 Dataplane", [draft-ietf-lsr-isis-srv6-extensions-18](#) (work in progress), October 2021.
- [ISO10589] International Organization for Standardization, "Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)", Nov 2002.

10.2. Informative References

- [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>>.
- [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>>.

[I-D.ietf-lsr-flex-algo] Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", [draft-ietf-lsr-flex-algo-18](#) (work in progress), October 2021.

11. Acknowledgments

The authors would like to thank the following for their valuable contributions of this document:

TBD

Authors' Addresses

Weiqiang Cheng
China Mobile

Email: chengweiqiang@chinamobile.com

Wenying Jiang
China Mobile

Email: jiangwenying@chinamobile.com

Changwang Lin
New H3C Technologies

Email: linchangwang.04414@h3c.com

Mengxiao Chen
New H3C Technologies

Email: chen.mengxiao@h3c.com

Liyan Gong
China Mobile

Email: gongliyan@chinamobile.com

Yao Liu
ZTE Corporation

Email: liu.yao71@zte.com.cn

