Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 20 January 2026

F. Duan
B. Song
S. Chen
Huawei Technologies

19 July 2025

Simplified MVPN for BIER and IR
draft-duan-bess-simplified-mvpn-for-bier-and-ir-04

Abstract

Per RFC6513 and RFC6514, seven MCAST-VPN NLRIs and relevant procedures are defined to build multicast forwarding tree over the service provider backbone. RFC8556 introduces that MVPN can use BIER as PMSI tunnel to perform optimal multicast forwarding. However, the complicated NLRI exchange and the switching from I-PMSI to S-PMSI tunnel is not necessary for BIER and IR tunnel. The architectural advantages of BIER and IR cannot be fully utilized. Therefore, a new simplified MVPN for BIER and IR is proposed to substitute current NLRIs exchange and procedures. This document would like to discuss the value of the MVPN simplification and provide suggestive solution.

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 20 January 2026.

Table of Contents

1. Introduction

In [RFC4364], IP Virtual Private Networks (VPNs) are proposed to forward unicast traffic from one VPN site to another. Afterwards, [RFC6037] firstly combined VPN with IP Multicast and multicast forwarding tree can be built over the provider backbone. PIM was the only protocol to establish PMSI tunnels. [RFC6513] and [RFC6514] then improved the MVPN procedure. On the one hand, more flexible tunnel type such as P2MP and IR are specified. On the other hand, seven MCAST-VPN NLRIs are defined to advertise the information of MVPN members, tunnels, source location and join/prune messages. MVPN solutions usually started with instantiate inclusive PMSI to build the multicast distribution trees over the provider network.

In order to optimize the bandwidth utilization of the provider backbone network, S-PMSI A-D Route is designed so that selective multicast can be performed when the traffic of (C-S,C-G) exceeds the

preset threshold. Switching from I-PMSI to S-PMSI is an inevitable action for selective multicast when the tunnel type is mLDP or RSVP-TE. Because new underlay tunnel establishing procedures are necessary for these two tunnels. The switching results in the complicated NLRI exchanging procedures.

[RFC8556] introduces that MVPN can use BIER to conduct optimal multicast forwarding. The complicated NLRI exchanging procedures are still maintained while those are unnecessary for BIER and Ingress Replication Tunnel. There are several problems in current MVPN procedures:

a. Even though per-flow multicast state is not maintained in the P routers, ingress root PE still follows the traditional process of building multicast tunnel. Root PE also needs to check whether the amount of multicast flow exceeds the preset threshold at any time so that it can initiate the switching from I-PMSI to S-PMSI. The exchange of control-plane and data-plane are still very complicated.

b. There are two types of NLRIs involved in the process of customer’s routes advertisement. Besides, four types of NLRIs are leveraged to collect tunnel informations. The exchange of NLRIs between each router is complicated.

The architectural advantages of BIER and IR are that they can intrinsically support explicit tracking at the ingress PE. When LDP and RSVP-TE tunnels are deployed, new MPLS labels or Opaque value are assigned along each branches of the multicast tunnels when the S-PMSI tunnels are initialized, which means new forwarding table are constructed along each relevant routers. When underlay tunnel is BIER or IR, S-PMSI tunnel can directly use the same forwarding table of I-PMSI tunnel on each router. The only way to differentiate these two tunnels is explicit tracking. Inress PE use explicit tracking to specify different leaves in the multicast packet. Each leaf PE of BIER and IR is globally unique from the perspective of ingress PE. Therefore, S-PMSI tunnel can be constructed directly at first and switching from I-PMSI to S-PMSI tunnel will no longer needed.

On the other hand, segment routing is widely discussed and implemented nowadays and it is regarded as a simplification of MPLS. SR-MPLS, SR-BIER and SR-IR are simplification of existing tunnel types in a sense. With SR, current MVPN architecture and NLRI exchanges seem to be too heavy. Under these circumstances, a light-weight architecture of MVPN needs to be considered. In that way, the feature of explicit tracking can also be fully utilized.

One possible method is proposed in this document to simplify the MVPN procedure for BIER and IR. There would be no inclusive PMSI tunnel. Two new multicast routes and procedures are proposed to substitute the existing seven NLRIs.

2. Terminology

The terminology used in this document is the terminology defined in[RFC6513], [RFC6514] and [RFC8556].

For convenience of description, the abbreviations used in this document is listed below.

