

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 30 August 2025

Y. Liu
China Mobile
Y. Zhu
China Telecom
26 February 2025

IPv6 Address Assignment for SRv6
draft-liu-srv6ops-sid-address-assignment-01

Abstract

This document provides detailed SRv6 SID assignment considerations, which could be a comprehensive guide for optimizing SRv6 SID allocation in diverse deployment scenarios.

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] [RFC8174].

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 30 August 2025.

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. SRv6 SID Block Assignment	3
3. SRv6 SID Assignment for P2P and P2MP	4
4. SRv6 Node ID Assignment	4
4.1. Node ID for SRv6 Compression	4
4.1.1. Assignment for REPLACE-CSID	5
4.1.2. Assignment for NEXT-CSID	5
5. SRv6 Function ID Assignment	6
5.1. Function ID for SRv6 Compression	6
5.1.1. Function ID for REPLACE-CSID	6
5.1.2. Function ID for NEXT-CSID	6
5.2. Dynamic and Static Assignment	6
6. IANA Considerations	7
7. Security Considerations	7
8. Acknowledgements	7
9. References	7
9.1. Normative References	7
9.2. Informative References	7
Authors' Addresses	8

1. Introduction

[RFC8986](Section 3.2 "SID Allocation within an SR Domain") provides basic principle and practice for allocating SRv6 SIDs within SRv6 networks. These practices typically involve assigning large IPv6 prefixes to the SR domain and further subdividing them into smaller prefixes for individual nodes. This approach ensures efficient address utilization and simplifies network management. However, existing work primarily focuses on basic SRv6 deployments without considering the complexities introduced by advanced features like SRv6 compression and diverse service provider requirements. This document aims to bridge the gap by comprehensively detailing SRv6 SID assignment, which could provide service providers and network engineers with a comprehensive and practical guide for optimizing SRv6 SID allocation in diverse deployment scenarios.

2. SRv6 SID Block Assignment

Prior to SRv6, service providers typically allocate IPv6 addresses based on "administrative divisions" (e.g., province, city) and "network types" (e.g., IP network, wireless network, transport network). This approach assigns distinct unicast addresses (e.g., interface and loopback addresses) for network device. However, introducing SRv6 with independent allocations within each administrative division or network type creates challenges:

1) Fragmented SRv6 Space: Independent allocations result in scattered SRv6 address blocks across the provider's network, hindering SRv6 SID aggregation. Aggregation simplifies network management and allows efficient use of address space.

2) Edge Filtering Complexity: With fragmented SRv6 space, filtering SRv6 traffic at network edges becomes significantly more complex due to the dispersed nature of the addresses. This complicates network security and policy enforcement.

To address these challenges, it's recommended to allocate a "dedicated IPv6 address block" for SRv6 across the entire service provider network.

Example:

Scenario 1: Integrated SRv6 SIDs Planning in 1 Service Operator

Assume Block A:A:X:X::/24 is allocated for SRv6.

It only requests to apply a single policy to this prefix in edge configuration, such as: deny A:A:X:X::/24

Scenario 2: Separate SIDs Planning in different administrative division

Each administrative division has its own SRv6 block.

It requires multiple policies for each discrete prefix in edge configuration, such as:

deny A:B:X:X:C1:D:/48

deny A:B:Z:Z:C2:D:/48

...

deny A:B:X:X:Cn:D:/48

Here, C1...Cn represent n administrative divisions, and D represents the SRv6 SIDs allocated to each division.

The difference between A:A and A:B indicates whether ordinary IPv6 network addresses and SRv6 addresses are distinguished from the high-order address bits: A:A represents a separate SRv6 prefix, while A:B represents an ordinary IPv6 network address prefix.

It is showed that integrated SRv6 SIDs planning simplifies edge configuration by requiring only a single policy for the dedicated SRv6 prefix. This approach improves network management efficiency and reduces configuration complexity.

3. SRv6 SID Assignment for P2P and P2MP

Based on existing SR P2MP solutions, both SRv6 P2P and P2MP utilize unicast IPv6 addresses. While considering the distinction between the service of unicast and multicast, it may request separate address pools for SRv6 P2P and P2MP, for the following 2 aspects of considerations:

1) It's crucial to avoid allocating SRv6 SIDs for both P2P and P2MP connections under the same Node ID. This prevents address space contention and simplifies traffic management.

2) Independent Locator Advertisement for P2MP SIDs

So, it is suggested to dedicate distinct SID ranges or Node IDs for P2P and P2MP traffic flows within a service provider's network to ensure clear differentiation.

