

SRv6 Operations
Internet-Draft
Intended status: Standards Track
Expires: 23 April 2026

J. Horn
K. Michielsen
Cisco Systems
N. Morris
Verizon
M. Horneffer
Deutsche Telekom
C. Hassen
Bell Canada
S. Li
China Mobile
20 October 2025

Addressing Recommendations for SRv6
draft-horn-srv6ops-srv6addressing-00

Abstract

This document describes SRv6 addressing for NEXT-CSID locator format. It introduces concepts of Blocks, Sets, and Node IDs, and explains how summarization boundaries and flexible algorithm support can be implemented for both small and large networks.

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 23 April 2026.

Copyright Notice

Copyright (c) 2025 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	2
2. Terminology	3
3. SRv6 Addressing principles and considerations	3
4. Locator format F3216	4
4.1. Global and Local ID Blocks	4
4.2. Set and Node ID	5
5. Addressing	6
5.1. Area Addressing	6
5.2. Summarization	7
5.3. Small Networks	7
5.4. Large Networks	8
5.5. Flexible Algorithms	9
6. Other addressing	11
6.1. Loopbacks	11
6.2. Interfaces	12
7. Security Considerations	12
8. IANA Considerations	12
9. Acknowledgements	12
10. Contributors	12
11. References	13
11.1. Normative References	13
11.2. Informative References	13
Authors' Addresses	14

1. Introduction

Segment Routing over IPv6 (SRv6) enables the creation of source routing policies by encoding instructions in the form of IPv6 addresses, known as Segment Identifiers (SIDs). To efficiently manage SRv6 deployments at scale, it is critical to design a structured and summarizable addressing plan. This document focuses on SRv6 deployments using the NEXT-CSID compression format (RFC9800) which is the most efficient and widely deployed SRv6 compression. Addressing structure is based on an SRv6 SID structure with 32-bit Locator-Block length and 16-bit CSID length, or F3216 locator format for short, and provides operational guidelines for Service Providers network of different sizes.

2. Terminology

This document leverages the terms defined in [RFC8402], [RFC8754], [RFC8986], [RFC9350], and [RFC9800], in particular segment, segment list, Segment Identifier (SID), SID list, SR policy, prefix segment, adjacency segment, SRH, SR domain, SR source node, SR segment endpoint node, transit node, SRv6 endpoint behavior, flavor, SID block, locator, function, and argument, IGP Flexible Algorithm (FA), Locator-Block, Locator-Node, Compressed-SID (CSID), CSID container, CSID length, compressed SID list, Global Identifier Block (GIB), and Locator Identifier Block (LIB). The reader is assumed to be familiar with this terminology.

This document introduces the following new terms:

Node ID: Identifier for a specific network node within a block or set

F3216: A short-hand notation that refers to the format of NEXT-CSID flavor SID with a 32-bit Locator-Block and a 16-bit CSID. NEXT-CSID ([RFC9800]) implementations must support this format.

Set ID: Identifier for subdivision of the block providing summarization boundaries

SID Space: Common prefix covering all Block prefixes in a network

3. SRv6 Addressing principles and considerations

SRv6 locator is a routable prefix allocated to a node and associated with a single algorithm. An SRv6 locator is often called simply "locator".

Simplicity: A simple consistent addressing plan eases allocation and operations

Hierarchical and aggregatable: A hierarchical addressing plan enables simple and efficient summarization which is fundamental for SRv6 scalability.

SID list compression efficiency: The addressing plan should offer efficient SID list compression, minimizing the number of CSID containers required for any TE policy.

Extensibility: The addressing plan must be able to adapt as the organization and the network evolves.

Infrastructure addresses and SRv6 locators serve different purposes and perform different roles in the network. Therefore, SRv6 locators must be allocated from a separate address range to facilitate network operation and security.

Since SRv6 locators have different allocation requirements and constraints, they should be allocated using a dedicated SRv6 addressing plan. An existing IPv6 addressing plan should not be a constraint for the SRv6 addressing plan.

This document uses the SRv6 SID address block 5f00::/16 ([RFC9602]), but the same principles are applicable when using ULA or GUA address ranges.

This document focuses on the F3216 locator format, currently the most widely deployed format due to its compression efficiency and simplicity. But all recommendations are applicable to any other format.

It is valuable to use nibble boundaries whenever possible to improve human readability.

4. Locator format F3216

[RFC9800] structures an SRv6 locator as the combination Locator-Block:Locator-Node. For simplicity, we will use the terms Block ID and CSID to refer to Locator-Block and Locator-Node respectively.

