Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 2, 2020

J. Xie
Huawei Technologies
L. Geng
China Mobile
M. McBride
Futurewei
R. Asati
Cisco
S. Dhanaraj
G. Yan
Y. Xia
Huawei

July 1, 2019

Encapsulation for BIER in Non-MPLS IPv6 Networks
draft-xie-bier-ipv6-encapsulation-02

Abstract

This document proposes a BIER IPv6 (BIERv6) encapsulation for Non-MPLS IPv6 Networks using the IPv6 Destination Option extension header.

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] and [RFC8174].

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 2, 2020.

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 any per-flow state by using a multicast-specific BIER header.

[RFC8296] defines a common BIER Header format for MPLS and Non-MPLS networks. It has defined two types of encapsulation methods using the common BIER Header, (1) BIER encapsulation in MPLS networks, here-in after referred as MPLS BIER Header in this document and (2) BIER encapsulation in Non-MPLS networks, here-in after referred as Non-MPLS BIER Header in this document. [RFC8296] also assigned

Ethertype=0xAB37 for Non-MPLS BIER Header packets to be directly carried over the Ethernet links.

This document proposes a BIER IPv6 encapsulation for Non-MPLS IPv6 Networks, defining a method to carry the standard Non-MPLS BIER header (as defined in [RFC8296]) in the native IPv6 header. A new IPv6 Option type - BIER Option is defined to encode the standard Non-MPLS BIER header and this newly defined BIER Option is carried under the Destination Options header of the native IPv6 Header [RFC8200].

This document details one of the proposed solutions for transporting BIER packets in an IPv6 network. To better understand the overall BIER IPv6 problem space, use cases and proposed solutions, refer to [I-D.ietf-bier-ipv6-requirements].

2. Terminology

Readers of this document are assumed to be familiar with the terminology and concepts of the documents listed as Normative References.

The following new terms are used throughout this document:

- BIERv6 - BIER IPv6.
- BIER Option - An Option type carried in IPv6 Destination Options Header which includes the standard Non-MPLS BIER Header.
- BIERv6 Header - An IPv6 Header with BIER Option.
- BIERv6 Packet - An IPv6 packet with BIERv6 Header. Such an IPv6 packet typically carries the user multicast payload and is forwarded by BFRs in the BIERv6 network towards the multicast receivers.
- BIER Multicast Address - A well-known multicast address used as a Destination Address in the BIERv6 Header to forward the packets to other BFRs in BIERv6 network.

3. BIER IPv6 Encapsulation

3.1. BIER Option in IPv6 Destination Options Header

Destination Options Header and the Options that can be carried under this extension header is defined in [RFC8200]. This document defines a new Option type - BIER Option, to encode the Non-MPLS BIER header. As specified in Section 4.2 [RFC8200], the BIER Option follows type-length-value (TLV) encoding format and the standard Non-MPLS BIER

header [RFC8296] is encoded in the value portion of the BIER Option TLV.

This BIER Option MUST be carried only inside the IPv6 Destination Options header and MUST NOT be carried under the Hop-by-Hop Options header.

Co-existence of Destination Options Header with BIER option TLV and other IPv6 extension headers MUST confirm to the general requirements defined in [RFC8200]. In addition to the requirements defined in [RFC8200], this document requires that the Destination Options Header with a BIER Option TLV MUST appear only after the Routing Header if the Routing Header is present in the IPv6 Header.

The BIER Option is encoded in type-length-value (TLV) format as follows:

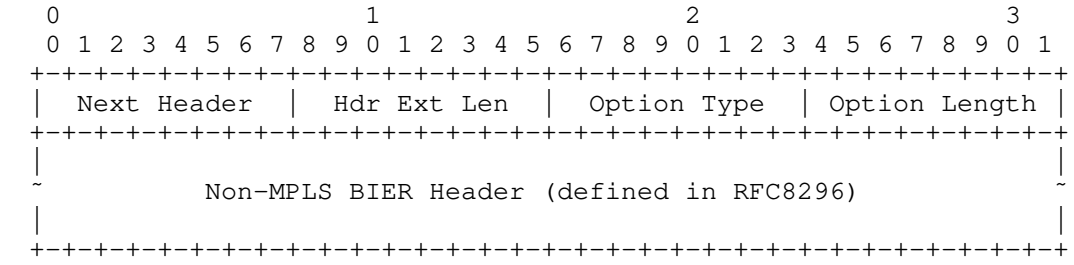

Next Header 8-bit selector. Identifies the type of header immediately following the Destination Options header.

Hdr Ext Len 8-bit unsigned integer. Length of the Destination Options header in 8-octet units, not including the first 8 octets.

Option Type To be allocated by IANA. See section 6.

