

BIER
Internet-Draft
Intended status: Standards Track
Expires: 3 September 2025

Z. Zhang
ZTE Corporation
Z. Zhang, Ed.
Juniper Networks
I. Wijnands
Individual
M. Mishra
Cisco Systems
H. Bidgoli
Nokia
G. Mishra
Verizon
2 March 2025

Supporting BIER in IPv6 Networks (BIERin6)
draft-ietf-bier-bierin6-11

Abstract

BIER is a multicast forwarding architecture that does not require per-flow state inside the network yet still provides optimal replication. This document describes how the existing BIER encapsulation specified in RFC 8296 works in a non-MPLS IPv6 network, which is referred to as BIERin6. Specifically, like in an IPv4 network, BIER can work over L2 links directly or over tunnels. In case of IPv6 tunneling, a new IP "Next Header" type is to be assigned for BIER.

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 3 September 2025.

Copyright Notice

Copyright (c) 2025 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. Introduction	2
1.1. BIER over L2/Tunnels	3
1.2. Considerations of Requirements for BIER in IPv6 Networks	4
2. IPv6 Header	5
3. BIER Header	6
4. IPv6 Encapsulation Advertisement	6
4.1. Format	6
4.2. Inter-area prefix redistribution	6
5. IANA Considerations	7
6. Security Considerations	7
7. Acknowledgement	7
8. References	7
8.1. Normative References	7
8.2. Informative References	8
Authors' Addresses	10

1. Introduction

BIER [RFC8279] is a multicast forwarding architecture that does not require per-flow state inside the network yet still provides optimal replication.


BIER forwarding with MPLS is IPv4/IPv6 agnostic. This document describes how BIER works in a non-MPLS IPv6 [RFC8200] environment using non-MPLS BIER encapsulation [RFC8296], with optional procedures specified for IPv6 specific features.

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

1.1. BIER over L2/Tunnels

[RFC8296] defines the BIER encapsulation format for MPLS and non-MPLS data planes. With a non-MPLS data plane, a BIER packet is the payload of an "outer" encapsulation, which could be a L2 link or a tunnel. The outer encapsulation has a "next header" field that is set to a value for "non-MPLS BIER". This "BIER over L2/Tunnel" model can be used as is in an IPv6 non-mpls environment, and is referred to as BIERin6.

If a BFR needs to tunnel BIER packets to another BFR, e.g. per [RFC8279] Section 6.9, while any type of tunnel will work, for best efficiency native 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 BIER.



Between two directly connected BFRs, a BIER header can directly follow link layer header, e.g., an Ethernet header (with the Ethertype set to 0xAB37). Optionally, IPv6 encapsulation can be used even between directly connected BFRs (i.e. one-hop IPv6 tunneling) in the following two cases:

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

1.2. Considerations of Requirements for BIER in IPv6 Networks

[I-D.ietf-bier-ipv6-requirements] lists mandatory and optional requirements for BIER in IPv6 Networks. As a solution based on the BIER over L2/tunnel model [RFC8296], BIERin6 satisfies all the mandatory requirements.

For the two optional requirements for fragmentation and Encapsulating Security Payload (ESP), they can be satisfied by one of two ways:

- * IPv6 based fragmentation/ESP: a BFIR encapsulates the payload in IPv6 with fragmentation and/or ESP header, and then the IPv6 packets are treated as BIER payload.
- * Generic Fragmentation/ESP
[I-D.zhang-intarea-generic-delivery-functions]: a BFIR does generic fragmentation and/or ESP (without using IPv6 encapsulation) and the resulting packets are treated as BIER payload.

BIERin6 can support SRv6 based overlay services (e.g. MVPN/EVPN) with one of the following methods:

- * An ingress PE (which is a BFIR) can encapsulate customer packets with an IPv6 header (with optional fragmentation and ESP extension headers). Any SRv6-based scheme can be used for service delimiting.
- * Alternatively, since only the IPv6 address in the above-mentioned outer IPv6 header is used for service delimiting purpose, a new value can be assigned for the Proto field in the BIER header to indicate that an SRv6 Service SID [RFC9252] (instead of an entire IPv6 header) is added between the BIER header and original payload. For example, an End.DT2/4/6 service SID could be used to route the original payload in a corresponding VRF.

