Network Working Group   M. McBride
Internet-Draft   Futurewei
Intended status: Informational   J. Xie
Expires: March 28, 2021   X. Geng
                                                              S. Dhanaraj
                                                                   Huawei
                                                                 R. Asati
                                                                    Cisco
                                                                    Y. Zhu
                                                            China Telecom
                                                                G. Mishra
                                                             Verizon Inc.
                                                                 Z. Zhang
                                                                  Juniper
                                                       September 24, 2020

                          BIER IPv6 Requirements
                     draft-ietf-bier-ipv6-requirements-08

Abstract

   There have been several proposed solutions with BIER being used in
   IPv6. But there hasn't been a document which describes the problem
   and lists the requirements. The goal of this document is to describe
   the general BIER IPv6 encapsulation problem, summarize the
   encapsulation modes of the proposed solutions, detail solution
   requirements, and assist the working group in the development of
   acceptable solutions.

Table of Contents

1. **Introduction**

   Bit Index Explicit Replication (BIER) [RFC8279] is an architecture
   that provides optimal multicast forwarding, without requiring
   intermediate routers to maintain per-flow state, through the use of a

multicast-specific BIER header. [RFC8296] defines two types of BIER
encapsulation: one is BIER MPLS encapsulation for MPLS environments,
the other is non-MPLS BIER encapsulation to run without MPLS. This
document describes non-MPLS BIER encapsulation in IPv6 environments.
We explain the requirements of transporting IPv4/IPv6 multicast
overlay payload through an IPv6 network underlay using BIER. The
solutions may require the use of IPv6 forwarding plane and may
include IPv6 encapsulation and/or generic IPv6 tunnelling.

1.1. 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 RFC 2119 [RFC2119].

1.2. Terminology

o BIER: Bit Index Explicit Replication. Provides optimal multicast
   forwarding through adding a BIER header and removing state in
   intermediate routers.

2. Problem Statement

The problem is how to transport multicast packets, with non-MPLS BIER
encapsulation, in an IPv6 environment. We need to determine where to
put the BIER header in this IPv6 environment. With IPv6
encapsulation being increasingly used for unicast services, such as
VPN or L2VPN, it may be desirable to have IPv6 encapsulation also
used in BIER deployments for multicast services such as MVPN. It may
also be desirable to not use IPv6 encapsulation except when IPv6
tunneling (native or GRE/UDP-like) is used to transport BIER packets
over BIER-incapable routers.

Below is a simple scenario that needs BIER IPv6-based forwarding:

```
        +--------------------------------------------+
        |                                            |
        |                                  +------+
        |                                  | BFER |
  +------+        +-------+        +-----+  +------+
  | BFIR |        |Non-BFR|        | BFR |     |
  +------+        +-------+        +-----+  +------+
        |                                  | BFER |
        |               IPv6 Network       +------+
        |                                            |
        +--------------------------------------------+
```
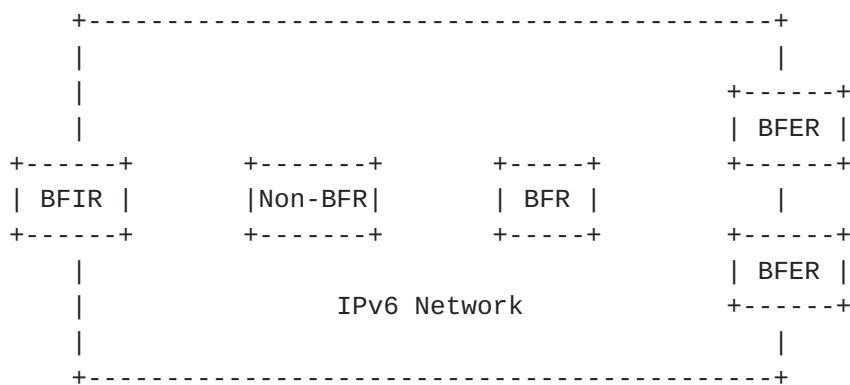
This scenario depicts the need to replicate BIER packets from a BFIR
to BFERs across an IPv6 Service Provider core. Inside the IPv6
network, the BIER header is used to direct the packet from one BFR to
the next BFRs, and either a IPv6 header or an L2/tunnel header is
used to provide reachability between BFRs. The IPv6 environment may
include a variety of link types, may be entirely IPv6, or may be dual
stack. There may be cases where not all routers are BFR capable in
the IPv6 environment but still want to deploy BIER. Regardless of
the environment, the problem is to deploy BIER, with non-MPLS BIER
encapsulation, in an IPv6 network.

