

BIER
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2020

Z. Zhang
ZTE Corporation
A. Przygienda
Juniper Networks, Inc.
I. Wijnands
Cisco Systems
H. Bidgoli
Nokia
M. McBride
Futurewei
July 8, 2019

BIER in IPv6 (BIERin6)
draft-zhang-bier-bierin6-03

Abstract

BIER is a new architecture for the forwarding of multicast data packets. This document defines native IPv6 encapsulation for BIER hop-by-hop forwarding or BIERin6 for short.

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 January 9, 2020.

Copyright Notice

Copyright (c) 2019 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. IPv6 Header | 3 |
| 2.1. IPv6 Options Considerations | 3 |
| 3. BIER Header | 4 |
| 4. IPv6 Encapsulation Advertisement | 4 |
| 4.1. Format | 4 |
| 4.2. Inter-area prefix redistribution | 5 |
| 5. IANA Considerations | 5 |
| 6. Security Considerations | 5 |
| 7. Acknowledgement | 5 |
| 8. References | 5 |
| 8.1. Normative References | 6 |
| 8.2. Informative References | 6 |
| Authors' Addresses | 7 |

1. Introduction

BIER [RFC8279] is a new architecture for the forwarding of multicast data packets. It provides optimal forwarding through a "multicast domain" and it does not necessarily precondition construction of a multicast distribution tree, nor does it require intermediate nodes to maintain any per-flow state.

This document specifies non-MPLS BIER forwarding in an IPv6 [RFC8200] environment, referred to as BIERin6, using non-MPLS BIER encapsulation specified in [RFC8296].

MPLS BIER forwarding in IPv6 is outside the scope of this document.

This document uses terminology defined in [RFC8279] and [RFC8296].

[RFC8296] defines the BIER encapsulation format in MPLS and non-MPLS environment. In case of non-MPLS environment, a BIER packet is the payload of an "outer" encapsulation, which has a "next protocol" codepoint that is set to a value that means "non-MPLS BIER".

That can be used as is in a pure IPv6 non-mpls environment. Between two directly connected BFRs, a BIER header could directly follow link layer header, e.g., an Ethernet header (with the Ethertype set to 0xAB37). If a BFR needs to tunnel BIER packets to another BFR, e.g. per [RFC8279] Section 6.9, IPv6 encapsulation can be used, with the destination address being the downstream BFR and the Next Header field set to a to-be-assigned value for "non-MPLS BIER".

The IPv6 encapsulation could be used even between two directly connected BFRs in the following two cases:

- o An operator mandates all traffic to be carried in IPv6.
- o A BFR does not have BIER support in its "fast forwarding path" and relies on "slow/software forwarding path", e.g. in environments like [RFC7368] where high throughput multicast forwarding performance is not critical.

2. IPv6 Header

Whenever IPv6 encapsulation is used for BIER forwarding, The Next Header field in the IPv6 Header (if there are no extension headers), or the Next Header field in the last extension header is set to TBD, indicating that the payload is a BIER packet.

If the neighbor is directly connected, The destination address in IPv6 header SHOULD be the neighbor's link-local address on this router's outgoing interface, the source destination address SHOULD be this router's link-local address on the outgoing interface, and the IPv6 TTL MUST be set to 1. Otherwise, the destination address SHOULD be the BIER prefix of the BFR neighbor, the source address SHOULD be this router's BIER prefix, and the TTL MUST be large enough to get the packet to the BFR neighbor.

The Flow-ID in the IPv6 packet SHOULD be copied from the entropy field in the BIER encapsulation.

2.1. IPv6 Options Considerations

RFC 8200 section 4, defines the IPv6 extension headers. Currently there are two defined extension headers, Hop-by-Hop and Destination options header, which can carry a variable number of options. These extension headers are inserted by the source node.

For directly connected BIER routers, IPv6 Hop-by-Hop or Destination options are irrelevant and SHOULD NOT be inserted by BFIR on the BIERin6 packet. In this case IPv6 header, Next Header field should be set to TBD. Any IPv6 packet arriving on BFRs and BFRs, with multiple extension header where the last extension header has a Next Header field set to TBD, SHOULD be discard and the node should transmit an ICMP Parameter Problem message to the source of the packet (BFIR) with an ICMP code value of TBD10 ('invalid options for BIERin6').

This also indicates that for disjoint BIER routers using IPv6 encapsulation, there SHOULD NOT be any IPv6 Hop-by-Hop or Destination options be present in a BIERin6 packet. In this case, if additional traffic engineering is required, IPv6 tunneling (i.e. BIERin6 over SRv6) can be implemented.

3. BIER Header

The BIER header MUST be encoded per Section 2.2 of [RFC8296].

