BESS
Internet Draft
Intended status: Standards Track
Expires: September 18, 2025

Yisong Liu
China Mobile
C. Lin
New H3C Technologies
Y. Liu
ZTE

March 18, 2025

SRv6 Service SID Anycast Flag
draft-liu-bess-srv6-service-sid-anycast-flag-02

Abstract

In some multihoming SRv6 L3VPN and EVPN scenarios, there are requirements for the egress PE to advertise both unicast and anycast SRv6 Service SIDs for the same service. This document defines the Anycast-flag for SRv6 Service SIDs carried in BGP messages.

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/1id-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 18, 2025.

Copyright Notice

Table of Contents

1. Introduction

[RFC9252] defines procedures and messages for SRv6-based BGP services, including Layer 3 Virtual Private Network (L3VPN), Ethernet VPN (EVPN), and Internet services. In some multihoming scenarios, there are requirements for the egress PE to advertise both unicast and anycast SRv6 Service SIDs for the same service. And those anycast SIDs need to be identified in the BGP messages.

This document defines the Anycast-flag for SRv6 Service SIDs carried in BGP messages.

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. Anycast Service SID

In some multihoming scenarios, there are requirements for the egress PE to advertise both unicast and anycast SRv6 Service SIDs for the

same service. It is required to identify which Service SID is anycast and which Service SID is unicast, when both two SIDs are advertised in BGP messages.

IGP has Anycast-flag for SRv6 locator, but the IGP Anycast-flag can be lost due to summarization. This document defines the Anycast-flag for SRv6 Service SIDs carried in BGP messages. Below are two application scenarios where both anycast service SID and unicast service SID are advertised simultaneously.

2.1. Use Case 1

In the multihoming SRv6 L3VPN and EVPN scenarios, anycast Service SID may be used to advertise the same service at different egress PEs, which can improve service reliability and load balancing.

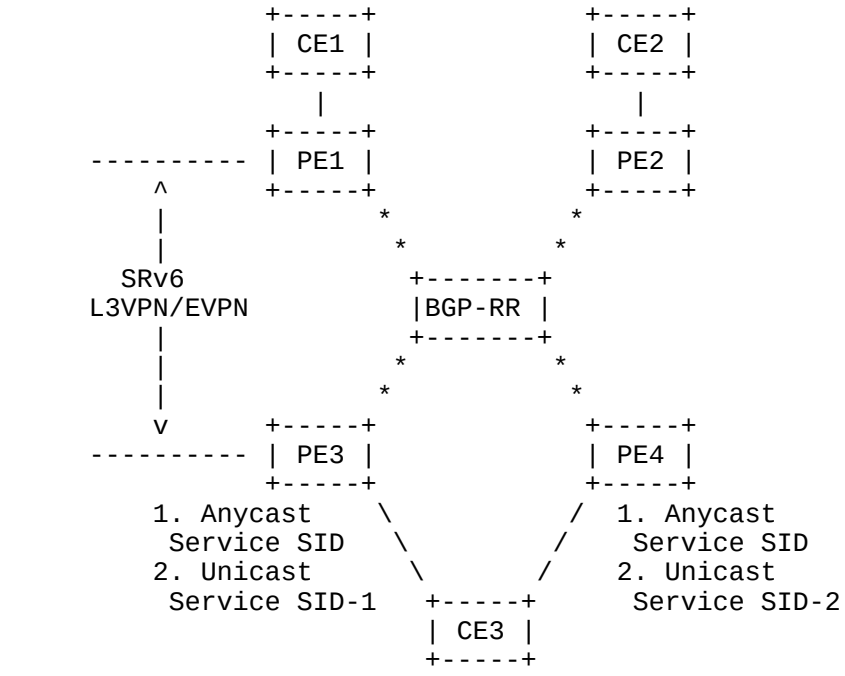

PE1:
  VPN Traffic Policy:
    PE3 & PE4 Load Balancing
  FIB Entry for VPN Traffic:
    Next-hop: Anycast Service SID

PE2:
  VPN Traffic Policy:
    PE3 Active, PE4 Backup
  FIB Entry for VPN Traffic:
    Primary Next-hop: Unicast Service SID-1
    Backup Next-hop: Unicast Service SID-2

Figure 1

As shown in Figure 1, PE3 and PE4 use the same anycast SRv6 Service SID for the VPN service of CE3. The ingress PE1 encapsulates the payload in an outer IPv6 header where the destination address is that anycast SRv6 Service SID. The packets from CE1 can reach CE3 through either PE3 or PE4. Assume that the path from PE1 to PE3 and the path from PE1 to PE4 have the same cost. The traffic flows will be load balanced between PE3 and PE4.

