SPRING
Internet Draft
Intended status: Standards Track
Expires: June 05, 2025

W. Cheng
R. Han
China Mobile
C. Lin
Y. Qiu
New H3C Technologies

December 04, 2025

Distribute SRv6 Locator by IPv6 Stateless Address Autoconfiguration
draft-cheng-spring-stateless-nd-srv6-locator-02

Abstract

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 for assigning SRv6 locators to SRv6 Segment Endpoint Nodes through IPv6 stateless address autoconfiguration.

Status of this Memo

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.

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 <= 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. The SRv6 Segment Endpoint Node then allocates SIDs with various behaviors based on its 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. In such scenarios, SIDs can only be configured manually on CPEs, and SRv6 Locator routes can only be statically distributed.

To address this issue, this document proposes a method of dynamically advertising SRv6 locators to SRv6 Segment Endpoint Nodes through IPv6 stateless address configuration method. It follows the existing process of IPv6 stateless address configuration, simplifying the allocation of SRv6 locators and route distribution.

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. Scenario for SRv6 Locator

The application scenario for obtaining SRv6 Locator through IPv6 stateless address autoconfiguration is similar to that of [I-D.ietf-spring-dhc-distribute-srv6-locator-dhcp].

In the IP backbone network, Telecom providers can use its IP Metro and Backbone networks to establish connectivity between access users who are located in different regions.

As shown in Figure 1, access network devices (CPE) are deployed for access users in different regions. This deployment assumes that all of the relevant components in Figure 1 are part of a single trusted SR domain. The CPE must be operator-managed and is only applicable when different arms of the same company operate their portions of the network separately, but must trust each other.

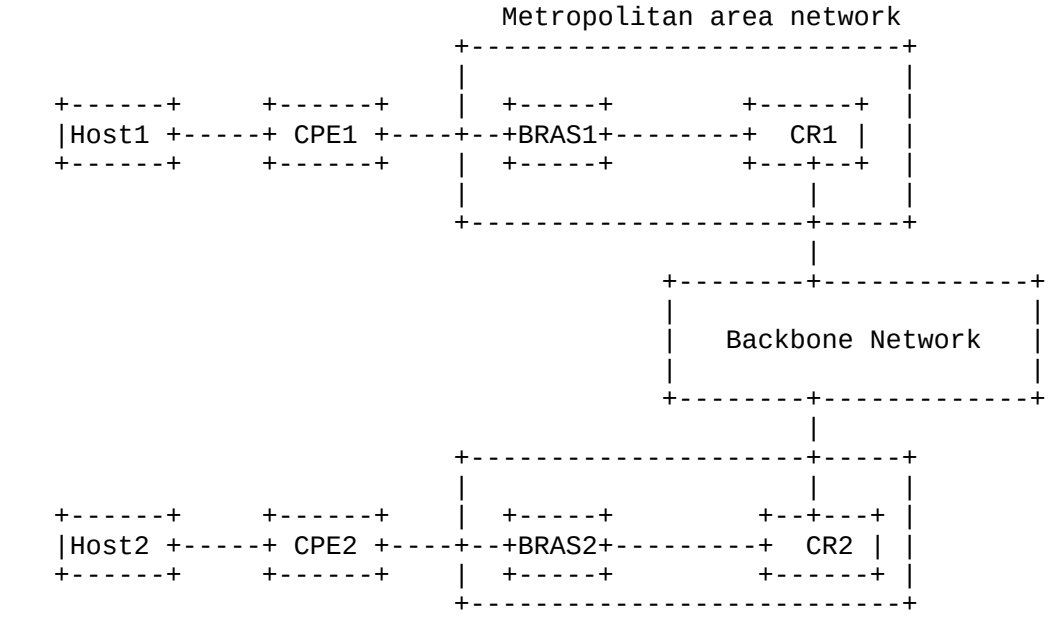

Figure 1: Telecom IPv6 Network

CPEs for access users are connected to the local metropolitan area network (MAN) in various ways. CPEs are responsible for assigning addresses to access users, so CPEs usually apply for IPv6 subnet prefix through DHCPv6 or stateless address autoconfiguration from BRAS.

In this network, operators hope to achieve interconnection between access users through End-to-End SRv6 tunnels. Taking the service traffic from Host1 to Host2 as an example, CPE1 is the SRv6 ingress node and CPE2 is the SRv6 egress node. The SRv6 locator should be configured on CPE. Other devices in the network learn the SRv6 locator route of the CPE.

