

BESS
Internet-Draft
Intended status: Standards Track
Expires: 22 October 2026

Z. Zhang
W. Lin
HPE
J. Rabadan
Nokia
A. Sajassi
Cisco Systems
20 April 2026

Multicast in L3VPNs Signaled by EVPN SAFI
draft-zzhang-bess-mcast-in-evpn-signaled-l3vpn-02

Abstract

RFC9136 specifies an EVPN SAFI Type-5 route that can be used to signal L3VPNs. This document specifies procedures for multicast in such an L3VPN.

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.

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 22 October 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 (<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. Terminology | 2 |
| 2. Introduction | 3 |
| 2.1. Optimized Inter-Subnet Multicast for EVPN | 3 |
| 2.2. Using [RFC6514] Procedures | 4 |
| 2.3. Using [RFC6037] Procedures | 5 |
| 2.4. Adapted [RFC6514] Procedures | 5 |
| 3. Specifications | 6 |
| 4. Security Considerations | 6 |
| 5. References | 6 |
| 5.1. Normative References | 6 |
| 5.2. Informative References | 7 |
| Authors' Addresses | 8 |

1. Terminology

It is expected that the audience is familiar with EVPN and MVPN concepts and terminologies. For convenience, the following terms are briefly explained.

- * PMSI: P-Multicast Service Interface - a conceptual interface for a PE to send customer multicast traffic to all or some PEs in the same VPN.
- * I-PMSI: Inclusive PMSI - to all PEs in the same VPN.
- * S-PMSI: Selective PMSI - to some of the PEs in the same VPN.
- * Leaf A-D routes: For explicit leaf tracking purposes. Triggered by S-PMSI A-D routes and targeted at triggering route's originator.
- * IMET A-D route: Inclusive Multicast Ethernet Tag A-D route. The EVPN equivalent of MVPN Intra-AS I-PMSI A-D route.
- * SMET A-D route: Selective Multicast Ethernet Tag A-D route. The EVPN equivalent of MVPN Leaf A-D route, but unsolicited and untargeted.

2. Introduction

Traditionally, an L3VPN is signaled with BGP "MPLS-labeled VPN address" SAFI and uses MPLS as the provider tunnel, as specified in [RFC4364>]. Multicast support in such an L3VPN is specified in [RFC6513] and [RFC6514].

[RFC9136] specifies another way of signaling L3VPN via EVPN SAFI Type-5 routes for two reasons:

- * VXLAN tunnels can be used, either for deployment scenarios where MPLS is not desired or for the purpose of better ECMP hashing.
- * In an environment where EVPN is already needed for L2VPN, an operator may prefer just using an additional EVPN route type to signal L3VPN routes, instead of using another SAFI for unicast reachability.

[RFC9136] does not define procedures for multicast. This document provides three options for different deployment scenarios.

2.1. Optimized Inter-Subnet Multicast for EVPN

If all multicast senders and receivers are in an EVPN domain (including both intra-DC and inter-DC cases), the Optimized Inter-Subnet Multicast (OISM) procedures defined in [RFC9625] is the best and preferred option. The advantages are that no new procedures are needed and Any Source Multicast (ASM) does not need PIM Rendezvous Point (RP) procedures.

This does require that, if not all BDs are presented on every PE, then a Supplemental Bridge Domain (SBD) needs to be configured on every PE. Since the "Interface-less IP-VRF-to-IP-VRF Model" defined in Section 4.4.1 of [RFC9136] does not use SBD, for multicast purpose it is better to not use the Interface-less model.

Additionally, in case of inter-DC, the SBD needs be stretched across DCs even if regular BDs are not stretched. If the number of PEs in all DCs becomes very large, segmentation procedures defined in [RFC9572] and further enhanced in [I-D.zzhang-bess-mvpn-evpn-cmcast-enhancements] and [I-D.rabnic-bess-evpn-mcast-eeg] can be used. Alternatively, the MVPN procedures defined in [RFC6514] can be used/adapted for an L3VPN signaled by EVPN Type-5 routes, as described in the following two sections.