PE3 and PE4 also have unicast SRv6 Service SIDs, which are SID-1 and SID-2, for the VPN service of CE3. The ingress PE2 uses SID-1 as the primary SRv6 Service SID, and SID-2 as backup. The packets from CE2

will be forwarded to CE3 through PE3. If any failure occurs on the path to PE3, service will be switched to PE4.

Since ingress PE1 and PE2 have different strategies for the control of VPN traffics, egress PE3 and PE4 each need to advertise two SRv6 Service SIDs, an anycast SID for ingress PE1 and a unicast SID for ingress PE2. Local export policy may be used by egress PE3 and PE4 to control which SID is advertised to ingress PE1 and which is advertised to ingress PE2. However, if BGP Route Reflector is deployed, both the anycast Service SID and the unicast Service SID will be advertised to RR and reflected to ingress PEs, and the receiver has to choose which Service SID to use.

2.2. Use Case 2

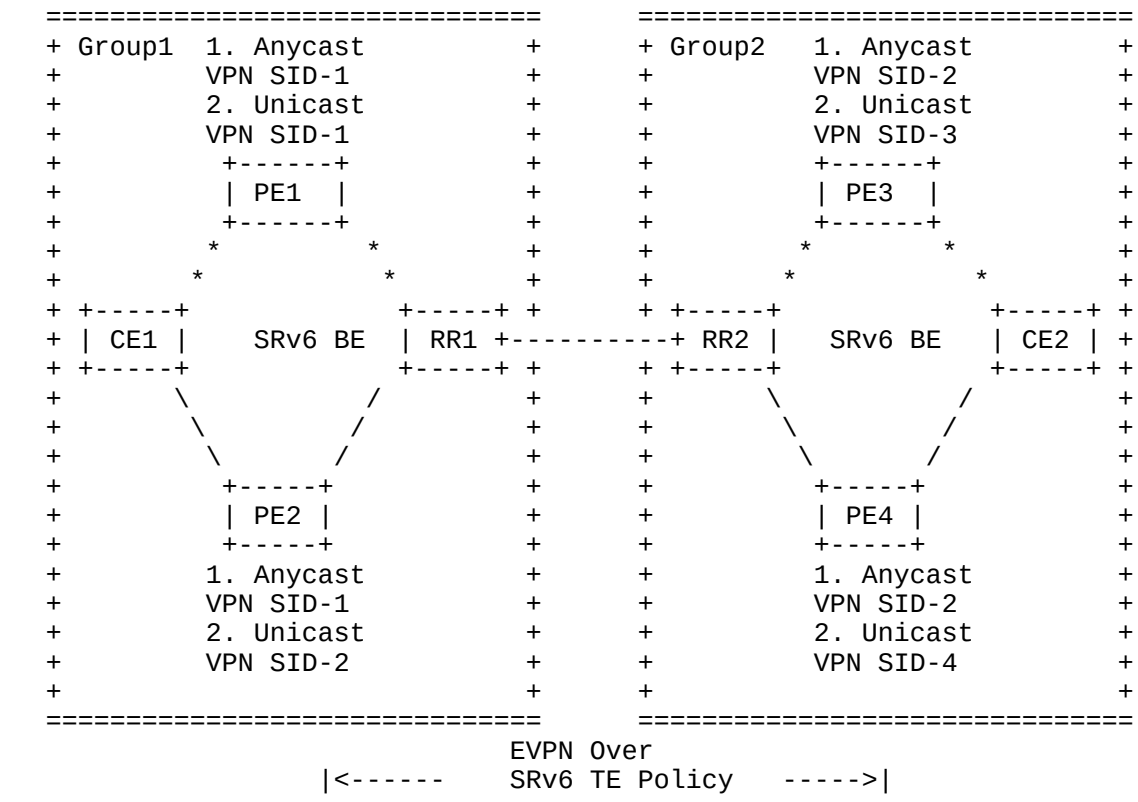

Figure 2

PE1 and PE2 belong to Group 1 and use the same Anycast IP 1. PE3 and PE4 belong to Group 2 and use the same Anycast IP 2. PEs from different groups use Anycast IP 1 and Anycast IP 2 as tunnel head nodes to deploy SRv6 TE policies, reducing the number of SRv6 TE

policies. Each PE device assigns two SRv6 VPN SIDs for the same VPN service: Anycast VPN SID and Unicast VPN SID. Anycast Service SIDs are used for forwarding between different groups. Within the same group, Unicast Service SIDs are used for forwarding between Multi-homed PE devices.

3. Extensions for BGP

