Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 11, 2021

M. McBride
Futurewei
J. Xie
S. Dhanaraj
Huawei
R. Asati
Cisco
Y. Zhu
China Telecom
G. Mishra
Verizon Inc.

July 10, 2020

BIER IPv6 Requirements
draft-ietf-bier-ipv6-requirements-05

Abstract

The BIER WG charter includes work on developing "a mechanism to use BIER natively in IPv6". There have been several proposed solutions in this area. But there hasn't been a document which describes the problem and lists the requirements. The goal of this document is to describe the BIER IPv6 requirements, summarize the encapsulation modes of the proposed solutions, guide the working group in understanding the benefits and drawbacks of the various solutions, and help in the development of acceptable solutions.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction 3
1.1. Requirements Language 3
1.2. Terminology 3
2. Problem Statement 3
3. Conceptual Models For BIER IPv6 Encapsulation and Forwarding 4
3.1. Transport-Independent Model 5
3.2. Native IPv6 Model 6
3.3. Encapsulation Approaches Considered 7
4. Requirements 7
4.1. Mandatory Requirements 8
4.1.1. L2 Agnostic 8
4.1.2. Support BIER architecture 8
4.1.3. Conform to existing IPv6 Spec 8
4.1.4. Support deployment with Non-BFR routers 8
4.1.5. Support inter-AS multicast deployment 8
4.1.6. Support Simple Encapsulation 9
4.1.7. Support Deployment Security 9
4.2. Optional Requirements 9
4.2.1. Support MVPN 9
4.2.2. Support OAM 9
4.2.3. Support IPSEC 9
4.2.4. Support Fragmentation 9
4.2.5. Support hardware fast path 10
5. IANA Considerations 10
6. Security Considerations 10
7. Acknowledgement 10
8. Normative References 10
Appendix A. Solutions Evaluation 12
A.1. BIER-ETH encapsulation in IPv6 networks 12
A.2. Encode Bitstring in IPv6 destination address 13
A.3. Add BIER header into IPv6 Extension Header 13

A.4. Transport BIER as IPv6 payload 14
A.5. Tunnelling BIER in a IPv6 tunnel 15
Authors' Addresses 15

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 to run on physical links: one is BIER MPLS encapsulation to run on various physical links that support MPLS, the other is non-MPLS BIER Ethernet encapsulation to run on ethernet links, with an ethertype 0xAB37. This document describes using BIER in non-MPLS IPv6 environments. We explain the requirements of transporting IPv4/IPv6 multicast payloads through an IPv6 network using "BIER natively in IPv6". As clarified in the working-group, "BIER natively in IPv6" means BIER not encapsulated in MPLS or Ethernet. This may include native IPv6 encapsulation and generic IPv6 tunnelling. The goal of this document is to help the BIER WG evaluate the BIER v6 requirements and solutions in order to begin adopting solution drafts.

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

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

- BUM: Broadcast, Unknown Unicast, Multicast. Term used to describe the three types of Ethernet modes that will be forwarded to multiple destinations

2. Problem Statement

The problem is the ability of the network to transport BUM packets, with BIER headers, in an IPv6 environment. In many IPv6 network deployments, non-MPLS encapsulation is used for unicast as the data-plane and it is likewise expected to have BIER IPv6 deployments which depend on these same unicast technologies.

One such case involves supporting a non-BFR router in a network as described in section 6.9 of RFC8279. In the context of this document, an IPv6 based unicast tunnel is needed to support such deployment where a non-BFR exists. Another case is to support inter-AS multicast deployment as illustrated in [I-D.geng-bier-ipv6-inter-domain]. In such deployment, there are non-BFR routers, or even an entire non-BIER network, that needs the ability to traverse from one BFR to another. [I-D.ietf-bier-use-cases] shows it is possible there are other cases where inter-AS multicast deployment is required.

As with IPv6, another problem of BIER IPv6 technology may be "Transition Mechanisms and Partial Deployments" which is listed as the No.1 charter item of BIER WG. Therefore, a basic requirement of BIER IPv6 is to leverage IPv6 reachability for incremental and inter-AS BIER deployment.

