

SPRING
Internet Draft
Intended status: Standards Track
Expires: September 03, 2025

L. Gong
China Mobile
C. Lin
New H3C Technologies
A. Lindem
LabN Consulting, L.L.C.
March 03, 2025

Advertise SRv6 Locator Information by IPv6 Neighbor Discovery
draft-gong-spring-nd-advertise-srv6-locator-01

Abstract

In an SRv6 network, each SRv6 segment endpoint has at least one SRv6 locator. Through the SRv6 locator routes, other SRv6 segment nodes can steer traffic to that node. This document describes a method of advertising SRv6 locator to neighboring SRv6 Segment Endpoint nodes through IPv6 Neighbor Discovery protocol.

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/lid-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 September 03, 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
1.1. Requirements Language.....	4
2. Terminology.....	4
3. Process of Advertising SRv6 Locator by IPv6 ND.....	5
4. Extension of IPv6 Neighbor Discovery Options.....	6
4.1. Source SRv6 Locator Information Option.....	6
4.2. Source SRv6 Capability Option.....	7
5. IANA Considerations.....	9
6. Security Considerations.....	9
7. Acknowledgements.....	9
8. References.....	9
8.1. Normative References.....	9
8.2. Informative References.....	10
Authors' Addresses.....	11

1. Introduction

Segment Routing (SR) [RFC8402] allows a node to steer a packet flow along any path. The headend is a node where the instructions for source routing (i.e., segments) are written into the packet. It hence becomes the starting node for a specific segment routing path. Intermediate per-path states are eliminated thanks to source routing. A Segment Routing Policy (SR Policy) [RFC8402] is an ordered list of segments (i.e., instructions) that represent a source-routed policy. The headend node is said to steer a flow into an SR Policy. The packets steered into an SR Policy have an ordered list of segments associated with that SR Policy written into them.

[RFC8402] defines an SRv6 Segment Identifier (SID) as an IPv6 address explicitly associated with the segment. When an SRv6 SID is in the Destination Address field of an IPv6 header of a packet, it is routed through transit nodes in an IPv6 network as an IPv6 address.

An SRv6 SID [RFC8986] is as consisting of 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). L, the locator length, is flexible, and an operator is free to use the locator length of their choice. F and A may be any value as long as $L+F+A \leq 128$. A locator may be represented as B:N where B is the SRv6 SID block (IPv6 prefix allocated for SRv6 SIDs by the operator) and N is the identifier of the parent node instantiating the SID. When the LOC part of the SRv6 SIDs is routable, it leads to the node, which instantiates the SID.

The SRv6 locator can be distributed to other IPv6 nodes within the SRv6 domain through IGP advertisement. This allows other nodes to learn the locator's route.

However, in the following two scenarios, the SRv6 endpoint can't advertise its SRv6 Locator information to other SRv6 nodes, resulting in other nodes being unable to dynamically generate routes to this node.

* Scenario 1: Deploying SRv6 on hosts.

As shown in Figure 1, the SRv6 network comprises Host, Access, Spin, and Core segments. Information is exchanged between Access and Spin through IGP, and between Spin and Core using BGP.

The SRv6 Locator information is transmitted between Access and Spin using the IGP protocol, and between Spin and Core using the BGP protocol. However, since Hosts generally do not support complex routing protocols, they cannot automatically transmit their SRv6 Locator information to the devices. This limitation complicates the deployment of SRv6 networks.




Figure 1

* Scenario 2: Static manual configuration of SRv6 Locator.

In IP network customer provider edge (CPE) devices often do not support an IGP protocol, which makes it impossible to advertise SRv6

locator routes for SRv6 Segment Endpoint Nodes through IGP. As shown in Figure 2, SIDs can only be configured manually on CPEs, and SRv6 Locator routes can only be statically distributed.