Format F3216 uses 32 bits for the Block ID and 16 bits for the CSID. This results in a 16-bit Locator-Node length.

This locator structure can be represented as follows, using one character for each 4 bits:

BBBB:BBBB:NNNN::/48

Where:

BBBB:BBBB = 32-bit Block ID

NNNN = 16-bit CSID

4.1. Global and Local ID Blocks

[RFC9800] specifies the Global ID Block (GIB) and Local ID Block (LIB) as two non-overlapping sub-spaces of the CSID numbering space.

As stated in [RFC9800]:

| "The CSID values that are allocated from the GIB have a global
| semantic within the Locator-Block, while those that are allocated
| from the LIB have a local semantic on an SR segment endpoint node
| and within the scope of the Locator-Block."

[RFC9800] states that, for a CSID allocated from the GIB:

| "The tuple (Locator-Block, CSID) identifies the same segment
| across all nodes of the SR domain."

And, for a CSID allocated from the LIB:

| "The tuple (Locator-Block, CSID) identifies a different segment on
| each node of the SR domain."

CSID values allocated from the GIB represent globally unique segments and can be included as part of the locator. On the other hand, CSID values allocated from the LIB represent segments that are locally significant to each node and, as a such, cannot be used as part of a locator. This separation is fundamental to ensuring fully deterministic behavior for TE.

The boundary between the GIB and LIB is flexible, but it must be consistent across all nodes of the SR domain.

The industry-accepted default sub-spaces for the F3216 format uses the uSID values 0x0001 to 0xdfff for GIB and the SIDs 0xe000 to 0xffff for LIB. This implies that for each Block there are 57,344 global SIDs (excluding CSID value 0, which is reserved to mark the end-of-container) and 8,192 local SIDs.

4.2. Set and Node ID

To organize address spaces more efficiently, the CSID numbering space within a block is divided into smaller, equally sized sections called sets. Each set has a unique Set ID. The remaining less significant bits of the CSID space are used for Node IDs, to uniquely identify each node within a set. The combination of the Block ID and Set ID creates a subnet prefix. This subnet can be assigned to a specific routing area or domain. This approach makes logical address assignment and address summarization much simpler and more efficient.

A Set ID can be any length, but for human readability it is recommended to use sizes that align with nibble (4-bit) boundaries. The most common size for a Set ID is 8 bits.

The locator structure is as follows:

BBBB:BBBB:SSNN::

Where:

BBBB:BBBB = 32-bit Block ID

SS = 8-bit Set ID

NN = 8-bit Node ID

This structure supports 256 different sets within each block. Each set can contain up to 256 unique Node IDs.

When using the default split between GIB and LIB, then Set IDs from 0x00 to 0xdf are used for locators (globally significant) and Set IDs from 0xe0 to 0xff are used for local functions.

5. Addressing

5.1. Area Addressing

An area is any part of the network with defined boundaries where address summarization can be performed (such as an ISIS level or an OSPF area). The number of devices in an area can vary, typically containing anywhere from a few hundred to several thousand devices.

Operators assign an appropriate number of sets to each area, based on the number of devices it needs to support.

Set Assignment Example:

Block ID: 5f00:0000::/32

Area 1 - 500 devices:

Assigned 2 Sets (IDs 0x00 and 0x01) providing a total of 511 global SIDs:

- * 5f00:0000:0000::/40 (255 locators)
- * 5f00:0000:0100::/40 (256 locators)

Area 2 - 1500 devices:

Assigned 6 Sets (IDs 0x02, 0x03, 0x04, 0x05, 0x11, 0x12) for a total of 1536 global SIDs:

- * 5f00:0000:0200::/40 (256 locators)
- * 5f00:0000:0300::/40 (256 locators)
- * 5f00:0000:0400::/40 (256 locators)
- * 5f00:0000:0500::/40 (256 locators)

- * 5f00:0000:1100::/40 (256 locators)
- * 5f00:0000:1200::/40 (256 locators)

Sets assigned to single area do not have to be consecutive. (See the Summarization section for more details.)

5.2. Summarization

Summarizing locators is essential for achieving virtually unlimited scalability in an SRv6 network. The need for summarization depends on the size of the network and the scalability of individual devices. In very small networks, summarization may not be necessary.

It proves to be operationally very beneficial to use a unified summarization scheme across the entire network. No matter the area size, always summarize at set boundaries. This approach keeps routing tables simple, clean, and structured, while still providing strong summarization benefits. It also eases the preservation of addressing space for future growth.