[3](). Requirements

There are several suggested requirements for BIER IPv6 solutions.

In this document, the requirements are divided into two levels:
Mandatory and Optional. The requirement levels are determined based
on the following factors:

> If the requirement is required for a feature that is likely to be
> a potential deployment, the requirement level will be considered
> mandatory.

> If the impact of not implementing the requirement may block BIER
> from been deployed, the requirement level will be considered
> mandatory.

[3.1](). Mandatory Requirements

Considering that these mandatory requirements are all well-known to
the working group, and practical in normal deployment, they will be
listed without a detailed description.

[3.1.1](). Support various L2 link types

The solution should support various kinds of L2 data link types.

[3.1.2](). Support BIER architecture

The solution must support the BIER architecture.

Supporting different multicast flow overlays, multiple sub-domains,
multi-topologies, multiple sets, multiple Bit String Lengths, and
deterministic ECMP are considered essential functions of BIER and
need to be supported.

3.1.3. Support deployment with Non-BFR routers

   The solution must support deployments with BIER-incapable routers.
   This is beneficial to the deployment of BIER, especially in early
   deployments when some routers do not support BIER forwarding but
   support IPv6 forwarding.

3.1.4. Support OAM

   BIER OAM should be supported, either directly using existing methods,
   or by specifying a new method for the same functionality. It may be
   considered essential as part of the BIER architecture in some cases.

3.2. Optional Requirements

   The requirements in this section are listed as optional, and each
   requirement is explained with a detailed scenario. Note that
   fragmentation and IPSEC ESP are not BIER functions, they are provided
   by the upper IP layer.

3.2.1. Support Fragmentation

   There are some cases where the Fragmentation/Assembly function is
   needed for BIER to work in an IPv6 network.

   For example, a customer IPv6 multicast packet may be 1280 bytes and
   is required to be transported through an IPv6 network using BIER.
   Every link of the IPv6 network is no less than the requisite 1280
   bytes [RFC8200], but the size of the payload that can be encapsulated
   in BIER (BIER-MTU) is less than 1280 bytes. In this case, it is not
   the appropriate action for a BFIR to drop the packet and advertise an
   MTU to the source [RFC8296]. Instead, the IPv6 transport mechanism,
   either integrated with or independent to BIER, need to provide the
   fragmentation and assembly function.

3.2.2. Support IPSEC ESP

   There are some cases where the IPSEC ESP function may be needed to
   transport c-multicast packets through an IPv6 network with
   confidentiality using BIER technology.

   A service provider may want to provide additional security SLA to its
   customer to ensure that the unencrypted c-multicast packet is not
   altered in the service provider's network. In this case, if the BIER
   technology is preferred for the multicast service, BIER with IPSEC
   ESP support may be a candidate solution. On the other hand, the
   traffic protection may be better provided via IPSEC or MACSEC at
   multicast flow overlay over and beyond the BIER domain.

4. IANA Considerations

   Some BIER IPv6 encapsulation proposals do not require any action from
   IANA while other proposals require new IPv6 Option codepoints from
   IPv6 sub-registries, new "Next header" values, or require new IP
   Protocol codes. This document, however, does not require anything
   from IANA.

5. Security Considerations

   There are no security issues introduced by this draft.

6. Acknowledgement

   Thanks to Eric Rosen for his listed set of initial requirements on
   the BIER WG mailing list.

