BIER   Z. Zhang
Internet-Draft   ZTE Corporation
Intended status: Standards Track   Z. Zhang, Ed.
Expires: December 16, 2021   Juniper Networks
                                                          I. Wijnands
                                                           Individual
                                                            M. Mishra
                                                        Cisco Systems
                                                           H. Bidgoli
                                                                Nokia
                                                       G. Mishra, Ed.
                                                              Verizon
                                                        June 14, 2021

Supporting BIER in IPv6 Networks (BIERin6)
draft-ietf-bier-bierin6-00

Abstract

   BIER is a new architecture for the forwarding of multicast data
   packets without requiring per-flow state inside the network. This
   document describes how the existing BIER encapsulation specified in
   RFC 8296 works in an IPv6 non-MPLS network, 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

   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 December 16, 2021.

Copyright Notice

Table of Contents

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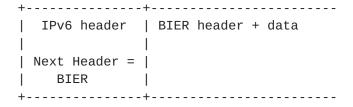
MPLS BIER forwarding in IPv6 is outside the scope of this document.

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

1.1. BIER over L2/Tunnels

[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 header" codepoint that is set to a value that means "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 "non-MPLS BIER".

```
+--------------+-----------------------
|  IPv6 header | BIER header + data
|              |
| Next Header = |
|    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:

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.

1.2. Considerations of Requirements for BIER in IPv6 Networks

[draft-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:

o 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.

o Generic Fragmentation/ESP
   [I-D.zzhang-tsvwg-generic-transport-functions]: a BFIR does
   generic fragmentation and/or ESP (without using IPv6
   encapsulation) and the resulting packets are treated as BIER
   payload.

Either way, the fragmentation/ESP is handled by a layer outside of
BIER and then the resulting packets are treated as BIER payload.

BIERin6 does support SRv6 based overlay services (e.g. MVPN/EVPN).
One of the following methods can be used (relevant overlay signaling
will be specified separately):

o An ingress PE (which is a BFIR) can encapsulate customer packets
   with an IPv6 header (with optional fragmentation and ESP extension
   headers). The destination address is a multicast locator plus the
   Fucn/Arg portion that identifies the service. That IPv6 packet is
   then treated as BIER payload. An egress PE (which is a BFER) uses
   the standard SRv6 procedures to forward the IPv6 packet that is
   exposed after the BIER header is decapsulated.

o Alternatively, since only the destination IPv6 address in the
   above-mentioned 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 IPv6 address (instead of an entire
   IPv6 header) is added between the BIER header and original
   payload.

BIERin6 being a solution based on [RFC8279] [RFC8296], ECMP is
inherently supported by BFRs using the 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 (instead of using local 5-tuple input key to a
hash function to locally generate the stateless 20-bit flow label) so
that routers along the tunnel can do ECMP based on Flow Labels. 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:

o BIER OAM specified in [I-D.ietf-bier-ping]

o BIER BFD specified in [I-D.ietf-bier-bfd]

o BIER Performance Measurement specified in [I-D.ietf-bier-pmmm-oam]

o BIER Path Maximum Transmission Unit Discovery specified in
   [I-D.ietf-bier-path-mtu-discovery]

o BIER IOAM specified in [I-D.xzlnp-bier-ioam]

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 label" field in the IPv6 packet SHOULD be copied from the
entropy field in the BIER encapsulation.

2.1. IPv6 Options Considerations

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 BFERs, 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.ietf-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.ietf-bier-lsr-ethernet-extensions](#)] is needed.

Otherwise, a node that requires IPv6 encapsulation MUST advertise the
BIER IPv6 transportation 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.

