

BIER Penultimate Hop Popping
draft-ietf-bier-php-06

Abstract

Bit Index Explicit Replication (BIER) can be used as provider tunnel for Multicast Virtual Private Network (MVPN). Global Table Multicast [[RFC7716](#)] or Ethernet Virtual Private Network (EVPN). It is possible that not all routers in the provider network support BIER and there are various methods to handle BIER-incapable transit routers. However those methods assume the MVPN/EVPN Provider Edges (PEs) are BIER-capable. This document specifies a method to allow BIER-incapable routers to act as MVPN/EVPN PEs with BIER as the transport, by having the upstream BIER Forwarding Router (BFR) that is connected directly or indirectly via a tunnel to a BIER-incapable PE remove the BIER header and send the payload to the PE.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119](#).

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 February 25, 2021.

Copyright Notice

Copyright (c) 2020 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 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
2.	Specifications	3
3.	Security Considerations	4
4.	IANA Considerations	4
5.	Acknowledgements	5
6.	References	5
6.1.	Normative References	5
6.2.	Informative References	6
	Author's Address	7

[1.](#) Introduction

The BIER architecture includes three layers: the "routing underlay", the "BIER layer", and the "multicast flow overlay". The multicast flow overlay is responsible for the BIER Forwarding Egress Routers (BFERs) to signal to BIER Forwarding Ingress Routers (BFIRs) that they are interested in receiving certain multicast flows so that BFIRs can encode the correct bitstring for BIER forwarding by the BIER layer.

MVPN and EVPN are two similar overlays where BGP Auto-Discovery routes for MVPN/EVPN are exchanged among all PEs to signal which PEs need to receive multicast traffic for all or certain flows. Typically the same provider tunnel type is used for traffic to reach all receiving PEs.

Consider an MVPN/EVPN deployment where enough provider routers are BIER-capable for BIER to become the preferred choice of provider tunnel. However, some PEs cannot be upgraded to support BIER forwarding. While there are ways to allow an ingress PE to send traffic to some PEs with one type of tunnel and send traffic to some

other PEs with a different type of tunnel, the procedure becomes complicated and forwarding is not optimized.

One way to solve this problem is to use Penultimate Hop Popping (PHP) so that the upstream BFR can pop the BIER header and send the payload "natively" (note that the upstream BFR can be connected directly or indirectly via any type of tunnel to the PE). This is similar to Multi-Protocol Label Switching (MPLS) PHP though it is the BIER header that is popped.

The transition of an existing MVPN/EVPN deployment with traditional provider tunnels to using BIER with some PEs not capable of receiving BIER packets can be incremental. All PEs are first upgraded to support BIER at least in the control plane, with those not capable of BIER forwarding requesting PHP. Then BIER-capable ingress PEs independently and incrementally switch to BIER transport.

While the above text uses MVPN/EVPN as example, BIER PHP is applicable to any scenario where the multicast flow overlay edge router does not support BIER, as long as the edge router does not need to know the transmitting BFIR or participate in BIER OAM procedures.

This works well if a BIER-incapable PE only needs to receive multicast traffic. If it needs to send multicast traffic as well, then it must Ingress Replicate to a BIER-capable helper PE, who will in turn relay the packet to other PEs. The helper PE is either a Virtual Hub as specified in [\[RFC7024\]](#) for MVPN and [\[I-D.keyupate-bess-evpn-virtual-hub\]](#) for EVPN, or an AR-Replicator as specified in [\[I-D.ietf-bess-evpn-optimized-ir\]](#) for EVPN.

2. Specifications

The BIER Penultimate Hop Popping is intended only for the scenario where a multicast flow overlay router for a BIER domain does not support BIER forwarding, either entirely or just for some particular BitStringLengths. In the latter case, PHP is only for BIER packets with those BitStringLengths. The flow overlay router would be a BFER if it did support BIER forwarding, and PHP would not be done by its penultimate hop.

The procedures in this section apply only if, by means outside the scope of this document, it is known that the payload after BIER header is one of the following:

- o MPLS packets with downstream-assigned label at top of stack (i.e., the Proto field in the BIER header is 1). For example, a label

from a Domain-wide Common Block (DCB) is used as specified in [I-D.ietf-bess-mvpn-evpn-aggregation-label].

- o IPv4/IPv6 multicast packets for which Reverse Path Forwarding check is disabled.

A BIER-incapable router, if acting as a multicast flow overlay router, MUST signal its BIER information as specified in [[RFC8401](#)] or [[RFC8444](#)] or [[I-D.ietf-bier-idr-extensions](#)], with a PHP sub-sub-TLV included in the BIER sub-TLV attached to the BIER-incapable router's BIER prefix to request BIER PHP from other BFRs. The sub-sub-TLV's type is TBD, and the length is 0.