7. Normative References

   [I-D.pfister-bier-over-ipv6]
             Pfister, P. and I. Wijnands, "An IPv6 based BIER
             Encapsulation and Encoding", draft-pfister-bier-over-
             ipv6-01 (work in progress), October 2016.

   [I-D.xie-bier-ipv6-encapsulation]
             Xie, J., Geng, L., McBride, M., Asati, R., Dhanaraj, S.,
             Zhu, Y., Qin, Z., Shin, M., Mishra, G., and X. Geng,
             "Encapsulation for BIER in Non-MPLS IPv6 Networks", draft-
             xie-bier-ipv6-encapsulation-08 (work in progress), July
             2020.

   [I-D.zhang-bier-bierin6]
             Zhang, Z., Zhang, Z., Wijnands, I., Bidgoli, H., and M.
             McBride, "BIER in IPv6 (BIERin6)", draft-zhang-bier-
             bierin6-07 (work in progress), July 2020.

   [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>.

   [RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in
             IPv6 Specification", RFC 2473, DOI 10.17487/RFC2473,
             December 1998, <https://www.rfc-editor.org/info/rfc2473>.
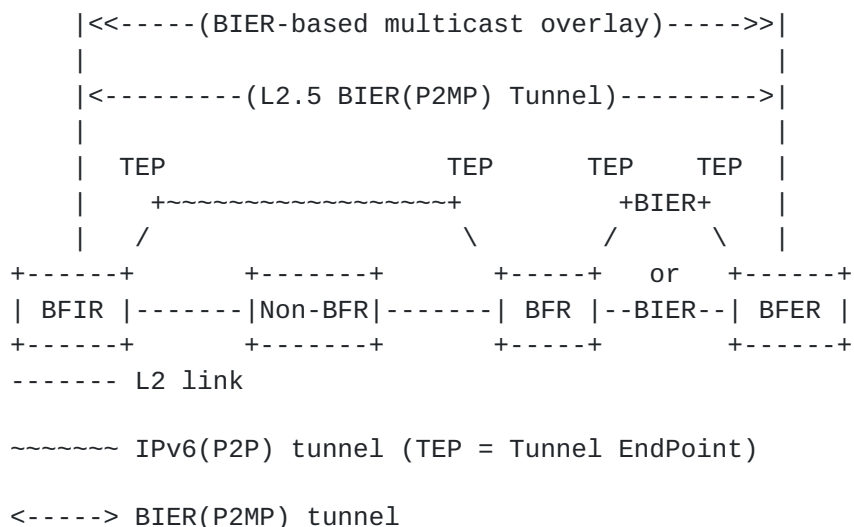
   [RFC8200]   Deering, S. and R. Hinden, "Internet Protocol, Version 6
               (IPv6) Specification", STD 86, RFC 8200,
               DOI 10.17487/RFC8200, July 2017,
               <https://www.rfc-editor.org/info/rfc8200>.

   [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>.

Appendix A. Conceptual Models For BIER IPv6 Encapsulation and Forwarding

   This analysis introduces two conceptual models for BIER in IPv6
   networks based on the experience and solutions discussed in the IETF
   community.

A.1. Independent Model

   The first conceptual model is an Independent Model, where IPv6 is
   nothing special to BIER but a transportation means that may be used
   just like other transportation means, and BIER is nothing special to
   IPv6 but a payload type just like other payload types.

```
        |<<-----(BIER-based multicast overlay)----->>|
        |                                            |
        |<---------(L2.5 BIER(P2MP) Tunnel)--------->|
        |                                            |
        |   TEP                 TEP     TEP     TEP  |
        |     +~~~~~~~~~~~~~~~~~+           +BIER+   |
        |    /                  \         /      \   |
   +------+       +-------+        +-----+   or   +------+
   | BFIR |-------|Non-BFR|-------| BFR |--BIER--| BFER |
   +------+       +-------+        +-----+        +------+
   ------- L2 link

   ~~~~~~~ IPv6(P2P) tunnel (TEP = Tunnel EndPoint)