At the same time, SRv6 policies needs to be configured on CPEs to steer the service traffic between CPEs to the specified SRv6 forwarding path. The SRv6 policy can be manually configured statically or issued through the controller, and its specific configuration method is out of the scope of this document.

However, in Metro network, the number of CPEs is very large and widely distributed geographically. Moreover, the mobility requirements of CPE are relatively high, and the access location of the same CPE often changes, so its IPv6 address cannot be fixed.

At present, an SRv6 locator can only be configured on each CPE through a controller or the Command Line Interface (CLI), which increases the configuration complexity.

To solve the difficulties this document proposes a method to allocate SRv6 locators to CPE through IPv6 stateless address autoconfiguration.

4. Extension of IPv6 Neighbor Discovery Options

The SRv6 Locator option is used to specify the SRv6 locators that are used for stateless address autoconfiguration.

The terms Locator Block and Locator Node correspond to the B and N parts, respectively, of the SRv6 Locator that is defined in Section 3.1 of [RFC8986].

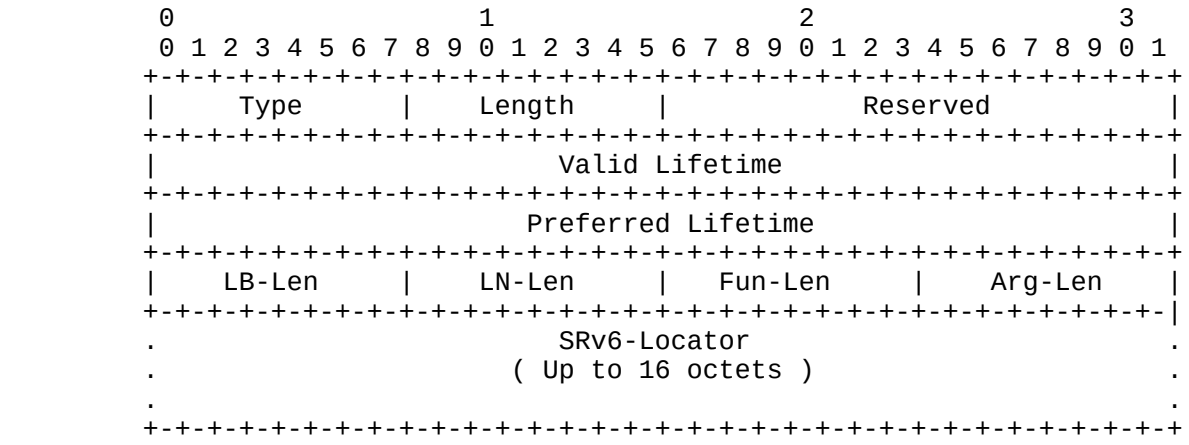

Fields:

* Type: 8-bit 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.

* Reserved: 16-bit unused field. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

* Valid Lifetime: 32-bit unsigned integer. The valid lifetime for the SRv6 locator in the option, expressed in units of seconds

(relative to the time the packet is sent). A value of all one bits (0xffffffff) represents infinity.

* Preferred Lifetime: 32-bit unsigned integer. The preferred lifetime for the SRv6 locator in the option, expressed in units of seconds (relative to the time the packet is sent). A value of all one bits (0xffffffff) represents infinity. Note that the value of this field MUST NOT exceed the Valid Lifetime field to avoid SRv6 locator that are no longer valid.

* LB-Len: 8-bit unsigned integer. SRv6 SID Locator Block (LB) length in bits.

* LN-Len: 8-bit unsigned integer. SRv6 SID locator Node (LN) length in bits.

* Fun-Len: 8-bit unsigned integer. SRv6 SID function (FUNCT) length in bits.

* Arg-Len: 8-bit unsigned integer. SRv6 SID arguments (ARG) length in bits.

* SRv6-Locator: 1-16 octets. This field encodes the SRv6 Locator. The SRv6 Locator is encoded in the minimal number of octets for the given number of bits. Trailing bits MUST be set to zero and ignored when received.

The option only may appear in the Router Advertisement message.

