SPRING
Internet-Draft
Intended status: Standards Track
Expires: 29 August 2026

L. Gong
China Mobile
C. Lin
New H3C Technologies
A. Lindem
Arrcus, Inc.

25 February 2026

Advertise SRv6 Locator Information by IPv6 Neighbor Discovery
draft-gong-spring-nd-advertise-srv6-locator-04

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 for an SRv6 endpoint (e.g., a host or a customer provider edge (CPE)) to advertise its SRv6 locator to a neighboring SRv6-aware router using extensions to the IPv6 Neighbor Discovery (ND) protocol. This approach eliminates the need to run a full routing protocol stack on simple endpoints, facilitating SRv6 deployment in controlled, trusted domains such as data centers and managed access networks.

Table of Contents

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.

The network programming paradigm for SRv6 is specified in [RFC8986]. It describes how any behavior can be bound to a SID and how any network program can be expressed as a combination of SIDs. It also describes several well-known behaviors that can be bound to SRv6 SIDs.

In an SRv6 network, each SRv6 Segment Endpoint Node must be assigned an SRv6 Locator, and segment IDs are generated within the address space of this SRv6 Locator. This document describes a method of advertising SRv6 locator to neighboring SRv6 Segment Endpoint nodes through IPv6 Neighbor Discovery protocol.

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 [RFC2119] (Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997) and [RFC8174] (Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017).

1.2. Terminology

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

2. Motivation and Applicability

The SRv6 locator is typically distributed within a network domain using IGP or BGP protocols, enabling other nodes to steer traffic towards the endpoint. However, there are deployment scenarios where an SRv6 segment endpoint (e.g., a server host, a lightweight CPE device) cannot or should not run a complex routing protocol suite. Manually configuring static routes for each such endpoint on the network side is operationally burdensome and does not scale. The following two typical scenarios illustrate this challenge:

* Scenario 1: Deploying SRv6 on Hosts in a Data Center

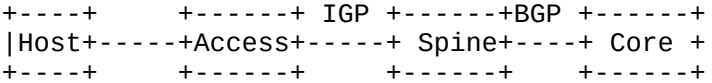

Figure 1

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

The SRv6 Locator information is transmitted between Access and Spine using the IGP protocol, and between Spine 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.

* Scenario 2: Static Manual Configuration at Customer Edges

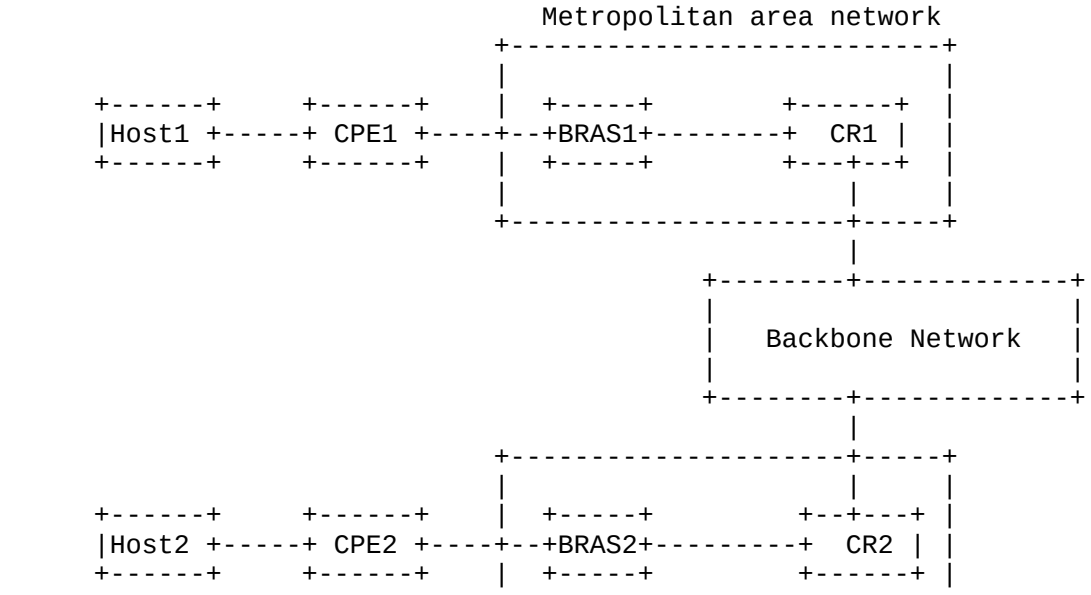

Figure 2

In IP networks, 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, leading to operational overhead.