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-sub-TLV of BIER Ethernet
Encapsulation sub-TLV defined in OSPFv3, and a new sub-sub-sub-TLV of
BIER Ethernet Encapsulation sub-sub-TLV defined in ISIS, as per
[[I-D.ietf-bier-lsr-ethernet-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     |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

o Type: For OSPF, value TBD1 (prefer 1) is used to indicate it is
   the IPv6 transportation sub-TLV. For ISIS, value TBD2 (prefer 1)
   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.ietf-bier-lsr-ethernet-extensions] or redistributed across
protocol boundaries per [I-D.ietf-bier-prefix-redistribute], the BIER
IPv6 transportation 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 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-sub-
TLV" type in the "OSPFv3 BIER Ethernet Encapsulation sub-TLV"
Registry.

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

IANA is requested to allocate a new value "IPv6 Address" from "BIER
Next Protocol Identifiers" registry to indicate that an IPv6 address
is inserted between the BIER header and BIER payload.

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

[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>.

[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>.

8.2. Informative References

[I-D.ietf-bier-bar-ipa]
            Zhang, Z., Przygienda, A., Dolganow, A., Bidgoli, H.,
            Wijnands, I., and A. Gulko, "BIER Underlay Path
            Calculation Algorithm and Constraints", draft-ietf-bier-
            bar-ipa-07 (work in progress), September 2020.

   [I-D.ietf-bier-bfd]
              Xiong, Q., Mirsky, G., Hu, F., and C. Liu, "BIER BFD",
              draft-ietf-bier-bfd-01 (work in progress), April 2021.

   [I-D.ietf-bier-idr-extensions]
              Xu, X., Chen, M., Patel, K., Wijnands, I., and A.
              Przygienda, "BGP Extensions for BIER", draft-ietf-bier-
              idr-extensions-07 (work in progress), September 2019.

   [I-D.ietf-bier-ipv6-requirements]
              McBride, M., Xie, J., Geng, X., Dhanaraj, S., Asati, R.,
              Zhu, Y., Mishra, G., and Z. Zhang, "BIER IPv6
              Requirements", draft-ietf-bier-ipv6-requirements-09 (work
              in progress), September 2020.

   [I-D.ietf-bier-lsr-ethernet-extensions]
              Dhanaraj, S., Yan, G., Wijnands, I., Psenak, P., Zhang,
              Z., and J. Xie, "LSR Extensions for BIER over Ethernet",
              draft-ietf-bier-lsr-ethernet-extensions-02 (work in
              progress), December 2020.

   [I-D.ietf-bier-non-mpls-bift-encoding]
              Wijnands, I., Mishra, M., 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-03 (work in progress), November 2020.

   [I-D.ietf-bier-ospfv3-extensions]
              Psenak, P., Nainar, N. K., and I. Wijnands, "OSPFv3
              Extensions for BIER", draft-ietf-bier-ospfv3-extensions-03
              (work in progress), November 2020.

   [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", draft-ietf-bier-path-mtu-
              discovery-10 (work in progress), March 2021.

   [I-D.ietf-bier-ping]
              Kumar, N., Pignataro, C., Akiya, N., Zheng, L., Chen, M.,
              and G. Mirsky, "BIER Ping and Trace", draft-ietf-bier-
              ping-07 (work in progress), May 2020.

   [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", draft-ietf-bier-
              pmmm-oam-10 (work in progress), March 2021.

   [I-D.ietf-bier-prefix-redistribute]
              Zhang, Z., Wu, B., Zhang, Z., Wijnands, I., and Y. Liu,
              "BIER Prefix Redistribute", draft-ietf-bier-prefix-
              redistribute-00 (work in progress), August 2020.

   [I-D.xzlnp-bier-ioam]
              Min, X., Zhang, Z., Liu, Y., Nainar, N. K., and C.
              Pignataro, "Bit Index Explicit Replication (BIER)
              Encapsulation for In-situ OAM (IOAM) Data", draft-xzlnp-
              bier-ioam-01 (work in progress), January 2021.

   [I-D.zhang-bier-babel-extensions]
              Zhang, Z. and T. Przygienda, "BIER in BABEL", draft-zhang-
              bier-babel-extensions-04 (work in progress), November
              2020.

   [I-D.zzhang-tsvwg-generic-transport-functions]
              Zhang, Z., Bonica, R., and K. Kompella, "Generic Transport
              Functions", draft-zzhang-tsvwg-generic-transport-
              functions-00 (work in progress), November 2020.

   [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 (editor)
Verizon

EMail: gyan.s.mishra@verizon.com