   <-----> BIER(P2MP) tunnel
```
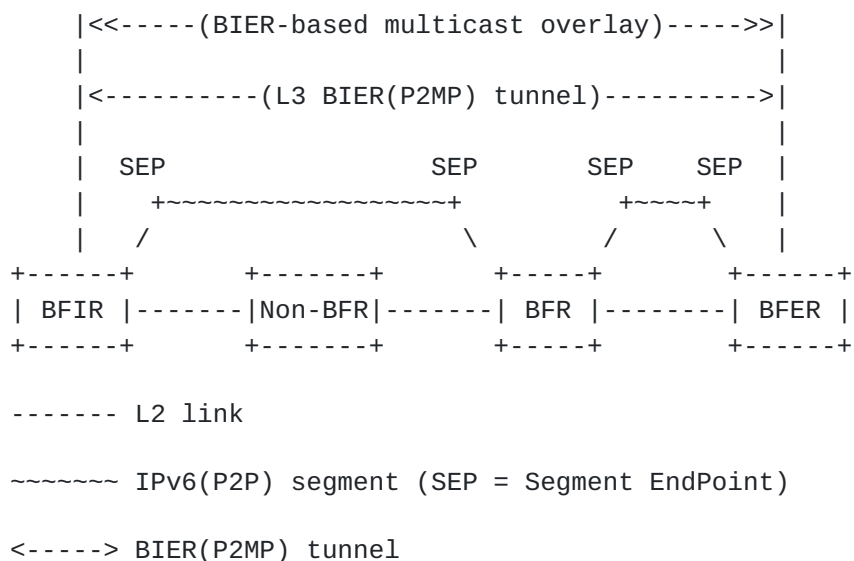
In this model, an IPv6 tunnel works as a link-layer of BIER, and BIER
works as a layer-2.5 over tunnels or L2 links. Between two BFRs,
either a L2 link can be used directly or any tunnel (IPv6 or not) can
be used for BIER transport. In the tunnel case, the transmitting BFR
adds tunnel encapsulation (e.g. IPv6 header) and the receiving BFR
removes the tunnel encapsulation.

General consideration of this model is to keep BIER and IPv6
independent of each other. The BIER header is not part of the IPv6
header but comes after the transport header (L2 or tunnel header) and
before BIER payload.

[A.2](). Integrated Model

The second conceptual model is an Integrated Model that integrates
BIER as part of the IPv6 data plane, making it a "Layer-3 BIER"
approach.

```
        |<<-----(BIER-based multicast overlay)----->>|
        |                                            |
        |<----------(L3 BIER(P2MP) tunnel)---------->|
        |                                            |
        |   SEP               SEP        SEP    SEP   |
        |     +~~~~~~~~~~~~~~~~~+         +~~~~+       |
        |    /                  \       /     \       |
  +------+         +-------+       +-----+        +------+
  | BFIR |-------|Non-BFR|-------| BFR |--------| BFER |
  +------+         +-------+       +-----+        +------+


    ------- L2 link

    ~~~~~~~ IPv6(P2P) segment (SEP = Segment EndPoint)

    <-----> BIER(P2MP) tunnel
```

In this model, BIER works as part of the IPv6 data plane. The BFIR
and BFERs work as IPv6 (P2MP) tunnel endpoints, and BFRs work as IPv6
segment endpoints. The BIER header is processed on each segment
endpoint and there is no decapsulation, or re-encapsulation, on the
segment endpoints.

This model typically needs an IPv6 extension header to carry the BIER
header. and processing of the BIER header (e.g., the BitString) will
be implemented as part of the IPv6 extension header processing. The
IPv6 source address is the BIER packet source-origin identifier, and
is unchanged through the BIER domain from BFIR to BFERs.

General consideration of this model is to use the IPv6 capabilities
integrated, in addition to normal BIER function, to facilitate new
requirements that may emerge in an IPv6 network.

Appendix B. List of Solutions

There have been some proposed solutions for BIER in IPv6
environments. Some solutions propose encoding while others propose
encapsulation. It is recommended for the wg to evaluate these
solutions, against the requirements listed previously, in order to
make informed decisions on solution readiness.

This section lists these solutions categorizing in the two conceptual
models.

B.1. Integrated mode approach

One example of this model is defined in [I-D.pfister-bier-over-ipv6],
where the information required for BIER forwarding, e.g., the
BitString, is encoded in the low-order bits of the IPv6 destination
address of each packet. The high-order bits of the IPv6 destination
address are used by intermediate routers for unicast forwarding,
deciding whether a packet is a BIER packet, and if so, to identify
the BIER Sub-Domain, Set Identifier and BitString length. The BIER
function is integrated in the IPv6 header and its forwarding
procedure, and the BIER payload is encapsulated as the IPv6 payload.