The BIFT-id is either encoded per [I-D.ietf-bier-non-mpls-bift-encoding] or per advertised by BFRs, as specified in [I-D.dhanaraj-bier-lsr-ethernet-extensions].

4. IPv6 Encapsulation Advertisement


When IPv6 encapsulation is not required between directly connected BFRs, no signaling in addition to that specified in [I-D.dhanaraj-bier-lsr-ethernet-extensions] is needed.

Otherwise, a node that requires IPv6 encapsulation MUST advertise the BIER IPv6 transportation sub-TLV/sub-sub-TLV according to local configuration or policy in the BIER domain to request other BFRs to always use IPv6 encapsulation.

In presence of multiple encapsulation possibilities hop-by-hop it is a matter of local policy which encapsulation is imposed and the receiving router MUST accept all encapsulations that it advertised.

4.1. Format

The BIER IPv6 transportation is a new sub-TLV of BIER defined in OSPF [RFC8444], and a new sub-sub-TLV of BIER Info sub-TLV defined in ISIS [RFC8401].



- o Type: For OSPF, value TBD1 (prefer 12) is used to indicate it is the IPv6 transportation sub-TLV. For ISIS, value TBD2 (prefer 3) is used to indicate it is the IPv6 transportation sub-sub-TLV.
- o Length: 0.

4.2. Inter-area prefix redistribution

When BFR-prefixes are advertised across IGP areas per [I-D.dhanaraj-bier-lsr-ethernet-extensions] or redistributed across protocol boundaries per [I-D.zwzw-bier-prefix-redistribute], the BIER IPv6 transportation sub-TLV or sub-sub-TLV MAY be re-advertised/re-distributed as well.

5. IANA Considerations

IANA is requested to assign a new "BIER" type for "Next Header" in the "Assigned Internet Protocol Numbers" registry.

IANA is requested to assign a new "BIERin6" type for "invalid options" in the "ICMP code value" registry.

IANA is requested to assign a new "BIER IPv6 transportation Sub-TLV" type in the "OSPFv2 Extended Prefix TLV Sub-TLVs" Registry.

IANA is requested to set up a new "BIER IPv6 transportation Sub-sub-TLV" type in the "IS-IS BIER Info sub-TLV" Registry.

6. Security Considerations

General IPv6 and BIER security considerations apply.

7. Acknowledgement

The authors would like to thank Jeffrey Zhang for his review and valuable contributions.

8. References

8.1. Normative References

- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [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>>.

8.2. Informative References

- [I-D.dhanaraj-bier-lsr-ethernet-extensions]
Dhanaraj, S., Wijnands, I., Psenak, P., Zhang, Z., Yan, G., and J. Xie, "LSR Extensions for BIER over Ethernet", draft-dhanaraj-bier-lsr-ethernet-extensions-00 (work in progress), January 2019.
- [I-D.ietf-bier-bar-ipa]
Zhang, Z., Przygienda, T., Dolganow, A., Bidgoli, H., Wijnands, I., and A. Gulko, "BIER Underlay Path Calculation Algorithm and Constraints", draft-ietf-bier-bar-ipa-04 (work in progress), May 2019.
- [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-06 (work in progress), January 2019.

- [I-D.ietf-bier-non-mpls-bift-encoding]
Wijnands, I., Xu, X., and H. Bidgoli, "An Optional Encoding of the BIFT-id Field in the non-MPLS BIER Encapsulation", draft-ietf-bier-non-mpls-bift-encoding-01 (work in progress), October 2018.
- [I-D.zhang-bier-babel-extensions]
Zhang, Z. and T. Przygienda, "BIER in BABEL", draft-zhang-bier-babel-extensions-01 (work in progress), June 2017.
- [I-D.zwzw-bier-prefix-redistribute]
Zhang, Z., Bo, W., Zhang, Z., and I. Wijnands, "BIER Prefix Redistribute", draft-zwzw-bier-prefix-redistribute-02 (work in progress), March 2019.
- [RFC7368] Chown, T., Ed., Arkko, J., Brandt, A., Troan, O., and J. Weil, "IPv6 Home Networking Architecture Principles", RFC 7368, DOI 10.17487/RFC7368, October 2014, <<https://www.rfc-editor.org/info/rfc7368>>.

Authors' Addresses

Zheng (Sandy) Zhang
ZTE Corporation

EEmail: zhang_ietf@hotmail.com

Tony Przygienda
Juniper Networks, Inc.

EEmail: prz@juniper.net

IJsbrand Wijnands
Cisco Systems

EEmail: ice@cisco.com

Hooman Bidgoli
Nokia

EEmail: hooman.bidgoli@nokia.com

Mike McBride
Futurewei

EMail: mmcbride@futurewei.com