This document addresses this gap by specifying a lightweight method for SRv6 locator advertisement using the IPv6 Neighbor Discovery protocol.The key advantages are:

o Zero Additional Protocol Stack: Leverages the existing, mandatory IPv6 ND stack present on all IPv6 nodes, imposing no new protocol implementation burden on simple endpoints.

o Operational Simplicity: Allows network operators to bring SRv6 endpoints into the routing domain without configuring routing protocols or numerous static routes on the access router.

o Scalable Signaling: ND provides an efficient, link-local signaling mechanism suitable for a large number of endpoints on a shared segment.

Applicability Scope: The mechanism defined in this document is designed for use within a single, trusted, and administratively controlled SR domain, as illustrated in the scenarios above. Primary use cases include:

o Data center networks (Figure 1) where servers (hosts) are under the control of the same operator as the network fabric.

o Managed access or enterprise networks (Figure 2) where CPE devices are provisioned and secured by the network operator.

It is NOT RECOMMENDED for use in environments where attached nodes are untrusted (e.g., general public Internet access).

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 locators 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.

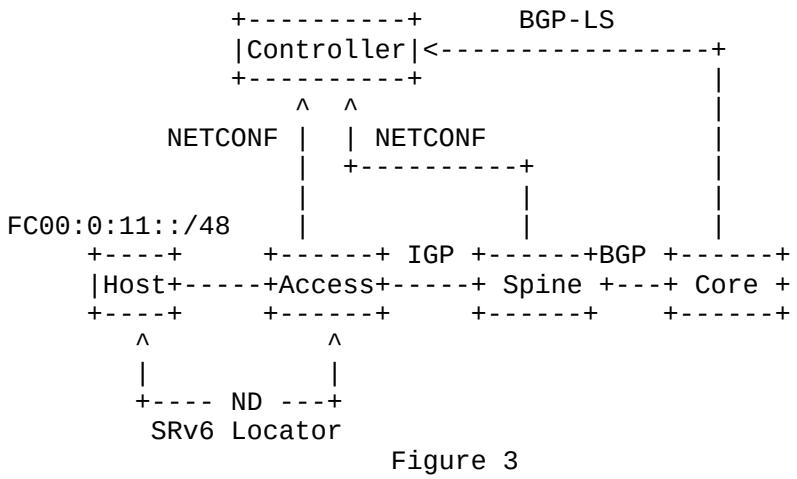

Figure 3

The process, illustrated in Figure 3, is as follows: (1)An SRv6 locator is configured on the Host.

(2)The Host advertises its SRv6 locator(FC00:0:11::/48) and capability in IPv6 ND packets (NS/NA) to the Access router.

(3)The Access router receives the ND packets, extracts and validates the SRv6 locator information(FC00:0:11::/48).

(4)Using IGP, the Access router re-advertises the learned SRv6 locator to the Spine device.To enhance scalability in deployments with a large number of endpoints, the Access router MAY aggregate multiple fine-grained locators (e.g., into a shorter prefix) before advertising.This design ensures that the fine-grained locator information from many endpoints is contained at the edge, while the core routing infrastructure deals only with aggregated prefixes.

(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.