```
    +---------------+-------------------
    |  IPv6 header  | payload
    | (BitString in |
    | DA lower bits)|
    |  Next Header  |
    +---------------+-------------------
```

Another example of this model is defined in
[I-D.xie-bier-ipv6-encapsulation], where information required for
BIER forwarding, e.g., the BIER header, is encoded in an Option TLV
(indicated by an Option Type to be allocated by IANA) of the IPv6
Destination Option Header. The third-highest-order bit of the Option
Type is set to 1 to allow Option Data (e.g., the BitString) change en
route. The BIER function is integrated in IPv6 extension header and
its forwarding procedure, and the BIER payload is encapsulated as the
IPv6 payload.

```
+--------------+----------------+------------------
|  IPv6 header | IPv6 Ext header | payload
|              | (BIER header in |
|              |  TLV Type = X)  |
| Next Header  |   Next Header   |
+--------------+----------------+------------------
```

B.2. Independent model approach

One example of this model is defined in [I-D.zhang-bier-bierin6],
where the BIER header and the payload following it are L2 payload
when feasible (e.g. when two BFRs are directly connected) or IPv6
payload when IPv6 transport is needed/desired (e.g. when two BFRs are
not directly connected). This is indicated by either a 0xAB37
Ethertype allocated to BIER or a new IPv6 Next-Header value to be
allocated by IANA.

```
+--------------+----------------+------------------
|   Ethernet   |   BIER header  | payload
|  (ethType =  | (BIFT-id, ...) |
|    0xAB37)   |                |
|              |   Next Header  |
+--------------+----------------+------------------


+--------------+----------------+------------------
|  IPv6 header | IPv6 Ext header | BIER Hdr + payload
|              |     (optional)  | as IPv6 payload
|              |                 |
| Next Header  | Next Header = X |
+--------------+----------------+------------------
```

While not specified in [I-D.zhang-bier-bierin6], any other tunnel
types supported by the IPv6 environment could be used, e.g. IPv6
GRE/UDP:

```
+--------------+----------------+-----------+---------------
|  IPv6 header | IPv6 Ext header | GRE header |
|              |     (optional)  |           | BIER Hdr +
|              |                 |           | payload as GRE
| Next Header  |   Next Header   |Proto=0xAB37| Payload
+--------------+----------------+-----------+---------------


+--------------+----------------+-----------+---------------
|  IPv6 header | IPv6 Ext header | UDP header |
|              |     (optional)  |           | BIER Hdr +
|              |                 |           | payload as UDP
| Next Header  |Next Header =UDP | DPort=TBD | Payload
+--------------+----------------+-----------+---------------
```

Authors' Addresses

   Mike McBride
   Futurewei

   Email: michael.mcbride@futurewei.com

   Jingrong Xie
   Huawei

   Email: xiejingrong@huawei.com

   Xuesong Geng
   Huawei

   Email: gengxuesong@huawei.com

   Senthil Dhanaraj
   Huawei

   Email: senthil.dhanaraj@huawei.com

   Rajiv Asati
   Cisco

   Email: rajiva@cisco.com

   Yongqing Zhu
   China Telecom

   Email: zhuyq8@chinatelecom.cn

   Gyan Mishra
   Verizon Inc.

   Email: gyan.s.mishra@verizon.com

   Zhaohui Zhang
   Juniper

   Email: zzhang@juniper.net