This document defines a new flag in the SRv6 Service SID Flags field of the SRv6 SID Information Sub-TLV [RFC9252]:

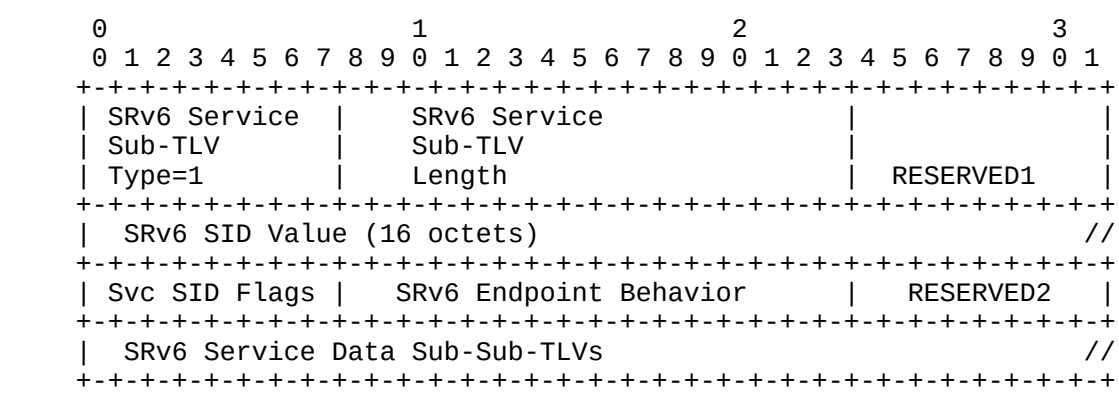

Svc SID Flags:

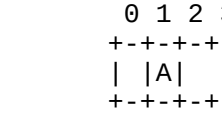

o A-flag: Anycast flag. When set, the associated SID is anycast.

The new-defined flag can be used for the SRv6 Service SIDs of L3 and L2 services, such as End.DX4, End.DT4, End.DX6, End.DT6, End.DT46. End.DX2, End.DX2V, End.DT2U, etc.

4. Backward Compatibility

According to [RFC9252],

o Any unknown flags in the SRv6 Service SID Flags field MUST be ignored by the receiver.

- When multiple SRv6 SID Information Sub-TLVs are present, the ingress PE SHOULD use the SRv6 SID from the first instance of the Sub-TLV.

When the egress PE advertises multiple service SIDs, the unicast service SID SHOULD be carried in the first instance of Sub-TLV. If there are PE routers not supporting the flag defined in this document, the egress PE MAY expect those routers to use the first SID and ignore the new-defined flag.

5. Security Considerations

TBD.

6. IANA Considerations

This document defines the following bit in the SRv6 Service SID Flags field of SRv6 SID Information Sub-TLV:

TLV Code Point	Value
TBD	A-flag

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017

[RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <https://www.rfc-editor.org/info/rfc9252>.

Authors' Addresses

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

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

Yao Liu
ZTE
China
Email: liu.yao71@zte.com.cn