

Internet Engineering Task Force (IETF)
Request for Comments: 9602
Category: Informational
ISSN: 2070-1721

S. Krishnan
Cisco
October 2024

Segment Routing over IPv6 (SRv6) Segment Identifiers in the IPv6 Addressing Architecture

Abstract

Segment Routing over IPv6 (SRv6) uses IPv6 as the underlying data plane. Thus, Segment Identifiers (SIDs) used by SRv6 can resemble IPv6 addresses and behave like them while exhibiting slightly different behaviors in some situations. This document explores the characteristics of SRv6 SIDs and focuses on the relationship of SRv6 SIDs to the IPv6 Addressing Architecture. This document allocates and makes a dedicated prefix available for SRv6 SIDs.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9602>.

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
2. Terminology
3. SRv6 SIDs and the IPv6 Addressing Architecture
4. Special Considerations for Compressed SIDs
5. Allocation of a Prefix for SIDs
6. IANA Considerations
7. Security Considerations
8. References
 - 8.1. Normative References
 - 8.2. Informative References

1. Introduction

Segment Routing over IPv6 (SRv6) [RFC8754] uses IPv6 as the underlying data plane. In SRv6, SR source nodes initiate packets with a Segment Identifier (SID) in the Destination Address of the IPv6 header, and SR segment endpoint nodes process a local segment present in the Destination Address of an IPv6 header. Thus, SIDs in SRv6 can, and do, appear in the Destination Address of IPv6 datagrams by design. This document explores the characteristics of SRv6 SIDs and focuses on the relationship of SRv6 SIDs to the IPv6 Addressing Architecture [RFC4291]. This document allocates and makes a dedicated prefix available for SRv6 SIDs.

2. Terminology

The following terms are used as defined in [RFC8402].

- * Segment Routing (SR)
- * SR Domain
- * Segment
- * Segment Identifier (SID)
- * SRv6
- * SRv6 SID

The following terms are used as defined in [RFC8754].

- * Segment Routing Header (SRH)
- * SR Source Node
- * Transit Node
- * SR Segment Endpoint Node

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.

3. SRv6 SIDs and the IPv6 Addressing Architecture

[RFC8754] defines the Segment List of the SRH as a contiguous array of 128-bit IPv6 addresses; further, it states that each of the elements in this list are SIDs. But all of these elements are not necessarily made equal. Some of these elements may represent a local interface as described in Section 4.3 of [RFC8754] as "A FIB entry that represents a local interface, not locally instantiated as an SRv6 SID". It follows that not all the SIDs that appear in the SRH are SRv6 SIDs as defined by [RFC8402].

As stated above, the non-SRv6-SID elements that appear in the SRH SID list are simply IPv6 addresses assigned to local interfaces, and they need to conform to [RFC4291]. So, the following discussions are applicable solely to SRv6 SIDs that are not assigned to local interfaces.

One of the key questions to address is how these SRv6 SIDs appearing

as IPv6 Destination Addresses are perceived and treated by "transit nodes" (that are not required to be capable of processing a Segment or the Segment Routing Header).

Section 3.1 of [RFC8986] describes the format of an SRv6 SID as being composed of three parts, LOC:FUNCT:ARG, where a locator (LOC) is encoded in the L most significant bits of the SID followed by F bits of function (FUNCT) and A bits of arguments (ARG). If $L+F+A < 128$, the ARG is followed by enough zero bits to fill the 128-bit SID. Such an SRv6 SID is assigned to a node within a prefix defined as a Locator of length L. When an SRv6 SID occurs in the IPv6 Destination Address of an IPv6 header, only the longest matching prefix corresponding to the Locator [BCP198] is used by the transit node to forward the packet to the node identified by the Locator.

It is clear that this format for SRv6 SIDs is not compliant with the requirements set forth in [RFC4291] for IPv6 addresses, but it is also clear that SRv6 SIDs are not intended for assignment onto interfaces on end hosts. They look and act like other mechanisms that use IPv6 addresses with different formats, such as those described in "IPv6 Addressing of IPv4/IPv6 Translators" [RFC6052] and "An IPv6 Prefix for Overlay Routable Cryptographic Hash Identifiers Version 2 (ORCHIDv2)" [RFC7343].

While looking at the transit nodes, it becomes apparent that these addresses are used purely for forwarding and not for packet delivery to end hosts. Hence, the relevant specification to apply here is [BCP198], which requires implementations to support the use of variable-length prefixes in forwarding while explicitly decoupling IPv6 routing and forwarding from the IPv6 address/prefix semantics described in [RFC4291]. Please note that [BCP198] does not override the rules in [RFC4291]: it merely limits where their impact is observed.

Furthermore, in the SRv6 specifications, all SIDs assigned within a given Locator prefix are located inside the node identified by Locator. Therefore, there does not appear to be a conflict with Section 2.6.1 of [RFC4291] since subnet-router anycast addresses are neither required nor useful within a node.

4. Special Considerations for Compressed SIDs

[CSID] introduces an encoding for Compressed-SIDs (C-SIDs), and describes how to use a single entry in the Segment List as a container for multiple SIDs. A node taking part in this mechanism accomplishes this by using the ARG part [RFC8986] of the Destination Address of the IPv6 header to derive a new Destination Address. That is, the Destination Address field of the packet changes at a segment endpoint in a way similar to how the address changes as the result of processing a segment in the SRH.

