

BESS
Internet Draft
Intended status: Standards Track
Expires: October 27, 2026

Yisong Liu
China Mobile
C. Lin
New H3C Technologies
Y. Liu
ZTE
J. Rabadan
Nokia
April 25, 2026

SRv6 Anycast VPN Service
draft-liu-bess-srv6-anycast-vpn-service-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 anycast flavor 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/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 October 27, 2026.

Copyright Notice

Copyright (c) 2026 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction.....	2
1.1. Requirements Language.....	2
2. Anycast Service SID.....	3
2.1. Use Case 1.....	3
2.2. Use Case 2.....	6
3. Solution.....	6
4. Extensions for SRv6 Endpoint Behaviors.....	8
4.1. End.DT4.Anycast : End.DT4 with Anycast.....	8
4.2. End.DT6.Anycast: End.DT6 with Anycast.....	8
4.3. End.DT46.Anycast : End.DT46 with Anycast.....	9
4.4. End.DX4.Anycast: End.DX4 with Anycast.....	9
4.5. End.DX6.Anycast: End.DX6 with Anycast.....	9
4.6. End.DT2U.Anycast : End.DT2U with Anycast.....	9
4.7. End.DX2.Anycast : End.DX2 with Anycast.....	10
5. Security Considerations.....	10
6. IANA Considerations.....	10
7. References.....	11
7.1. Normative References.....	11
7.2. Informative References.....	12
Authors' Addresses.....	12

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 behavior for SRv6 Service SIDs carried in BGP update messages.

1.1. Requirements Language

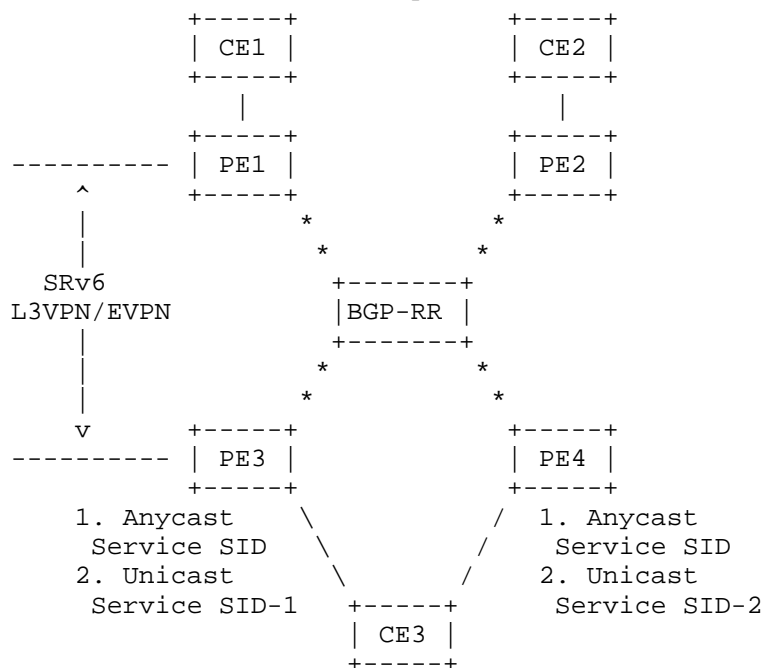
The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and

"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

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. For VPN-IPv4 and VPN-IPv6 BGP SAFIs [RFC4364] or EVPN IP Prefix routes [RFC9136][I-D.rabnag-bess-evpn-anycast-aliasing], traffic flows from PE1 are load-balanced between PE3 and PE4, assuming the paths from PE1 to PE3 and PE4 to each have the same cost. When EVPN [RFC7432][RFC9252] is used, and PE3 and PE4 are attached to the same Ethernet Segment operating in Anycast

multi-homing mode [I-D.rabnag-bess-evpn-anycast-aliasing], unicast traffic is similarly load balanced across 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. In this example, the operator prefers that PE2 sends traffic only to PE3, with PE4 acting as a backup. For VPN-IPv4/VPN-IPv6 or EVPN IP Prefix routes, the ingress PE2 uses SID-1 as the primary SRv6 Service SID and SID-2 as the backup. Similarly, for EVPN, if CE3's Ethernet Segment operates in single-active mode and PE3 is the Designated Forwarder (DF) [RFC7432], PE2 uses SID-1 as the primary path and SID-2 as the backup path to reach CE3. In the event of a failure on PE3, traffic is automatically switched to the path to PE4.