With MPLS encapsulation, the BIER-incapable multicast flow overlay router MAY omit the BIER MPLS Encapsulation sub-sub-TLV, or MUST set the Label field in BIER MPLS Encapsulation sub-sub-TLV to Implicit Null Label [[RFC3032](#)].

With MPLS encapsulation, if a BFER does not support certain BSL, it MAY still advertise a corresponding BIER MPLS Encapsulation sub-TLV but set the Label field to Implicit Null Label.

If a BFR follows [section 6.9 of \[RFC8279\]](#) to handle BIER-incapable routers, it must treat a router as BIER-incapable if the label advertised by the router is Implicit Null, or if the router advertises a PHP sub-sub-TLV, so that the router is not used as a transit BFR.

If the downstream neighbor for a BIER prefix is the one advertising the prefix with a PHP sub-sub-TLV or with an Implicit Null Label in the Label field in its BIER MPLS Encapsulation sub-sub-TLV, then when the corresponding BIRT or BIFT entry is created/updated, the forwarding behavior MUST be that the BIER header is removed and the payload be sent to the downstream router without the BIER header, either directly or over any type of tunnel.

3. Security Considerations

This specification does not introduce additional security concerns beyond those already discussed in BIER architecture and OSPF/ISIS/BGP extensions for BIER signaling.

4. IANA Considerations

This document requests a new sub-sub-TLV type value from the "Sub-sub-TLVs for BIER Info Sub-TLV" registry in the "IS-IS TLV Codepoints" registry:

Type	Name
----	----
TBD	BIER PHP Request

This document also requests a new sub-TLV type value from the OSPFv2 Extended Prefix TLV Sub-TLV registry:

Type	Name
----	----
TBD	BIER PHP Request

5. Acknowledgements

The author wants to thank Eric Rosen and Antonie Przygienda for their review, comments and suggestions. The author also wants to thank Senthil Dhanaraj for his suggestion of requesting PHP if a BFER does not support certain BSL.

6. References

6.1. Normative References

- [I-D.ietf-bess-mvpn-evpn-aggregation-label]
Zhang, Z., Rosen, E., Lin, W., Li, Z., and I. Wijnands,
"MVPN/EVPN Tunnel Aggregation with Common Labels", [draft-ietf-bess-mvpn-evpn-aggregation-label-03](#) (work in progress), October 2019.
- [I-D.ietf-bier-evpn]
Zhang, Z., Przygienda, T., Sajassi, A., and J. Rabadan,
"EVPN BUM Using BIER", [draft-ietf-bier-evpn-03](#) (work in progress), April 2020.
- [I-D.ietf-bier-idr-extensions]
Xu, X., Chen, M., Patel, K., Wijnands, I., and T. Przygienda,
"BGP Extensions for BIER", [draft-ietf-bier-idr-extensions-07](#) (work in progress), September 2019.
- [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>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", [RFC 8279](#), DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.

- [RFC8296] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks", [RFC 8296](#), DOI 10.17487/RFC8296, January 2018, <<https://www.rfc-editor.org/info/rfc8296>>.
- [RFC8401] Ginsberg, L., Ed., Przygienda, T., Aldrin, S., and Z. Zhang, "Bit Index Explicit Replication (BIER) Support via IS-IS", [RFC 8401](#), DOI 10.17487/RFC8401, June 2018, <<https://www.rfc-editor.org/info/rfc8401>>.
- [RFC8444] Psenak, P., Ed., Kumar, N., Wijnands, IJ., Dolganow, A., Przygienda, T., Zhang, J., and S. Aldrin, "OSPFv2 Extensions for Bit Index Explicit Replication (BIER)", [RFC 8444](#), DOI 10.17487/RFC8444, November 2018, <<https://www.rfc-editor.org/info/rfc8444>>.

6.2. Informative References

- [I-D.ietf-bess-evpn-optimized-ir]
Rabadan, J., Sathappan, S., Lin, W., Katiyar, M., and A. Sajassi, "Optimized Ingress Replication solution for EVPN", [draft-ietf-bess-evpn-optimized-ir-07](#) (work in progress), July 2020.
- [I-D.keyupate-bess-evpn-virtual-hub]
Patel, K., Sajassi, A., Drake, J., Zhang, Z., and W. Henderickx, "Virtual Hub-and-Spoke in BGP EVPNs", [draft-keyupate-bess-evpn-virtual-hub-02](#) (work in progress), September 2019.
- [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>>.
- [RFC7024] Jeng, H., Uttaro, J., Jalil, L., Decraene, B., Rekhter, Y., and R. Aggarwal, "Virtual Hub-and-Spoke in BGP/MPLS VPNs", [RFC 7024](#), DOI 10.17487/RFC7024, October 2013, <<https://www.rfc-editor.org/info/rfc7024>>.

[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>>.

Author's Address

Zhaohui Zhang
Juniper Networks

EMail: zzhang@juniper.net