The details of fragmentation/ESP support and service delimitation are all outside the scope of this document.

BIERin6 being a solution based on [RFC8279] and [RFC8296], ECMP is inherently supported by BFRs using the 20-bit entropy field in the BIER header for the load balancing hash. When a BIER packet is transported over an IPv6 tunnel, the entropy value is copied into the 20-bit IPv6 Flow Label [RFC6437], so that routers along the tunnel can do ECMP based on Flow Labels (instead of hashing based on 5-tuple of an IP packet). For a router along the tunnel doing deep packet inspection for ECMP purpose, if it understands BIER header it can go past the BIER header to look for the 5-tuple input key to a hash

function. Otherwise, it stops at the BIER header. In either case the router will not mistake the BIER header as an IP header so no misordering should happen.

BIER has its own OAM functions independent of those related to the underlying links or tunnels. With BIERin6 following the "BIER over L2/tunnel" model, IPv6 OAM function and BIER OAM functions are used independently for their own purposes.

Specifically, BIERin6 works with all of the following OAM methods, or any future methods that are based on the "BIER over L2/tunnel" model:

- * BIER OAM specified in [I-D.ietf-bier-ping]
- * BIER BFD specified in [I-D.ietf-bier-bfd]
- * BIER Performance Measurement specified in [I-D.ietf-bier-pmmmm-oam]
- * BIER Path Maximum Transmission Unit Discovery specified in [I-D.ietf-bier-path-mtu-discovery]
- * BIER IOAM specified in [I-D.xzlnp-bier-ioam]

2. IPv6 Header

If IPv6 encapsulation is used to tunnel BIER packets (whether to a direct or indirect BIER neighbor), 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 TTL MUST be set to 1.

If the neighbor is not directly connected, 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 label" field in the IPv6 packet SHOULD be copied from the entropy field in the BIER encapsulation.

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.ietf-bier-lsr-non-mpls-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.ietf-bier-lsr-non-mpls-extensions] is needed.

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

4.1. Format

The BIER IPv6 Encapsulation is a new sub-sub-TLV of OSPFv3 BIER non-MPLS Encapsulation sub-TLV, and a new sub-sub-sub-TLV of ISIS BIER non-MPLS Encapsulation sub-sub-TLV as per [I-D.ietf-bier-lsr-non-mpls-extensions].

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type           |   Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

* Type: For OSPF, value TBD1 indicates it is the IPv6 Encapsulation sub-TLV. For ISIS, value TBD2 indicates it is the IPv6 encapsulation sub-sub-TLV.

* Length: 0.

4.2. Inter-area prefix redistribution

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

5. IANA Considerations

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

IANA is requested to assign a "BIER IPv6 encapsulation Sub-sub-TLV" type in the "OSPFv3 BIER non-MPLS Encapsulation sub-TLV" Registry.

IANA is requested to assign a "BIER IPv6 encapsulation Sub-sub-sub-TLV" type in the "IS-IS BIER non-MPLS Encapsulation sub-sub-TLV" Registry.

IANA is requested to allocate a value "SRv6 Service" from "BIER Next Protocol Identifiers" registry to indicate that BIER payload starts with an SRv6 Service SID.

6. Security Considerations

General IPv6 and BIER security considerations apply.

7. Acknowledgement

The authors would like to thank Tony Przygienda, Nagendra Kumar for their review and valuable comments.

8. References

8.1. Normative References

- [I-D.ietf-bier-lsr-non-mpls-extensions]
Dhanaraj, S., Yan, G., Wijnands, I., Psenak, P., Zhang, Z. J., and J. Xie, "LSR Extensions for BIER non-MPLS Encapsulation", Work in Progress, Internet-Draft, draft-ietf-bier-lsr-non-mpls-extensions-03, 7 February 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-lsr-non-mpls-extensions-03>>.
- [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>>.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <<https://www.rfc-editor.org/info/rfc6437>>.