Below is a simple scenario that needs BIER IPv6 encapsulation and forwarding:

```
         +--------------------------------------------+
         |                                            |
         |                                         +------+
         |                                         | BFER |
     +------+       +-------+       +-----+        +------+
     | BIER |       |Non-BFR|       | BFR |           |
     +------+       +-------+       +-----+        +------+
         |                                         | BFER |
         |                IPv6 Network             +------+
         |           (intra-AS or inter-AS)           |
         +--------------------------------------------+
```

This scenario depicts the need to replicate bier packets from a BFIR to BFERs across an IPv6 core. The IPv6 environment may include a variety of link types, may be entirely IPv6, may be dual stack or any type of combination which includes IPv6. Regardless of the environment, there are times when a BIER header, including the BIER BitString used to determine the set of BIER forwarding egress routers, will need to traverse a IPv6 domain. The ways in which BIER will function in an IPv6 environment is the problem that needs to be solved.

3. Conceptual Models For BIER IPv6 Encapsulation and Forwarding

This analysis introduces two conceptual models for BIER IPv6 encapsulation and forwarding based on the experience and examples that have been seen in the IETF community.

3.1. Transport-Independent Model

The first conceptual model is a Transport-Independent Model that views IP tunnels as links of BIER, and views BIER as an independent "Layer-2.5".

```
         |<----------(L2.5 BIER(P2MP) Tunnel)-------->|
         |                                            |
         |    +~~~~~~~~~~~~~~~~~~+          +~~~~+    |
         |   /                    \        /      \   |
     +------+       +-------+       +-----+        +------+
     | BFIR |-------|Non-BFR|-------| BFR |--------| BFER |
     +------+       +-------+       +-----+        +------+

     ------- physical link

     ~~~~~~~ IPv6(P2P) tunnel

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

In this model, an IPv6 tunnel works as a link-layer of BIER, and BIER works as a transport-independent layer (or layer-2.5) over a virtual-link (IPv6 tunnel). On each BFR, the IPv6 tunnel of the receiving packet is decapsulated, and a new IPv6 tunnel is encapsulated before sending the packet to the next-hop BFR neighbour.

From the view of the IPv6 layer, the BIER header is a kind of Upper-layer header (Layer-4). From the view of the BIER layer, the IPv6 encapsulation is a tunnel working as a "link" of BIER. With an End-to-End view, the tunnel from BFIR to BFERs is a Layer-2.5 BIER (P2MP) tunnel, and the BFIR-id is the BIER packet source-origin identifier, and is unchanged through the BIER domain from BFIR to BFERs.

This model is similar to the "MPLS over IP" [RFC4023] or "MPLS over UDP" [RFC7510] approach. A more general output of such approach in IETF is "MPLS Segment Routing over IP" [RFC8663]. It makes use of IPv4/IPv6 tunnel, IPv4/IPv6 UDP tunnel and IPv4/IPv6 GRE tunnel to encapsulate the MPLS-based instructions. In fact, BIER-MPLS could use this approach directly since BIER-MPLS is based on MPLS.

There may be, however, in certain cases some difficulty with allocation of an MPLS label and advertisement through the control-plane. For example, a simple inter-AS BIER deployment may want to use the auto-configuration of BIFT-id using Non-MPLS BIER encapsulation [RFC8296] as illustrated in [I-D.geng-bier-ipv6-inter-domain]. This brings the need of a new "Next Header" value to indicate the "Non-MPLS" BIER header.

For IPv4/IPv6 GRE, the "Next Header" is the 16-bit "Protocol Type" field, and has adequate space for such requirement.

For IPv4/IPv6 UDP, the "Next Header" is the 16-bit "Destination Port" field, and has adequate space for such requirement.

For IPv4/IPv6, the "Next Header" is a 8-bit value and needs to be allocated from the "Assigned Internet Protocol Numbers" registry.

Reassembly/Re-fragmentation of a packet has to be executed on each BFR in such case. This may be common and even friendly for a protocol stack in a BFR software implementation, but it may impose cost for a BFR hardware implementation.

IPv6 functions that are expected to be executed from BFIR to BFER are assumed to be broken on the BFRs, for example, IPv6 Fragmentation/Assembly or IPSEC ESP. This is because the "IPv6 tunnel" and all its functions is "terminated" on the BFRs. These functions, if desired, may need to be re-designed in the "Layer-2.5" BIER mode.

For deployment security, it is necessary to ensure the "BIER" packet is only using the allowed IPv6 tunnel.

3.2. Native IPv6 Model

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

```
          |<----------(L3 BIER(P2MP) tunnel)--------->|
          |                                           |
     +-------+      +--------+      +------+       +-------+
     | BFIR |-------|Non-BFR|-------| BFR |--------| BFER |
     +-------+      +--------+      +------+       +-------+

     ------- physical link

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

In this model, BIER works as part of the IPv6 data plane. BFIR and BFERs work as IPv6 (P2MP) tunnel endpoints, and BFRs work as IPv6 segment endpoints. On each BFR, the segment endpoint behaviour of IPv6 data plane is executed, and there is no decapsulation of receiving IPv6 tunnel and encapsulation of new IPv6 tunnel for sending.