NLRI: Network Layer Reachability Information

UMH: Upstream Multicast Hop

PMSI: P-Multicast Service Interface

VPN: Virtual Private Network

MVPN: Multicast VPN

RD: Route Distinguisher

IR: Ingress Replication

3. Specification

3.1. Simplification of Type 1 and 3 NLRI

Type 1 and 3 NLRIs may be replaced by the eligible UMH route. The eligible UMH route was initially introduced in [RFC6513]. It contains Source AS Extended Community and VRF Route Import Extended Community. In this document, MS-ID and BIER attributes are added into the eligible UMH route so that type 1 and 3 NLRIs are no longer needed. When the leaf PE receives the eligible UMH routes, it will import the unicast route into its local instance. Simultaneously, the MS-ID will be used to generate the correspondence between the MS-id and local instance. When the leaf PE receives the join or prune messages, it will find the multicast source or RP in the unicast routing-table of corresponding instance. The underlay BIER attribute of the unicast route will be used.

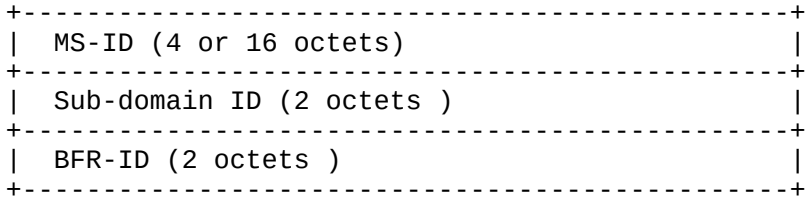

Figure 1: New MVPN Eligible UMH Route

3.2. Simplification of Type 4, 6 and 7 NLRIs

When leaf PE receives igmp membership report or pim join messages, it will check whether the sub-domain-id inside the BIER attribute of the unicast route is same as its local sub-domain-id. If the two IDs are same, leaf PE will advertise a BGP multicast route to root PE. The BGP multicast route is proposed in this document to replace Type 4, 6 and 7 NLRI. It contains RD, originator IP, source address and group address. Additionally, it includes one-octet field called 'Flag'. Flag is used to distinguish (C-*,C-G) Join, (C-S,C-G) Join and (C-S,C-G,rpt) Prune. The route also includes BIER sub-domain-id and BFR-id of leaf PE. The conventional Join and Prune of c-multicast route are substituted by the update and withdraw of this BGP multicast route. Moreover, Source AS Extended Community and VRF Route Import Extended Community are also carried by the BGP multicast route.

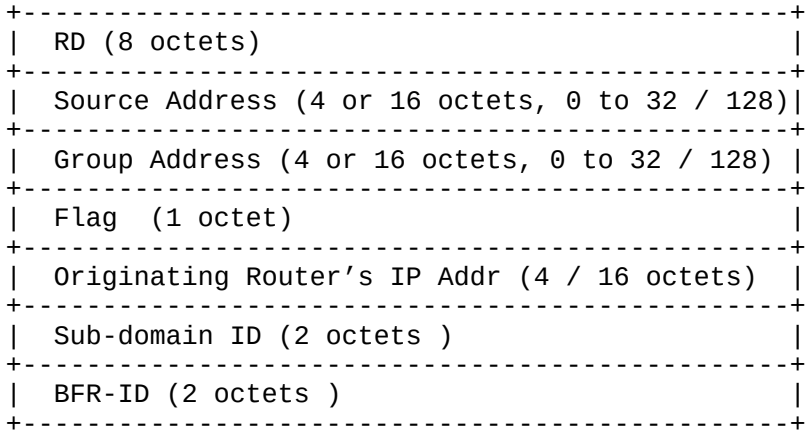

Figure 2: New BGP Multicast Route

4. Segmentation scenario

Adaption about Inter-AS I-PMSI A-D route has not been mentioned yet. We are working on solution for tunnel segmentation scenario and relevant solutions will be updated in later version.

5. Back compatibility

Back compatibility is a significant issue and will be discussed in the future.

6. Security Considerations

//TODO

7. IANA Considerations

//TODO

8. Acknowledgements

//TODO

9. References

9.1. 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/info/rfc2119>.

[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>.

[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>.

[RFC8556] Rosen, E., Ed., Sivakumar, M., Przygienda, T., Aldrin, S., and A. Dolganow, "Multicast VPN Using Bit Index Explicit Replication (BIER)", RFC 8556, DOI 10.17487/RFC8556, April 2019, <https://www.rfc-editor.org/info/rfc8556>.

9.2. Informative References

[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>.

Authors' Addresses

Fanghong Duan
Huawei Technologies
Email: duanfanghong@huawei.com

Bo Song
Huawei Technologies
Email: sunbird.song@huawei.com

Siyu Chen
Huawei Technologies
Email: chensiyu27@huawei.com