Since ingress PE1 and PE2 have different strategies for the control of VPN traffic, 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.

When both SIDs are advertised in BGP update messages, it is necessary to distinguish which Service SID is anycast and which is unicast.

IGP has Anycast-flag for SRv6 locator, but the IGP Anycast-flag can be lost due to summarization. This document defines the anycast behavior for SRv6 Service SIDs carried in BGP update messages.

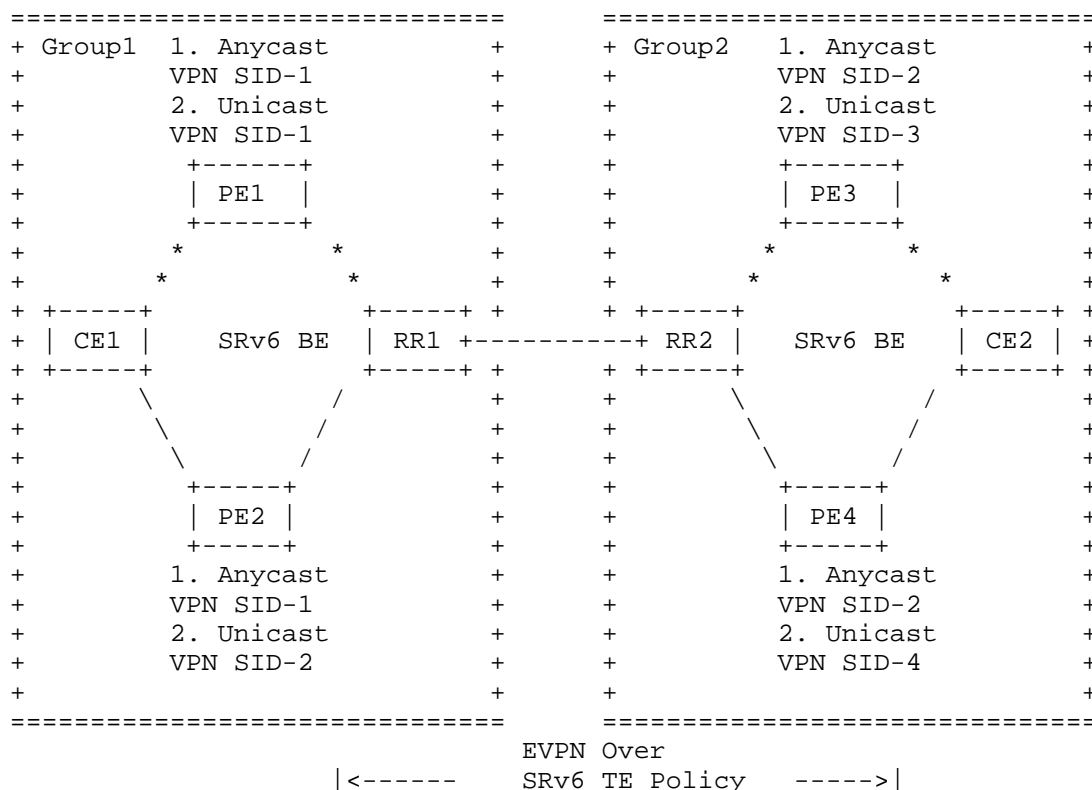


Figure 2

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. Solution

To enable the advertisement of anycast SRv6 behaviors together with the behaviors used by VPN-IPv4, VPN-IPv6, and EVPN SAFIs [RFC9252]

Internet-Draft SRv6 Anycast VPN Service April 2026
in the same BGP update message, this document defines seven Anycast
Endpoint behaviors, as follows:

The "End.DT4 with Anycast" behavior ("End.DT4.Anycast" for short) is a variant of the End.DT4 behavior.
The "End.DT6 with Anycast" behavior ("End.DT6.Anycast" for short) is a variant of the End.DT6 behavior.
The "End.DT46 with Anycast" behavior ("End.DT46.Anycast" for short) is a variant of the End.DT46 behavior.
The "End.DX4 with Anycast" behavior ("End.DX4.Anycast" for short) is a variant of the End.DX4 behavior.
The "End.DX6 with Anycast" behavior ("End.DX6.Anycast" for short) is a variant of the End.DX6 behavior.
The "End.DT2U with Anycast" behavior ("End.DT2U.Anycast" for short) is a variant of the End.DT2U behavior.
The "End.DX2 with Anycast" behavior ("End.DX2.Anycast" for short) is a variant of the End.DX2 behavior.

Continuing with the network example shown in Figure 1, the BGP VPN-IPv4/VPN-IPv6 and EVPN IP-Prefix routes advertised by PE3 and PE4 carry an SRv6 L3 Service TLV, as follows:

```
BGP Route by PE3:
  VPN Prefix of CE3:
    BGP Prefix SID Attr:
      SRv6 L3 Service TLV:
        SRv6 SID Information sub-TLV:
          SID: SID-1
          Behavior: End.DT46
        SRv6 SID Information sub-TLV:
          SID: SID-3
          Behavior: End.DT46.Anycast
```

```
BGP Route by PE4:
  VPN Prefix of CE3:
    BGP Prefix SID Attr:
      SRv6 L3 Service TLV:
        SRv6 SID Information sub-TLV:
          SID: SID-2
          Behavior: End.DT46
        SRv6 SID Information sub-TLV:
          SID: SID-3
          Behavior: End.DT46.Anycast
```

The detailed working process is as follows:

- o Egress PE3 and PE4 advertise both the anycast service SID End.DT46.Anycast and the unicast service SID End.DT46 through BGP.

- o PE1 uses the End.DT46.Anycast, and the traffic can be forwarded to PE3 and PE4 in a load-balanced manner.
- o Egress PE3 and Egress PE4 access each other via Unicast Service SID.
- o On PE3, when a failure occurs between PE3 and CE3, traffic originally destined for CE3 SHOULD be redirected. Specifically, PE3 MUST forward this traffic to PE4 using PE4's Unicast Service SID.

This solution supports multi-homing deployments for both L3VPN and L2VPN. The described processing applies equally to other Anycast-enabled behaviors including End.DT4.Anycast, End.DT6.Anycast, End.DX4.Anycast, End.DX6.Anycast, End.DT2U.Anycast, and End.DX2.Anycast.

4. Extensions for SRv6 Endpoint Behaviors

4.1. End.DT4.Anycast : End.DT4 with Anycast

The "End.DT4 with Anycast" behavior ("End.DT4.Anycast" for short) is a variant of the End.DT4 behavior.

The End.DT4.Anycast behavior extends End.DT4 with identical forwarding semantics while providing anycast functionality. This enables load balancing across multi-homed peers through advertisement of both End.DT4.Anycast and End.DT4 SIDs. Ingress nodes SHOULD use End.DT4.Anycast for anycast traffic. Egress nodes MUST use End.DT4 SID for intra-multi-homed communication and SHOULD advertise both SID types when implementing anycast services.

4.2. End.DT6.Anycast: End.DT6 with Anycast

The "End.DT6 with Anycast" behavior ("End.DT6.Anycast" for short) is a variant of the End.DT6 behavior.

The End.DT6.Anycast behavior extends End.DT6 with identical forwarding semantics while providing anycast functionality. This enables load balancing across multi-homed peers through advertisement of both End.DT6.Anycast and End.DT6 SIDs. Ingress nodes SHOULD use End.DT6.Anycast for anycast traffic. Egress nodes MUST use End.DT6 SID for intra-multi-homed communication and SHOULD advertise both SID types when implementing anycast services.

4.3. End.DT46.Anycast : End.DT46 with Anycast

The "End.DT46 with Anycast" behavior ("End.DT46.Anycast" for short) is a variant of the End.DT46 behavior.