2.2. Using [RFC6514] Procedures

If the OISM procedure cannot be used for any of the following situations that use L3VPN signaled by EVPN Type-5 routes:

- * There are senders/receivers not on a BD of an EVPN domain and OISM cannot extend to connect them.
- * Stretching SBD across a DCI is not desired, as described in the previous section.
- * It's a pure L3VPN scenario (only using EVPN IP Prefix routes instead of VPN-IP families), where EVPN-based multicast does not add any value or [RFC6514] procedures are desired.

MVPN procedures defined in [RFC6514] (often referred to as BGP-MVPN) can be used as is as long as:

- * The MVPN procedures treat EVPN Type-5 routes the same as routes signaled with "MPLS-labeled VPN address" when it comes to UMH selection.
- * The EVPN Type-5 routes to C-RP or C-src carry the VRF Route Import Extended Community and Source AS Extended Community.

In other words, the only difference is that the routes used for UMH selection now includes those signaled via EVPN Type-5 routes, and they MUST carry the two ECs mentioned above. The rest of [RFC6514] procedures are unchanged.

The EVPN Type-2 signaled IP routes may be used as well, though from MVPN point of view, they're no different from "local" routes associated with IRB interfaces.

In the case of non-MPLS data plane, the following options are available:

- * If PIM tunnels are used without tunnel aggregation (i.e., traffic in different VPNs sharing the same PMSI tunnel), GRE encapsulation can be used as currently specified in [RFC6514].
- * If Ingress Replication (IR) is used, or PIM tunnels are used with tunnel aggregation, the I/S-PMSI routes MUST carry a BGP Encapsulation Extended Community to indicate the encapsulation type (e.g., VXLAN or NVGRE), as specified in [RFC8365] and clarified in [RFC9012]. The label field in the PMSI Tunnel Attribute is set to the VNI. For simplicity, in the non-IR case, the VNI MUST be a global VNI.

Note that, the above procedures for VXLAN/NVGRE encapsulation MAY be used even if the L3VPN is not signaled with EVPN SAFI.

2.3. Using [RFC6037] Procedures

The historic RFC 6037 describes the legacy PIM-based MVPN (often referred to as Rosen/PIM-MVPN). While the BGP-MVPN specified in [RFC6514] is widely used and deemed more scalable and more versatile, the legacy PIM/Rosen-MVPN is still used by some operators, and in case of EVPN-signaled L3VPN, it can also be used, perhaps with little implementation change, especially if PIM-ASM-based Multicast Distribution Tree (MDT, or provider tunnel) is appropriate or desired.

It must be pointed out that, if PIM-SSM or other types of MDTs are desired, or if Inter-AS MDTs are needed, [RFC6037] requires an MDT SAFI to be used. In that case, the BGP-MVPN approach as discussed in the previous section is recommended (since a new SAFI is needed anyway, even with PIM-MVPN in this case).

2.4. Adapted [RFC6514] Procedures

Notice that an operator may have chosen to use EVPN Type-5 routes to signal L3VPN because they wanted to avoid signaling another BGP SAFI. Using [RFC6514] procedures as described in the previous section defeats that purpose because a new MCAST-VPN SAFI has to be used.

That can be resolved by adapting the [RFC6514] procedures with EVPN SAFI, as described below. This is included for sake of completeness, and may be removed in a future revision.

RFC6514 uses 7 route types, and only the Source Active route does not already have a corresponding EVPN route type:

| MVPN | | EVPN | |
|------|-------------------|------|-------------------------------|
| Type | Name | Type | Name |
| ---- | ---- | ---- | ---- |
| 1 | Intra-AS I-PMSI | 3 | IMET |
| 2 | Inter-AS I-PMSI | 9 | Per-Region I-PMSI |
| 3 | S-PMSI | 10 | S-PMSI |
| 4 | Leaf | 11 | Leaf |
| 5 | Source Active | TBD | Source Active (this document) |
| 6 | (*,G) C-Multicast | 6 | SMET |
| 7 | (S,G) C-Multicast | 7 | SMET |

As pointed out in [zzhang-bess-mvpn-evpn-cmcast-enhancements], the MVPN Type-6/7 C-multicast routes don't have leaf tracking semantics while EVPN SMET route has built-in leaf tracking semantics. Both

have pros and cons depending on the situation. This document will specify when SMET routes used for MVPN do or do not need leaf tracking semantics and the corresponding procedures.