In this mode, BIER is integrated into the IPv6 data plane. The IPv6 source address is the BIER packet source-origin identifier, and is unchanged through the BIER domain from BFIR to BFERs.

This model is similar to many examples emerging in the IETF community which soley use the IPv6 data plane. SRv6 introduced in [RFC8754] and [I-D.ietf-spring-srv6-network-programming] is an example. The benefits of such approach includes reducing the number of encapsulation layers, capability of deployment with non-capable routers in a network, extending the technology in a wider inter-AS scope using IP reachability, and capability of integrating the functions of the IPv6 data plane.

This model typically needs an extension to IPv6 data plane, with an IPv6 extension header or Option introduced.

IPv6 functions that are expected to be executed from BFIR to BFER is supported if correctly designed, for example, IPv6 Fragmentation/ Assembly or IPSEC ESP.

For deployment security, it is necessary to ensure the "BIER" packet is in a trusted IPv6-based domain.

3.3. Encapsulation Approaches Considered

A number of approaches to the design of BIER-IPv6 encapsulation were investigated by the BIER Working Group and were discussed in IETF meetings and on the BIER list. This section divides these approaches into the two conceptual models.

Transport-Independent Model approaches include:

Transport-Independent BIER [I-D.xu-bier-encapsulation].

BIERin6 [I-D.zhang-bier-bierin6].

Native IPv6 Model approaches include:

BIER-over-IPv6 [I-D.pfister-bier-over-ipv6].

BIERv6 [I-D.xie-bier-ipv6-encapsulation].

4. Requirements

There have been several suggested requirements, on the BIER email list and in meetings, which have been used to form BIER IPv6 requirements used to help the wg evaluate against the proposed

solutions. There is also many further discussions on the list about BIER IPv6 requirements from different scenarios.

Considering that the importance of requirement for BIER IPv6 solution is different, in this document, the requirements are divided into two groups: mandatory and optional. The requirements in the mandatory group are considered necessary for any BIER IPv6 solution, while the requirements in the optional group should be considered but are not mandatory.

4.1. Mandatory Requirements

4.1.1. L2 Agnostic

The solution must be agnostic to the underlying L2 data link type. The solution needs to support P2P ethernet links as well as shared media ethernet links without requiring the LAN switch to perform BIER snooping.

4.1.2. Support BIER architecture

The solution must support the BIER architecture.

Multiple sub-domains bound to one or many topologies or algorithms, multiple sets for more BFERs, multiple Bit String Lengths for different forwarding capabilities, and multiple BIFTs for ECMP are considered essential functions of BIER and need to be supported.

4.1.3. Conform to existing IPv6 Spec

The proposed encapsulation must conform to the IPv6 specification and guidelines as described in RFC8200. If new extensions to existing IPv6 specification are required, it needs to be discussed and reviewed by the 6man working-group.

4.1.4. Support deployment with Non-BFR routers

The solution must support deployments with Non-BFR 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. This is also the No.1 charter item, "Transition Mechanisms and Partial Deployments" of the BIER WG.

4.1.5. Support inter-AS multicast deployment

Inter-AS multicast support is needed for ease of provisioning the P2MP transport service to enterprises. This could greatly increase

the scalability of BIER, as it is usually considered to be suitable only for small intra-AS scenarios.

4.1.6. Support Simple Encapsulation

The solution must avoid requiring different encapsulation types. A solution needs to do careful trade-off analysis and select one encapsulation as its proposal for best coverage of various scenarios.

4.1.7. Support Deployment Security

The proposed solution must include careful security considerations, including all that is already considered in BIER architecture RFC8279 and RFC8296, and other security concerns that may raise due to the addition of IPv6.

4.2. Optional Requirements

4.2.1. Support MVPN

The solution MAY support MVPN services that is defined in [RFC6513]. When MVPN is supported, it should work in a "tunnel" mode, encapsulating IP or IPv6 multicast packet in an outer IPv6 header. When MVPN is supported, it is suggested to think about both intra-AS and inter-AS deployment.