- [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>>.
- [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>>.
- [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>>.

8.2. Informative References

- [I-D.ietf-bier-bfd]
Xiong, Q., Mirsky, G., hu, F., Liu, C., and G. S. Mishra, "BIER BFD", Work in Progress, Internet-Draft, draft-ietf-bier-bfd-08, 26 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-bfd-08>>.
- [I-D.ietf-bier-ipv6-requirements]
McBride, M., Xie, J., Geng, X., Dhanaraj, S., Asati, R., Zhu, Y., Mishra, G. S., and Z. J. Zhang, "BIER IPv6 Requirements", Work in Progress, Internet-Draft, draft-ietf-bier-ipv6-requirements-09, 28 September 2020, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-ipv6-requirements-09>>.
- [I-D.ietf-bier-non-mpls-bift-encoding]
Wijnands, I., Mishra, M. P., Xu, X., and H. Bidgoli, "An Optional Encoding of the BIFT-id Field in the non-MPLS BIER Encapsulation", Work in Progress, Internet-Draft,

draft-ietf-bier-non-mpls-bift-encoding-04, 30 May 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-non-mpls-bift-encoding-04>>.

[I-D.ietf-bier-path-mtu-discovery]

Mirsky, G., Przygienda, T., and A. Dolganow, "Path Maximum Transmission Unit Discovery (PMTUD) for Bit Index Explicit Replication (BIER) Layer", Work in Progress, Internet-Draft, draft-ietf-bier-path-mtu-discovery-17, 4 July 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-path-mtu-discovery-17>>.

[I-D.ietf-bier-ping]

Nainar, N. K., Pignataro, C., Chen, M., and G. Mirsky, "BIER Ping and Trace", Work in Progress, Internet-Draft, draft-ietf-bier-ping-15, 8 November 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-ping-15>>.

[I-D.ietf-bier-pmmm-oam]

Mirsky, G., Zheng, L., Chen, M., and G. Fioccola, "Performance Measurement (PM) with Marking Method in Bit Index Explicit Replication (BIER) Layer", Work in Progress, Internet-Draft, draft-ietf-bier-pmmm-oam-16, 7 November 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-pmmm-oam-16>>.

[I-D.ietf-bier-prefix-redistribute]

Zhang, Z., Wu, B., Zhang, Z. J., Wijnands, I., Liu, Y., and H. Bidgoli, "BIER Prefix Redistribute", Work in Progress, Internet-Draft, draft-ietf-bier-prefix-redistribute-08, 23 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-prefix-redistribute-08>>.

[I-D.xzlnp-bier-ioam]

Min, X., Zhang, Z., Liu, Y., Nainar, N. K., and C. Pignataro, "BIER Encapsulation for IOAM Data", Work in Progress, Internet-Draft, draft-xzlnp-bier-ioam-08, 29 July 2024, <<https://datatracker.ietf.org/doc/html/draft-xzlnp-bier-ioam-08>>.

[I-D.zzhang-intarea-generic-delivery-functions]

Zhang, Z. J., Bonica, R., Kompella, K., and G. Mirsky, "Generic Delivery Functions", Work in Progress, Internet-Draft, draft-zzhang-intarea-generic-delivery-functions-03, 11 July 2022, <<https://datatracker.ietf.org/doc/html/draft-zzhang-intarea-generic-delivery-functions-03>>.

[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
Email: zhang.zheng@zte.com.cn

Zhaohui Zhang (editor)
Juniper Networks
Email: zzhang@juniper.net

IJsbrand Wijnands
Individual
Email: ice@braindump.be

Mankamana Mishra
Cisco Systems
Email: mankamis@cisco.com

Hooman Bidgoli
Nokia
Email: hooman.bidgoli@nokia.com

Gyan Mishra
Verizon
Email: gyan.s.mishra@verizon.com