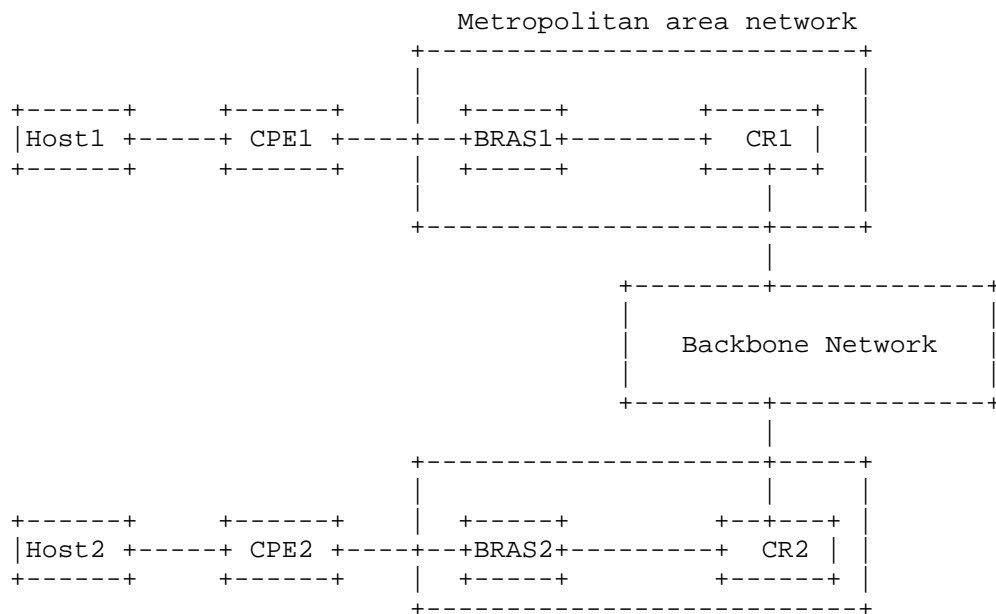


Figure 2

To address these issues, this document proposes a method of advertising SRv6 locators to neighboring SRv6 Segment Endpoint Nodes through IPv6 Neighbor Discovery packets.

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. Terminology

This document leverages the terms defined in [RFC4861] and [RFC8986]. The reader is assumed to be familiar with this terminology.

3. Process of Advertising SRv6 Locator by IPv6 ND

By extending the Neighbor Solicit (NS) and Neighbor Advertisement (NA) packets of the IPv6 Neighbor Discovery protocol to carry SRv6 locator information, SRv6 segment endpoint nodes can advertise their own SRv6 locator information to neighboring SRv6 nodes, and can also obtain the SRv6 locator information of neighboring SRv6 nodes, thereby achieving the exchange of SRv6 locator information and facilitating the deployment of SRv6.

Taking the scenario of deploying SRv6 on a host in Section 1 as an example, illustrate the process flow of exchanging SRv6 Locator information through IPv6 ND.

By extending the IPv6 ND protocol to carry SRv6 Locator information, Hosts and Access devices can exchange all SRv6 Locator information within an SRv6 network, facilitating the deployment of SRv6.

The srV6 locator are advertised in IPv6 ND NS and NA packets between the host and access, and are advertised via IGP within the domain. This information is collected by the controller using NETCONF or BGP-LS.

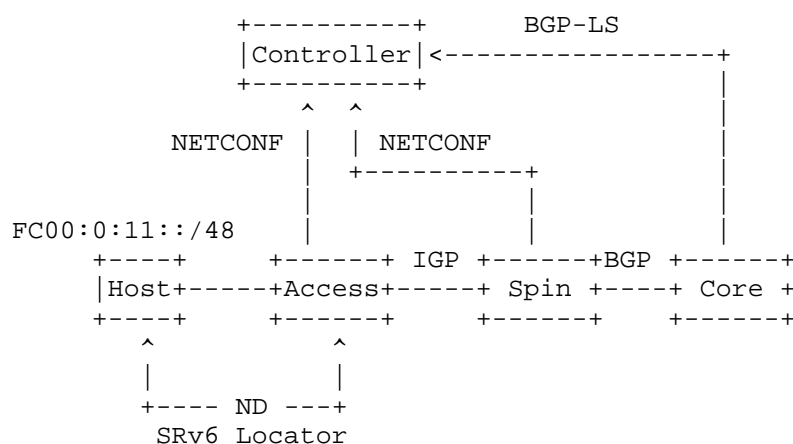


Figure 3

The process is as follows:


- 1) Manually configure SRv6 Locator statically on the host.
- 2) The Host advertises the SRv6 locator (For example, FC00:0:11::/48) and its SRv6 capability through IPv6 NS and NA packets.

- 3) The Access device learns the SRv6 locator (FC00:0:11::/48) via ND packets.
- 4) Using IGP, the Access device re-advertises the SRv6 locator (FC00:0:11::/48) obtained from ND packets to the Spine device.
- 5) The Spine device learns the SRv6 locator (FC00:0:11::/48) route through IGP.
- 6) A BGP neighbor relationship is established between the Spine and Core devices, and the learned SRv6 locator (FC00:0:11::/48) is sent to the Core using BGP-LS routes.
- 7) The Core device learns the SRv6 locator (FC00:0:11::/48) via BGP-LS.
- 8) A BGP neighbor relationship is established between the Core device and the controller, and the SRv6 locator (FC00:0:11::/48) is sent to the controller via BGP-LS routes.