Summarization Example:

Area 1 - summaries:

- * 5f00:0000:0000::/40
- * 5f00:0000:0100::/40

Area 2 - summaries:

- * 5f00:0000:0200::/40
- * 5f00:0000:0300::/40
- * 5f00:0000:0400::/40
- * 5f00:0000:0500::/40
- * 5f00:0000:1100::/40
- * 5f00:0000:1200::/40

It is technically possible to advertise larger blocks, such as 5f00:0000:0000::/39 for Area 1 or 5f00:0000:0200::/38 for Area 2. However, the summarization gain here is negligible, while increasing operational complexity.

5.3. Small Networks

A small network is any network where every device can be assigned a locator from a single block. Theoretically, a single block can support up to 57,344 devices. In practice it is suitable for networks with up to around 35,000 devices.

In this scenario, only one block is used to address all devices in the network. This is the optimal option for TE, because all SIDs in any policy will come from the same block. As a result, a headend can use the minimal possible number of containers in its TE policies.

A small network is divided into areas, depending on the scalability of devices and the routing protocol. Each area is assigned the appropriate number of sets based on its current size. Devices within each area are assigned locators from the sets of the area.

Example:

Block ID: 5f00:0000::/32

Area 1 - 500 devices - 2 sets 0x00 and 0x01

First device locator: 5f00:0000:0001::/48

Last device Locator: 5f00:0000:01ff::/48

5.4. Large Networks

A large network is any network with more than approximately 35,000 devices. In this case, a single block is not enough to assign a unique locator to each device. Therefore, the network is divided into regions, with each region able to contain up to about 35,000 devices. The number of regions depends on the total network size.

It is important to make each region as large as possible to minimize the number of block boundaries, which benefits efficient TE. Note that a "region" here is a logical grouping and is not limited to a geographical area. Within each region, sets are assigned following the same principles as in small network addressing.

For a network with up to approximately 500,000 SRv6 devices, 16 regions are needed. The addressing format is the following:

BBBB:BBBR:SSNN::

Where:

BBBB:BBB = 28-bit SRv6 Space

R = 4-bit Region ID

SS = 8-bit Set ID

NN = 8-bit Node ID

Example:

In a network using SRv6 Space 5f00:0000::/28, the locator of a device with Node ID 0x43 in region 0x5, set 0x75, is:

5f00:0005:7543::/48

For a very large network with more than 500,000 SRv6 devices, more regions are needed. In this case, allocate 8 bits (2 nibbles) for the Region ID, allowing to scale up to approximately 8 million SRv6 devices. The addressing format becomes:

BBBB:BBRR:SSNN::

Where:

BBBB:BB = 24-bit SRv6 Space

RR = 8-bit Region ID

SS = 8-bit Set ID

NN = 8-bit Node ID

Example:

In a network using SRv6 Space 5f00:0000::/24, the locator of a device with Node ID 0x43 in region 0x11, set 0x75, is:

5f00:0011:7543::/48

5.5. Flexible Algorithms

Using Flexible Algorithms (FAs) in SRv6 networks requires assigning multiple independent locators to single device, one for each algorithm. To ensure efficient summarization schema and operational simplicity, follow these guidelines:

Never use the same block for different algorithms. If two locators share the same block, they will share the same LIB, which reduces scalability. Therefore, the FA ID should always be encoded in the Block ID.

Encode FA in the highest-order bits of the network addressing plan. This approach makes summarization efficient. Locators for different algorithms cannot be summarized together, so summarization should happen at the level of sets (or at the region level in large networks).

For a given device, use the same Node ID across all algorithms. This simplifies operations and device management.

Typically, a single nibble (4 bits) is enough for the FA ID.

It is very unlikely that a service provider will use more than 16 flexible algorithms.