Router Advertisement messages can include zero or more SRv6 Locator options. If multiple SRv6 Locators need to be advertised to the same device, multiple SRv6 Locator options MUST be encapsulated in the same Router Advertisement message. The SRv6 Locator 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. Process of Advertising SRv6 Locator by Router Advertisement

This section describes router and host behavior of adverting SRv6 locator related to the Router Discovery portion of Neighbor Discovery. Router Discovery is used to locate neighboring routers as well as learn prefixes and configuration parameters related to stateless address autoconfiguration.

The IPv6 stateless autoconfiguration mechanism requires no manual configuration of hosts, minimal (if any) configuration of routers, and no additional servers. The stateless mechanism allows a host to generate its own addresses using a combination of locally available information and information advertised by routers. Routers advertise prefixes that identify the subnet(s) associated with a link, while hosts generate an "interface identifier" that uniquely identifies an interface on a subnet. An address is formed by combining the two. In the absence of routers, a host can only generate link-local addresses. However, link-local addresses are sufficient for allowing communication among nodes attached to the same link.

The stateless approach is used when a site is not particularly concerned with the exact addresses hosts use, so long as they are unique and properly routable.

Global addresses generated through stateless autoconfiguration mechanism are formed by appending an interface identifier to a prefix of appropriate length. Prefixes are obtained from Prefix Information options contained in Router Advertisements.

Router Advertisements are sent periodically to the all-nodes multicast address. To obtain an advertisement quickly, a host sends out Router Solicitations as described in [RFC4861].

When the router sends a Router Advertisement, it carries the SRv6 locator prefix information assigned to the host in the SRv6 Locator Option. After receiving the Router Advertisement, the host extracts the SRv6 Locator Option, obtains the SRv6 Locator prefix.

The detailed process of routers and hosts when advertising SRv6 locators during the IPv6 stateless autoconfiguration is as follows.

5.1. Router Behavior

The Router follows the specifications in Section 6.2 of [RFC4861] to send out Router Advertisement messages periodically, or in response to Router Solicitations.

When receiving a Router Solicitation from a host, if the router has already assigned an SRv6 locator to the host, it will include the SRv6 Locator Option in the Router Advertisement message of the responding host and advertise the SRv6 locator to the host.

In the scenario where all hosts under the advertising interface share the SRv6 Locator prefix, the SRv6 Locator option SHOULD be included in the Router Advertisement message sent periodically.

5.2. Host Behavior

On receipt of a Router Advertisement, the host follows the specifications in Section 6.3 of [RFC4861] to verify the packet. For valid Router Advertisement, continue to extract the SRv6 Locator option.

For each SRv6 Locator option, a host does the following:

* If the SRv6 locator is not already present in the SRv6 Locator List, and the SRv6 Locator Option’s Valid Lifetime field is non-zero, create a new entry for the SRv6 locator and initialize its invalidation timer to the Valid Lifetime value in the SRv6 Locator option.

* If the SRv6 locator is already present in the host’s SRv6 Locator List as the result of a previously received advertisement, reset its invalidation timer to the Valid Lifetime value in the SRv6 Locator option. If the new Lifetime value is zero, time-out the SRv6 locator immediately.

* If the SRv6 Locator option’s Valid Lifetime field is zero, and the SRv6 locator is not present in the host’s SRv6 Locator List, silently ignore the option.

Whenever the invalidation timer expires for a SRv6 Locator entry, that entry is discarded.

After processing the SRv6 Locator Option, the host records the SRv6 Locator prefix, and generates the corresponding SID based on the local configuration. The method of generating SID based on SRv6 locator is out of the scope of this document.

6. 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	SRv6 Locator Option	This document

Table 1

7. Security Considerations

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

8. Acknowledgements

TBD

9. References

9.1. 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>.

[I-D.ietf-spring-dhc-distribute-srv6-locator-dhcp] Cheng, W., Han, R., Lin, C., Qiu, Y., Zhang, G., "Distribute SRv6 Locator by DHCP", draft-ietf-spring-dhc-distribute-srv6-locator-dhcp-03 (work in progress), June 2024.

9.2. Informative References

TBD

Authors' Addresses

Weiqiang Cheng
China Mobile

Email: chengweiqiang@chinamobile.com

Ruibo Han
China Mobile

Email: hanruibo@chinamobile.com

Changwang Lin
New H3C Technologies

Email: linchangwang.04414@h3c.com

Yuanxiang Qiu
New H3C Technologies

Email: qiuyuanxiang@h3c.com