The End.DT46.Anycast behavior extends End.DT46 with identical forwarding semantics while providing anycast functionality. This enables load balancing across multi-homed peers through advertisement of both End.DT46.Anycast and End.DT46 SIDs. Ingress nodes SHOULD use End.DT46.Anycast for anycast traffic. Egress nodes MUST use End.DT46 SID for intra-multi-homed communication and SHOULD advertise both SID types when implementing anycast services.

4.4. End.DX4.Anycast: End.DX4 with Anycast

The "End.DX4 with Anycast" behavior ("End.DX4.Anycast" for short) is a variant of the End.DX4 behavior.

The End.DX4.Anycast behavior extends End.DX4 with identical forwarding semantics while providing anycast functionality. This enables load balancing across multi-homed peers through advertisement of both End.DX4.Anycast and End.DX4 SIDs. Ingress nodes SHOULD use End.DX4.Anycast for anycast traffic. Egress nodes MUST use End.DX4 SID for intra-multi-homed communication and SHOULD advertise both SID types when implementing anycast services.

4.5. End.DX6.Anycast: End.DX6 with Anycast

The "End.DX6 with Anycast" behavior ("End.DX6.Anycast" for short) is a variant of the End.DX6 behavior.

The End.DX6.Anycast behavior extends End.DX6 with identical forwarding semantics while providing anycast functionality. This enables load balancing across multi-homed peers through advertisement of both End.DX6.Anycast and End.DX6 SIDs. Ingress nodes SHOULD use End.DX6.Anycast for anycast traffic. Egress nodes MUST use End.DX6 SID for intra-multi-homed communication and SHOULD advertise both SID types when implementing anycast services.

4.6. End.DT2U.Anycast : End.DT2U with Anycast

The "End.DT2U with Anycast" behavior ("End.DT2U.Anycast" for short) is a variant of the End.DT2U behavior.

The End.DT2U.Anycast behavior extends End.DT2U with identical forwarding semantics while providing anycast functionality. This enables load balancing across multi-homed peers through advertisement of both End.DT2U.Anycast and End.DT2U SIDs. Ingress nodes SHOULD use End.DT2U.Anycast for anycast traffic. Egress nodes

4.7. End.DX2.Anycast : End.DX2 with Anycast

The "End.DX2 with Anycast" behavior ("End.DX2.Anycast" for short) is a variant of the End.DX2 behavior.

The End.DX2.Anycast behavior extends End.DX2 with identical forwarding semantics while providing anycast functionality. This enables load balancing across multi-homed peers through advertisement of both End.DX2.Anycast and End.DX2 SIDs. Ingress nodes SHOULD use End.DX2.Anycast for anycast traffic. Egress nodes MUST use End.DX2 SID for intra-multi-homed communication and SHOULD advertise both SID types when implementing anycast services.

5. Security Considerations

TBD.

6. IANA Considerations

This document introduces seven new Endpoint behaviors. This document requests IANA assign a seven new values and update the "SRv6 Endpoint Behaviors" subregistry under the top-level "Segment Routing" registry as follows:

Value	Hex	Endpoint Behavior	Reference
TBD	TBD	End.DT4.Anycast	This document
TBD	TBD	End.DT6.Anycast	This document
TBD	TBD	End.DT46.Anycast	This document
TBD	TBD	End.DX4.Anycast	This document
TBD	TBD	End.DX6.Anycast	This document
TBD	TBD	End.DT2U.Anycast	This document
TBD	TBD	End.DX2.Anycast	This document

Table 1: SRv6 Endpoint Behaviors Subregistry

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.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [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>>.
- [RFC9136] Rabadan, J., Ed., Henderickx, W., Drake, J., Lin, W., and A. Sajassi, "IP Prefix Advertisement in Ethernet VPN (EVPN)", RFC 9136, DOI 10.17487/RFC9136, October 2021, <<https://www.rfc-editor.org/info/rfc9136>>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.

[I-D.rabnag-bess-evpn-anycast-aliasing] Rabadan, J., Ed., Nagaraj, K., Nichol, A., Sajassi A. , Lin, W., and Tantsura, J., "EVPN Anycast Multi-Homing", draft-rabnag-bess-evpn-anycast-aliasing-04, work-in-progress, July 2025.

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

Jorge Rabadan
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085
United States of America
Email: jorge.rabadan@nokia.com