Also as pointed out in [zzhang-bess-mvpn-evpn-cmcast-enhancements], the MVPN Type-6/7 C-multicast routes are targeted while EVPN SMET routes are not. This document specifies that the EVPN SMET routes used for MVPN purpose will be targeted, except in a special case as mentioned in [zzhang-bess-mvpn-evpn-cmcast-enhancements].

With this, the MEG (MVPN/EVPN Gateway) [RFC9625] follows the adapted MVPN procedures as specified in this document instead of the [RFC6514] procedures on MVPN side.

3. Specifications

Details to be added.

4. Security Considerations

This document does not introduce new security risks. Whatever security aspects that are applicable to [RFC7432], [RFC6513], [RFC6514] and [RFC9136] apply here.

5. References

5.1. Normative References

- [I-D.rabnic-bess-evpn-mcast-eeg]
Rabadan, J., Dornon, O., Prabhu, V., Nichol, A., Zhang, Z. J., and W. Lin, "EVPN Multicast Forwarding for EVPN to EVPN Gateways", Work in Progress, Internet-Draft, draft-rabnic-bess-evpn-mcast-eeg-06, 20 October 2025, <<https://datatracker.ietf.org/doc/html/draft-rabnic-bess-evpn-mcast-eeg-06>>.
- [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/info/rfc2119>>.
- [RFC6037] Rosen, E., Ed., Cai, Y., Ed., and IJ. Wijnands, "Cisco Systems' Solution for Multicast in BGP/MPLS IP VPNs", RFC 6037, DOI 10.17487/RFC6037, October 2010, <<https://www.rfc-editor.org/info/rfc6037>>.

- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.
- [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/info/rfc8174>>.
- [RFC8365] Sajassi, A., Ed., Drake, J., Ed., Bitar, N., Shekhar, R., Uttaro, J., and W. Henderickx, "A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)", RFC 8365, DOI 10.17487/RFC8365, March 2018, <<https://www.rfc-editor.org/info/rfc8365>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [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>>.
- [RFC9572] Zhang, Z., Lin, W., Rabadan, J., Patel, K., and A. Sajassi, "Updates to EVPN Broadcast, Unknown Unicast, or Multicast (BUM) Procedures", RFC 9572, DOI 10.17487/RFC9572, May 2024, <<https://www.rfc-editor.org/info/rfc9572>>.
- [RFC9625] Lin, W., Zhang, Z., Drake, J., Rosen, E., Ed., Rabadan, J., and A. Sajassi, "EVPN Optimized Inter-Subnet Multicast (OISM) Forwarding", RFC 9625, DOI 10.17487/RFC9625, August 2024, <<https://www.rfc-editor.org/info/rfc9625>>.

5.2. Informative References

- [I-D.zzhang-bess-mvpn-evpn-cmcast-enhancements]
Zhang, Z. J., Kebler, R., Lin, W., and E. C. Rosen, "MVPN/ EVPN C-Multicast Routes Enhancements", Work in Progress, Internet-Draft, draft-zzhang-bess-mvpn-evpn-cmcast-enhancements-04, 17 March 2024, <<https://datatracker.ietf.org/doc/html/draft-zzhang-bess-mvpn-evpn-cmcast-enhancements-04>>.

- [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>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [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>>.

Authors' Addresses

Zhaohui Zhang
HPE
Email: zhaohui.zhang@hpe.com

Wen Lin
HPE
Email: wen.lin@hpe.com

Jorge Rabadan
Nokia
Email: jorge.rabadan@nokia.com

Ali Sajassi
Cisco Systems
Email: sajassi@cisco.com