4.2.2. Support OAM

BIER OAM MAY be supported, either directly using existing method, or specify some variant method for the same function. It may be considered essential as part of the BIER architecture in some cases.

4.2.3. Support IPSEC

IPSEC is optional to IPv6 and multicast. It is preferred to support IPSEC, including AH/ESP. If IPSEC is to be supported, it shouldn't require hop-by-hop encryption/decryption.

4.2.4. Support Fragmentation

As part of IPv6 specification [RFC8200], BIER IPv6 may support fragmentation on BFIR and assembly on BFER. Support of Fragmentation can enhance the capability of BIER leveraging the BIER-MTU as introduced in section 3 of [RFC8296]. If Fragmentation is to be supported, it shouldn't require fragmentation and re-assembly at each hop.

4.2.5. Support hardware fast path

If a proposed solution is intended for some scenarios like service-provider networks, it should enable the processing and forwarding of BIER packets in hardware fast path.

5. IANA Considerations

Some BIERv6 encapsulation proposals do not require any action from IANA while other proposals require new BIER Destination 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.

6. Security Considerations

There are no security issues introduced by this draft.

7. Acknowledgement

Thank you to Eric Rosen for his listed set of requirements on the bier wg list.

8. Normative References

[I-D.geng-bier-ipv6-inter-domain]
Geng, L., Xie, J., McBride, M., and G. Yan, "Inter-Domain Multicast Deployment using BIERv6", draft-geng-bier-ipv6-inter-domain-01 (work in progress), January 2020.

[I-D.ietf-bier-use-cases]
Nainar, N., Asati, R., Chen, M., Xu, X., Dolganow, A., Przygienda, T., Gulko, A., Robinson, D., Arya, V., and C. Bestler, "BIER Use Cases", draft-ietf-bier-use-cases-11 (work in progress), March 2020.

[I-D.ietf-spring-srv6-network-programming]
Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 Network Programming", draft-ietf-spring-srv6-network-programming-16 (work in progress), June 2020.

[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., and X. Geng, "Encapsulation for BIER in Non-MPLS IPv6 Networks", draft-xie-bier-ipv6-encapsulation-07 (work in progress), June 2020.

[I-D.xu-bier-encapsulation]
Xu, X., somasundaram.s@alcatel-lucent.com, s., Jacquenet, C., Raszuk, R., and Z. Zhang, "A Transport-Independent Bit Index Explicit Replication (BIER) Encapsulation Header", draft-xu-bier-encapsulation-06 (work in progress), September 2016.

[I-D.zhang-bier-bierin6]
Zhang, Z., Przygienda, T., Wijnands, I., Bidgoli, H., and M. McBride, "BIER in IPv6 (BIERin6)", draft-zhang-bier-bierin6-04 (work in progress), January 2020.

[RFC1112] Deering, S., "Host extensions for IP multicasting", STD 5, RFC 1112, DOI 10.17487/RFC1112, August 1989, <https://www.rfc-editor.org/info/rfc1112>.

[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>.

[RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, DOI 10.17487/RFC4023, March 2005, <https://www.rfc-editor.org/info/rfc4023>.

[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>.

[RFC7510] Xu, X., Sheth, N., Yong, L., Callon, R., and D. Black, "Encapsulating MPLS in UDP", RFC 7510, DOI 10.17487/RFC7510, April 2015, <https://www.rfc-editor.org/info/rfc7510>.

[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>.

[RFC8663] Xu, X., Bryant, S., Farrel, A., Hassan, S., Henderickx, W., and Z. Li, "MPLS Segment Routing over IP", RFC 8663, DOI 10.17487/RFC8663, December 2019, <https://www.rfc-editor.org/info/rfc8663>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

Appendix A. Solutions Evaluation

The following are solutions that have been proposed to solve 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.

As illustrated in these examples, the BIER header, or the BitString, may appear in the IPv6 Header, IPv6 Extension Header, IPv6 Payload, or IPv6 Tunnel Packet:

A.1. BIER-ETH encapsulation in IPv6 networks

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

BIER-ETH encapsulation (BIER header for Non-MPLS networks as defined in [RFC8296]) can be used to transport the multicast data in the IPv6 network by encapsulating the multicast user data payload within the BIER-ETH header. However, BIER-ETH in IPv6 networks is not considered to be a "BIER natively in IPv6" solution which utilizes the IPv6 header to forward the packet.

Mixed use of BIER-ETH in a native IPv6 solution is up to the solution and is outside the scope of this document.