Option Length 8-bit unsigned integer. Length of the option, in octets, excluding the Option Type and Option Length fields.

Non-MPLS BIER Header The Non-MPLS BIER Header defined in RFC8296. Fields in the Non-MPLS BIER Header MUST be encoded as below.

BIFT-id: The BIFT-id is a domain-wide unique value in Non-MPLS IPv6 encapsulation. See Section 2.2 of RFC 8296.

TC: SHOULD be set to binary value 000 upon transmission and MUST be ignored upon. See Section 2.2 of RFC 8296.

S bit: SHOULD be set to 1 upon transmission, and MUST be ignored upon reception. See Section 2.2 of RFC 8296.

TTL: MUST be set to 0 upon transmission, and MUST be ignored upon reception. The function of TTL is replaced by the Hop Limit field in IPv6 header.

Nibble: SHOULD be set to 0000 upon transmission, and MUST be ignored upon reception. See Section 2.2 of RFC 8296.

Ver: MUST be set to 0 upon transmission, and MUST be discarded when it is not 0 upon reception. See Section 2.2 of RFC 8296.

BSL: See Section 2.1.2 of RFC 8296.

Entropy: See Section 2.1.2 of RFC 8296.

OAM: See Section 2.1.2 of RFC 8296.

Rsv: See Section 2.1.2 of RFC 8296.

DSCP: SHOULD be set to binary value 000000 upon transmission and MUST be ignored upon reception. In IPv6 BIER encapsulation, uses highest 6-bit of Traffic Class field of IPv6 header to hold a Differentiated Services Codepoint [RFC2474].

Proto: SHOULD be set to 0 upon transmission and MUST be ignored upon reception. In IPv6 BIER encapsulation, the functionality of this 6-bit Proto field is replaced by the Next Header field in Destination Options header, which is the last IPv6 extension header, to indicate the BIER payload, which is also IPv6 payload.

For BIER Proto 1, indicating a Downstream-assigned MPLS payload, use Next Header value 137.

For BIER Proto 2, indicating an Upstream-assigned MPLS payload, there is no Next Header code currently. An upstream-assigned MPLS label within the context of special BFIR router, which in turn is represented by the BFIR-id and the Sub-domain indirectly indicated by the BIFT-id in a BIER-MPLS or BIER-ETH packet, can be replaced by an IPv6 source address in a BIER IPv6 encapsulation packet in a direct manner. In this case, use Next Header value 4 for IPv4 payload, or value 41 for IPv6 payload.

For BIER Proto 3, indicating an Ethernet payload, use Next Header value 97.

For BIER Proto 4, indicating an IPv4 payload, use Next Header value 4.

For BIER Proto 5, indicating a BIER-OAM payload, use Next Header value 58. How the BIER-PING is supported with BIER IPv6 encapsulation is outside the scope of this document.

For BIER Proto 6, indicating an IPv6 payload, use Next Header value 41.

BFIR-id: See Section 2.1.2 of RFC 8296.

BitString: See Section 2.1.2 of RFC 8296.

3.2. Multicast and Unicast Destination Address

BIER is generally a hop-by-hop and one-to-many architecture, and thus the IPv6 Destination Address (DA) being a Multicast Address is a proper approach for both the two paradigms in BIERv6 encapsulation.

This document proposes to use multicast address FF0X::AB37 (to be allocated and reserved by IANA - See Section 6.2) as the IPv6 destination address for the BIERv6 packets to be forwarded in the BIER domain.

All the interfaces of the BFRs supporting the BIERv6 encapsulation defined in this document MUST subscribe and listen to BIER multicast address FF0X::AB37 belong to scopes [1, 2, 3, 4, 5, E] defined in [RFC7346]. However it is RECOMMENDED to use Realm-Local scope (scope value 3), that is FF03:AB37 as a destination address while forwarding the BIERv6 packet, as this scope zone is exactly the BIERv6 Domain. The use of other scopes is outside the scope of this document.

Use of a Unicast Address as a IPv6 Destination Address is permissible and useful in certain cases.

1. Tunneling a BIERv6 packet over a non-BIER capable router.

2. Fast rerouting a BIERv6 packet using a unicast by-pass tunnel.

3. Forwarding a BIERv6 packet to one of the many BFR neighbors connected on a LAN.

4. Connecting BIER domains, for example Data Center domains, in an overlay manner.

The unicast address used in BIERv6 packet targeting a BFR SHOULD be the IPv6 BFR-Prefix advertised from this BFR. When a BFR advertises the BIER information with BIERv6 encapsulation capability, the IPv6 BFR-prefix of this BFR MUST be selected specifically for BIERv6 packet forwarding. Locally this "BIER Specific" IPv6 address is initialized in FIB with a flag of "BIER specific handling", represented as End.BIER function. For convenience, the indication in FIB share the same space as SRv6 Endpoints Behaviors defined in [I-D.ietf-spring-srv6-network-programming]. Apart from this sharing of code space, there is nothing dependent on SRv6. The co-existance of BIERv6 and SRv6 is outside the scope of this document.