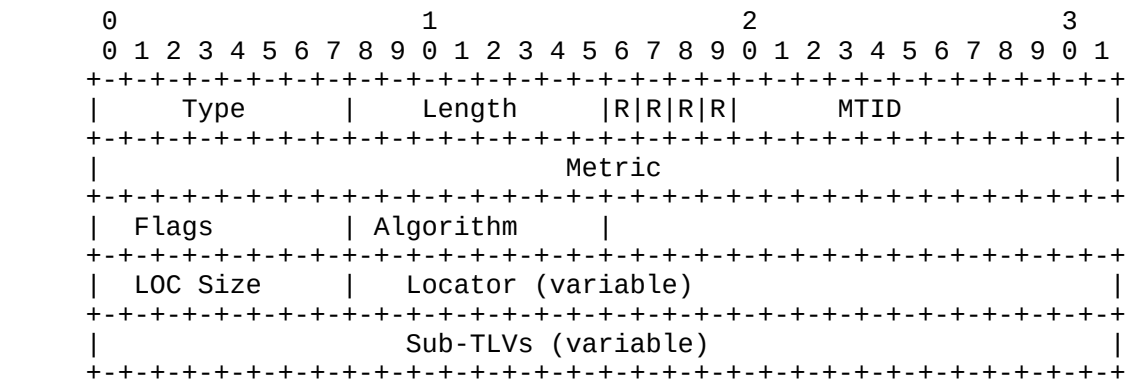

Where:

* 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.

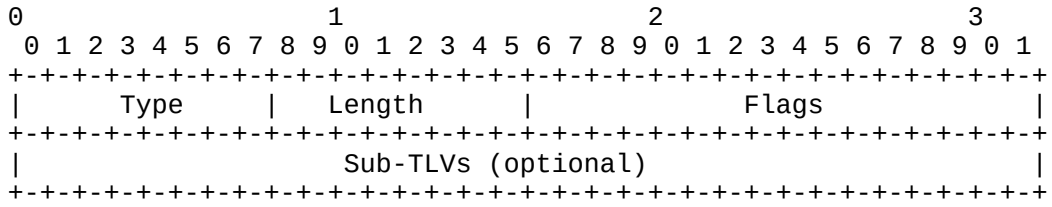

where:

* 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:

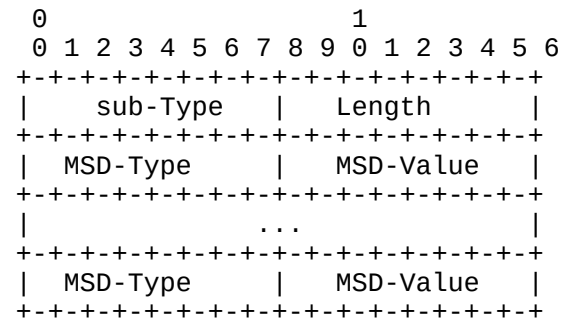

where:

* 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

The security considerations of IPv6 Neighbor Discovery [RFC4861] fully apply. This specification introduces new ND options that carry routing information (SRv6 locators), and therefore MUST be used within a well-defined trust boundary. The mechanism defined in this document is designed for use within a single, trusted, and administratively controlled SRv6 domain. This domain is a network segment or realm where:

o All nodes (including hosts or CPEs that advertise SRv6 locators) are under the control of a single administrative authority.

o The software and configuration of these endpoints are authorized and managed by the network operator (e.g., they are considered managed network edge devices, not untrusted end-user devices).

o Access to the link layer is controlled to prevent unauthorized nodes from attaching.

Primary applicable scenarios include operator-controlled data centers (Scenario 1 in Section 2) and managed access networks (Scenario 2 in Section 2).It is NOT RECOMMENDED for use in environments where attached nodes are not trusted (e.g., general Internet access points).

7. Normative 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/rfc/rfc2119>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/rfc/rfc4861>.

[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/rfc/rfc8174>.

[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/rfc/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/rfc/rfc8986>.

Authors' Addresses

Liyan Gong
China Mobile
China
Email: gongliyan@chinamobile.com

Changwang Lin
New H3C Technologies
China
Email: linchangwang.04414@h3c.com

Acee Lindem
Arrcus, Inc.
United States of America
Email: acee.ietf@gmail.com