A.2. Encode Bitstring in IPv6 destination address

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

As described in [I-D.pfister-bier-over-ipv6], The information required by BIER is stored in the destination IPv6 address. The BIER 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. No additional extension or encapsulation header is required. Instead of encapsulating the packet in IPv6, the payload is attached to the BIER IPv6 header and the IPv6 protocol number is set to the type of the payload. If the payload is UDP, the UDP checksum needs to change when the BitString in the IPv6 destination address changes.

A.3. Add BIER header into IPv6 Extension Header

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

According to [RFC8200] In IPv6, optional internet-layer information is encoded in separate headers that may be placed between the IPv6 header and the upper- layer header in a packet. There is a small number of such extension headers, each one identified by a distinct Next Header value. An IPv6 packet may carry zero, one, or more extension headers, each identified by the Next Header field of the preceding header. Extension headers (except for the Hop-by-Hop

Options header) are not processed, inserted, or deleted by any node along a packet's delivery path, until the packet reaches the node (or each of the set of nodes, in the case of multicast) identified in the Destination Address field of the IPv6 header. The Hop-by-Hop Options header is not inserted or deleted, but may be examined or processed by any node along a packet's delivery path, until the packet reaches the node (or each of the set of nodes, in the case of multicast) identified in the Destination Address field of the IPv6 header. The Hop-by-Hop Options header, when present, must immediately follow the IPv6 header. Its presence is indicated by the value zero in the Next Header field of the IPv6 header.

Two of the currently-defined extension headers are the Hop-by-Hop Options header and the Destination Options header which carry a variable number of type-length-value (TLV) encoded "options".

In [I-D.xie-bier-ipv6-encapsulation] an IPv6 BIER Destination Option is carried by the IPv6 Destination Option Header (indicated by a Next Header value 60). It is initialized in a packet sent by an IPv6 BFIR router to inform the following BFR routers in an IPv6 BIER domain to replicate to destination BFER routers hop-by-hop. BIER is generally a hop-by-hop and one-to-many architecture and it is required for a BIER IPv6 encapsulation to include the BIER Header ([RFC8296]) as an IPv6 Extension Header, to pilot the hop-by-hop BIER replication.

Hop by hop Options Headers may be considered. The Hop-by-Hop Options header is used to carry optional information that may be examined and processed by every node along a packet's delivery path. The Hop-by-Hop Options header is identified by a Next Header value of 0 in the IPv6 header.

Defining New Extension Headers and Options may also be considered, if the IPv6 Destination Option Header is not good enough and new extension headers can solve the problem better.

Such proposals may include requests to IANA to allocate a "BIER Option" code from "Destination Options and Hop-by-Hop Options", and/or a "BIER Option Header" code from "IPv6 Extension Header Types".

A.4. Transport BIER as IPv6 payload

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

There is a proposal for a transport-independent BIER encapsulation header which is applicable regardless of the underlying transport technology. As described in [I-D.xu-bier-encapsulation] and [I-D.zhang-bier-bierin6], the BIER header, and the payload following it, can be combined as an IPv6 payload, and be indicated by a new Upper-layer IPv6 Next-Header value. A unicast IPv6 destination address is used for the replication and changes when replicating a packet out to a neighbor.

Such proposals may include a request to IANA to allocate an IPv6 Next-Header code from "Assigned Internet Protocol Numbers".

A.5. Tunnelling BIER in a IPv6 tunnel

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

A generic IPv6 Tunnel could be used to encapsulate the bier packet within an IPv6 domain.

GRE is a mechanism by which any ethernet payload can be carried by an IP GRE tunnel due to the 16-bits 'Protocol Type' field. Both IPv4 and IPv6 can be used to carry GRE. The Ethernet type codepoint 0xAB37, defined for BIER, can be used in a GRE header to indicate the subsequent BIER header and payload in an IPv6 network.

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

UDP-based tunnelling is another mechanism which uses a specific UDP port to indicate a UDP payload format. Both IPv4 and IPv6 can support UDP. Such UDP-based tunnels can be used for BIER in a IPv6 network by defining a new UDP port to indicate the BIER header and payload.

Authors' Addresses

Mike McBride
Futurewei

Email: michael.mcbride@futurewei.com

Jingrong Xie
Huawei

Email: xiejingrong@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 S. Mishra
Verizon Inc.

13101 Columbia Pike

Silver Spring
,

MD 20904

United States of America

Phone:
301 502-1347

Email:
gyan.s.mishra@verizon.com