Locator Structure for small networks:

BBBB:BBBF:SSNN::

Where:

BBBB:BBB = 28-bit SRv6 space

F = 4-bit Flexible Algorithm ID

SS = 8-bit Set ID

NN = 8-bit Node ID

Locator Structure for large networks:

BBBB:BBFR:SSNN::

Where:

BBBB:BB = 24-bit SRv6 space

F = 4-bit Flexible Algorithm ID

R = 4-bit Region ID

SS = 8-bit Set ID

NN = 8-bit Node ID

Locator Structure for very large networks:

BBBB:BFRR:SSNN::

Where:

BBBB:B = 20-bit SRv6 space

F = 4-bit Flexible Algorithm ID

RR = 8-bit Region ID

SS = 8-bit Set ID

NN = 8-bit Node ID

Example Small Network:

In a network using the default algorithm (Algorithm 0) and two flexible algorithms, using SRv6 Space 5f00:0000::/28, the locators of a device with Node ID 0x43 in set 0x75 are:

- * Locator for default algorithm - 5f00:0000:7543::/48
- * Locator for FA ID 128 - 5f00:0001:7543::/48
- * Locator for FA ID 129 - 5f00:0002:7543::/48

Example Large Network:

In a network using the default algorithm and two flexible algorithms, using SRv6 Space 5f00:0000::/24, the locators of a device with Node ID 0x43 in region 0x5, set 0x75 are:

- * Locator for default algorithm - 5f00:0005:7543::/48
- * Locator for FA128 - 5f00:0015:7543::/48
- * Locator for FA129 - 5f00:0025:7543::/48

6. Other addressing

The addressing of interfaces and loopbacks is independent of locator addressing. If the already uses legacy IPv6 addressing, there is no need to change it. But some synergic addressing strategies can simplify network operations and increase scalability.

6.1. Loopbacks

Assigning loopback addresses from the node's locator space has several advantages:

- * Simpler summarization rules: all loopback addresses can be summarized with locators.
- * Smaller IGP database: fewer unique prefixes are advertised.
- * Simpler addressing plan: no separate loopback addressing plan required.
- * Using loopback as encapsulation source simplifies configuration and troubleshooting.

Although loopback addresses can technically be assigned from any algorithm's locator, it is most logical and common to use the default algorithm's locator.

Example:

For a device with Node ID 0x43 in set 0x75:

- * Locator for default algorithm: 5f00:0000:7543::/48
- * Loopback address: 5f00:0000:7543::1/128

6.2. Interfaces

Each IPv6 interface has a Link-Local unicast address ([RFC4291]). Link-local addresses are sufficient to establish routing adjacencies and build a fully functional IPv6 and SRv6 network. Therefore, it is possible to use only link-local addresses on infrastructure links between routers ([RFC7404]).

This simplifies the overall addressing plan and reduces the IGP database but comes with the drawback that interfaces with only link-local addresses cannot be reached remotely, which may cause operational challenges.

An elegant workaround to mitigate this limitation is to assign End.X SIDs to all interfaces. Then the interfaces are remotely pingable since End.X SIDs are remotely reachable over ICMP.

7. Security Considerations

This document does not introduce new security issues beyond those existing in SRv6 and IPv6.

8. IANA Considerations

This document makes no IANA requests.

9. Acknowledgements

The authors would like to acknowledge the contribution of Francois Clad, Clarence Filsfils and Ketan Talaulikar for their valuable input and review of this document.

10. Contributors

Daniel Voyer
Cisco Systems
Montreal
Canada
Email: davoyer@cisco.com

11. References

11.1. Normative References

- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC9800] Cheng, W., Ed., Filsfils, C., Li, Z., Decraene, B., and F. Clad, Ed., "Compressed SRv6 Segment List Encoding", RFC 9800, DOI 10.17487/RFC9800, June 2025, <<https://www.rfc-editor.org/info/rfc9800>>.

11.2. Informative References

- [RFC7404] Behringer, M. and E. Vyncke, "Using Only Link-Local Addressing inside an IPv6 Network", RFC 7404, DOI 10.17487/RFC7404, November 2014, <<https://www.rfc-editor.org/info/rfc7404>>.

[RFC9602] Krishnan, S., "Segment Routing over IPv6 (SRv6) Segment Identifiers in the IPv6 Addressing Architecture", RFC 9602, DOI 10.17487/RFC9602, October 2024, <<https://www.rfc-editor.org/info/rfc9602>>.

Authors' Addresses

Jakub Horn
Cisco Systems
Milpitas, CA 95035
United States of America
Email: jakuhorn@cisco.com

Kris Michielsen
Cisco Systems
De Kleetlaan 6a
1831 Diegem
Belgium
Email: kmichiel@cisco.com

Nicklous Morris
Verizon
United States of America
Email: nicklous.morris@verizonwireless.com

Martin Horneffer
Deutsche Telekom
Germany
Email: martin.horneffer@telekom.de

Clayton Hassen
Bell Canada
Vancouver
Canada
Email: clayton.hassen@bell.ca

Sen Li
China Mobile
Hong Kong SAR
China
Email: senli@cmi.chinamobile.com