4. Extension of IPv6 Neighbor Discovery Options

4.1. Source SRv6 Locator Information Option

The Source SRv6 Locator Information (SSLI) option is used to advertise the SRv6 locator information of the sending node to IPv6 neighbor.



Where:

- * Type: 8 bits, identifier of the type of option. The value is TBA by IANA.

- * Length: 8 bits, unsigned integer. The length of the option (including the type and length fields) in units of 8 octets. The value 0 is invalid. Nodes MUST silently discard an ND packet that contains an option with length zero.
- * R: 4 bits, reserved field. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
- * MTID: 12 bits, Multi-Topology Identifier.
- * Metric: 32 bits, Cost.
- * Flags: 8 bits, Reserved.
- * Algorithm: 8 bits Algorithm:
 - * - 0: Shortest Path First (SPF).
 - * - 1: Strict Shortest Path First (Strict SPF).
- * LOC Size: 8 bits, Locator Length.
- * Locator (variable): Variable length, indicates the advertised SRv6 Locator.
- * Sub-TLVs (variable): Variable length, contains Sub-TLVs such as SRv6 End SID Sub-TLV.


The option only may appear in Neighbor Solicit message and Neighbor Advertisement message.

Neighbor Solicit message and Neighbor Advertisement message can include zero or more Source SRv6 Locator Information options. If multiple SRv6 Locators need to be advertised, multiple Source SRv6 Locator Information options MUST be encapsulated in the same Neighbor Discovery message. The Source SRv6 Locator Information Option should be padded when necessary to ensure that it end on its natural 64-bit boundary.

Receivers MUST silently ignore the option if they can't recognize and continue processing the message.

4.2. Source SRv6 Capability Option

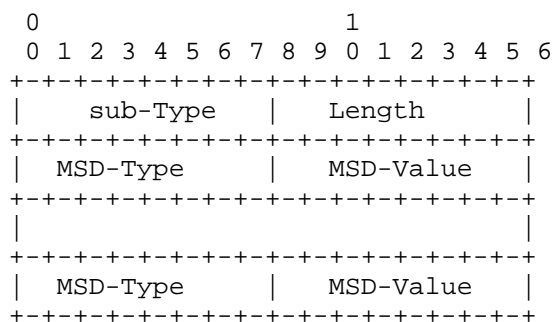
To support the SRv6 functionality, the source node also needs to advertise SRv6-related capabilities through IPv6 ND packets. The Source SRv6 Capability (SSC) option is used to advertise the SRv6 capability information of the sending node to IPv6 neighbor.



where:

- * Type: 8 bits, identifier of the type of option. The value is TBA by IANA.
- * Length: 8 bit, unsigned integer. The length of the option (including the type and length fields) in units of 8 octets. The value 0 is invalid. Nodes MUST silently discard an ND packet that contains an option with length zero.
- * Flags: 8 bits, Reserved.
- * Sub-TLVs: Variable length, contains Sub-TLVs such as MSD sub-TLV.

For optional MSD sub-sub-TLVs, the format is as follows:



where:

- * sub-Type: 8 bits, The sub-TLV Type value for MSD is TBA.
- * Length: 8 bits, variable.
- * MSD-Type: 8 bits.
- * MSD-Value: 8 bits.

The option only may appear in Neighbor Solicit message and Neighbor Advertisement message.

Neighbor Solicit message and Neighbor Advertisement message can only include a maximum of one Source SRv6 Capability option. The Source SRv6 Capability Option should be padded when necessary to ensure that it end on its natural 64-bit boundary.

Receivers MUST silently ignore the option if they can't recognize and continue processing the message.

5. IANA Considerations

IANA is asked to assign a new value for the "IPv6 Neighbor Discovery Option Formats" registry under the heading "Internet Control Message Protocol version 6 (ICMPv6) Parameters", as follows:

Value	Description	Reference
TBA	Source SRv6 Locator Information Option	This document
TBA	Source SRv6 Capability Option	This document

Table 1

6. Security Considerations

See Section 11 of [RFC4861] for the Neighbor Discovery security considerations.

7. Acknowledgements

TBD

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017.

[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>>.

8.2. Informative References

TBD

Authors' Addresses

Liyan Gong
China Mobile
China

Email: gongliyan@chinamobile.com

Changwang Lin
New H3C Technologies
China

Email: linchangwang.04414@h3c.com

Acee Lindem
LabN Consulting, L.L.C.
301 Midenhall Way
Cary, NC, 95110
United States of America
Email: acee.ietf@gmail.com