One key thing to note here is that the Locator Block at the beginning of the address does not get modified by the operations needed for supporting C-SIDs. As we have established that the SRv6 SIDs are being treated simply as routing prefixes on transit nodes within the SR Domain, this does not constitute a modification to the IPv6 data plane on such transit nodes: any changes are restricted to SR-aware nodes.

5. Allocation of a Prefix for SIDs

All of the SRv6-related specifications discussed above are intended to be applicable to a contained SR Domain or between collaborating SR Domains. Nodes either inside or outside the SR Domains that are not SR-aware will not perform any special behavior for SRv6 SIDs and will treat them solely as IPv6 routing prefixes.

As an added factor of security, it is desirable to allocate some address space that explicitly signals that the addresses within that space cannot be expected to comply with [RFC4291]. As described in Section 3, there is precedent for mechanisms that use IPv6 addresses in a manner different from that specified in [RFC4291]. This would be useful in identifying and potentially filtering packets at the edges of the SR Domains to make it simpler for the SR Domain to fail closed.

At the time of writing, global DNS [RFC9499] SHOULD NOT reference addresses assigned from this block. Further specifications are needed to describe the conventions and guidelines for the use of this newly allocated address block. The SRv6 operational community, which is the first intended user of this block, is requested to come up with such conventions and guidelines in line with their requirements.

6. IANA Considerations

IANA has assigned the following /16 address block for the purposes described in Section 5 and recorded the allocation in the "IANA IPv6 Special-Purpose Address Registry" [SPECIAL] as follows:

Address Block:
5f00::/16

Name:
Segment Routing (SRv6) SIDs

RFC:
RFC 9602

Allocation Date:
2024-04

Termination Date:
N/A

Source:
True

Destination:
True

Forwardable:
True

Globally Reachable:
False

Reserved-by-Protocol:
False

7. Security Considerations

The security considerations for the use of Segment Routing [RFC8402], SRv6 [RFC8754], and SRv6 network programming [RFC8986] apply to the use of these addresses. The use of IPv6 tunneling mechanisms (including SRv6) also brings up additional concerns such as those described in [RFC6169]. The usage of the prefix allocated by this document improves security by making it simpler to filter traffic at the edge of the SR Domains.

In case the deployments do not use this allocated prefix, additional care needs to be exercised at network ingress and egress points so that SRv6 packets do not leak out of SR Domains and do not

accidentally enter SR-unaware Domains. Similarly, as stated in Section 5.1 of [RFC8754], the SR Domain needs to be configured to filter out packets entering that use the selected prefix.

8. References

8.1. Normative References

- [BCP198] Best Current Practice 198,
<<https://www.rfc-editor.org/info/bcp198>>.
At the time of writing, this BCP comprises the following:
- Boucadair, M., Petrescu, A., and F. Baker, "IPv6 Prefix Length Recommendation for Forwarding", BCP 198, RFC 7608, DOI 10.17487/RFC7608, July 2015,
<<https://www.rfc-editor.org/info/rfc7608>>.
- [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>>.
- [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>>.
- [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] Filtsfils, 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] Filtsfils, 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] Filtsfils, 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>>.

8.2. Informative References

- [CSID] Cheng, W., Filtsfils, C., Li, Z., Decraene, B., and F. Clad, "Compressed SRv6 Segment List Encoding", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-srh-compression-18, 22 July 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-srh-compression-18>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <<https://www.rfc-editor.org/info/rfc6052>>.
- [RFC6169] Krishnan, S., Thaler, D., and J. Hoagland, "Security Concerns with IP Tunneling", RFC 6169, DOI 10.17487/RFC6169, April 2011, <<https://www.rfc-editor.org/info/rfc6169>>.
- [RFC7343] Laganier, J. and F. Dupont, "An IPv6 Prefix for Overlay

Routable Cryptographic Hash Identifiers Version 2
(ORCHIDv2)", RFC 7343, DOI 10.17487/RFC7343, September
2014, <<https://www.rfc-editor.org/info/rfc7343>>.

[RFC9499] Hoffman, P. and K. Fujiwara, "DNS Terminology", BCP 219,
RFC 9499, DOI 10.17487/RFC9499, March 2024,
<<https://www.rfc-editor.org/info/rfc9499>>.

[SPECIAL] IANA, "IANA IPv6 Special-Purpose Address Registry",
<<https://www.iana.org/assignments/iana-ipv6-special-registry>>.

Acknowledgments

The author would like to extend a special note of thanks to Brian Carpenter and Erik Kline for their precisely summarized thoughts on this topic that provided the seed of this document. The author would also like to thank Andrew Alston, Fred Baker, Ron Bonica, Nick Buraglio, Bruno Decraene, Dhruv Dhody, Darren Dukes, Linda Dunbar, Reese Enghardt, Adrian Farrel, Clarence Filsfils, Jim Guichard, Joel Halpern, Ted Hardie, Bob Hinden, Murray Kucherawy, Cheng Li, Acee Lindem, Jen Linkova, Gyan Mishra, Yingzhen Qu, Robert Raszuk, Alvaro Retana, Michael Richardson, John Scudder, Petr Spacek, Mark Smith, Dirk Steinberg, Ole Troan, Eduard Vasilenko, ric Vyncke, Rob Wilton, Jingrong Xie, Chongfeng Xie, and Juan Carlos Zuniga for their ideas and comments to improve this document.

Author's Address

Suresh Krishnan
Cisco
Email: suresh.krishnan@gmail.com