4. SRv6 Node ID Assignment

Node ID allocation could be flat or structured.

Flat allocation requests less resources but needs complex management, which is caused by the structural nature of administrative division in geography.

Node ID allocation could also be structured, enabling further administrative division refinement in SRv6 SIDs.

4.1. Node ID for SRv6 Compression

4.1.1.1. Assignment for REPLACE-CSID

If compressed and uncompressed Node IDs remain different:

- Benefit: Independent function planning becomes possible.
- Limitation: Separate Locator publication necessitates Node ID space wastage, which is not acceptable.

SoIf compressed and uncompressed Node IDs should remain consistent:

- Benefit: Only one Locator needs publication.
- Limitation: Available space for uncompressed Node IDs reduces proportionately.

It is necessary to mention that sharing the same Node ID for compressed and uncompressed functions inevitably leads to the joint occupation of the function space. This means that compressed and uncompressed functions must "share" this space, which can result in resource contention and limit the flexibility of function allocation.

Additionally, using the same Node ID indirectly restricts the Node ID length for non-compressed SIDs, considering the Node ID length limitation for compressed SIDs. This can be a significant constraint for networks that require a large number of non-compressed SIDs.

4.1.1.2. Assignment for NEXT-CSID

According to the definition of NEXT-CSID [I-D.ietf-spring-srv6-srh-compression], the argument value of the active SID is non-zero. If only a single Node ID is allocated to a node to support SRv6 compression and normal SRv6 without compression, it is difficult to distinguish in forwarding. It can be treated as a NEXT-CSID flavor END SID with non zero argument, but in the meantime, it is also associated with a 128-bit SRv6 SID.

Therefore, in order to support normal SRv6 and compressed SRv6, a node must be configured with two Node IDs. For the case of upgrading the network from SRv6 network to support NEXT-CSID, a new Node ID is required for each node.

As per [I-D.ietf-spring-srv6-srh-compression], a NEXT-CSID MUST be allocated from a Global ID Base(GIB). The space sharing between GIB and Local ID Base(LIB) depends on the configuration. Normally, operators can use the first nibble to distinguish the GIB space and LIB space. For example, 0x000-0x9FFF can be used for GIB, and 0xA000-0xFFFF can be used for LIB, where the 0-9 in the first nibble indicate the GIB, and A-F indicate the LIB space.

5. SRv6 Function ID Assignment

Function ID should have enough space for different functionalities and services.

5.1. Function ID for SRv6 Compression

5.1.1. Function ID for REPLACE-CSID

It should also be considered about balancing the length of Node ID and Function ID for Compressed SIDs. It means that due to the inherent length limitations of compressed SIDs, a trade-off must be made between the scope of manageable nodes and the range of network functions that each node can provide.

Even considering only path-related network functions like End and End.X, a certain amount of Function IDs(for example: End.X and links are related, with multiple Flavors: the number of Function IDs should at least exceed the link number * Flavor number) are required. Additional functions such as VPN services will further occupy the address space. Therefore, when planning SRv6 SID allocation, it is crucial to carefully consider the application scenarios and functional requirements of compressed SIDs to find the optimal trade-off solution.

5.1.2. Function ID for NEXT-CSID

As per [I-D.ietf-spring-srv6-srh-compression], a NEXT-CSID function ID is a fixed length ID, for example which can be 16 bits or 32 bits. As the function ID has the local significance on the node, therefore, it MUST be allocated from the LIB space to avoid mis-match in forwarding.

5.2. Dynamic and Static Assignment

There could be 2 types of Allocation:

1) Dynamic: Automatic device allocation; For example: End. X and VPN SIDs.

2) Static assignment: Manual configuration for easy management; It is suitable for functions with small number of SIDs. For example: End .

6. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

7. Security Considerations

TBD

8. Acknowledgements

TBD

9. References

9.1. Normative References

- [I-D.ietf-spring-srv6-srh-compression]
Cheng, W., Filsfils, C., Li, Z., Decraene, B., and F. Clad, "Compressed SRv6 Segment List Encoding (CSID)", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-srh-compression-23, 6 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-srh-compression-23>>.
- [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>>.

9.2. Informative References

- [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>>.

Authors' Addresses

Yisong Liu
China Mobile
China
Email: liuyisong@chinamobile.com

Yongqing Zhu
China Telecom
China
Email: zhuyq8@chinatelecom.cn