BFR Prefix is used only in control plane in BIER MPLS encapsulation but not used in data plane. While in BIERv6, BFR prefix is used in both control plane and data plane. The "BIER Specific" IPv6 address can be used for BIER MPLS in control plane too. So it is RECOMMENDED to use a "BIER specific" IPv6 address as BFR prefix when deploying BIER in IPv6 network from the scratch. One should be careful not use the IPv6 address selected as BFR prefix for other purpose like BGP session until the "BIER specific handling" can do more general process.

The following is an example of configuring a BIER specific IPv6 address and using this address as BFR prefix:

```
 # Config a BIER specific IPv6 address with 128-bit mask on loopback0.
 interface loopback0
   ipv6 address 2019::AB37 128 End.BIER

 # Config the BIER-specific IPv6 address on loopback0 as BFR Prefix.
 bier sub-domain 6 ipv6-underlay
   bfr-prefix interface loopback0
```

The address used as "BIER specific" IPv6 address can be from inside the scope of an SRv6 Locator or outside the scope of the SRv6 Locator(s) since it is a host prefix (128-bit prefix-length prefix).

Each "BIER specific" address can be used in one or many sub-domains as BFR-prefix, such that it can be associated with one or many Multi-Topologies (MTs) or algorithms.

More than one "BIER specific" address are also allowed as different BFR-prefix of more than one sub-domain, as described in section 2 of [RFC8279].

The following is an example pseudo-code of the End.BIER function:

```
 1. IF NH = 60 AND HopLimit > 0                                  ;;Ref1
 2.   IF (OptType1 = BIER) and (OptLength1 = HdrExtLen*8 + 4) ;;Ref2
 3.     Lookup the BIER Header inside the BIER option TLV.
 4.     Forward via the matched entry.
 5.   ELSE
 6.     Drop the packet.
 7. ELSE IF Last_NH = ICMPv6                                     ;;Ref3
 8.   Send to CPU.
 9. ELSE
10.   Drop the packet.
```

Ref1: Destination options header follows the IPv6 header directly and HopLimit is bigger than zero.

Ref2: The first TLV is BIER type and is the only one in Destination options header.

Ref3: An ICMPv6 packet using End.BIER as destination address.

3.3. BIERv6 Packet Format

As a multicast packet enters the BIER domain in a Non-MPLS IPv6 network, the multicast packet will be encapsulated with BIERv6 Header.

Typically a BIERv6 header would contain the Destination Options Header as the only Extensions Header besides IPv6 Header. However, it is allowed and possible for other extension headers to appear along with the Destination Options Header as long as the requirements listed in section 3.1 of this document is met.

Format of the multicast packet with BIERv6 encapsulation carrying only the Destination Options header is depicted in the below figure.

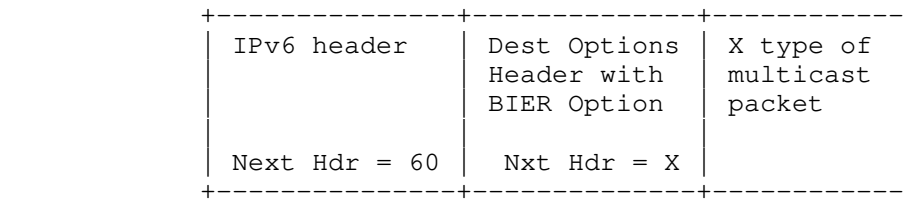

Format of the multicast packet with BIERv6 encapsulation carrying other extension headers along with Destination Options extension header is required to follow general recommendations of [RFC8200] and examples in other RFCs. [RFC6275] introduces how the order should be when other extension headers carries along with Home address option in a destination options header. Similar to this example, this

document requires the Destination Options Header carrying the BIER option MUST be placed as follows:

- After the routing header, if that header is present
- Before the Fragment Header, if that header is present
- Before the AH Header or ESP Header, if either one of those headers is present

Source Address field in the IPv6 header MUST be a routable IPv6 unicast address of the BFIR in any case.

BFIR encodes the Non-MPLS BIER header in the above mentioned encapsulation format and forwards the BIERv6 packet to the nexthop BFR following the local BIFT table.

BFRs in the IPv6 network, processes and replicates the packets towards the BFERs using the local BIFT table. The bit-string field in the Non-MPLS BIER header may be changed by the BFRs as they replicate the packet. BFRs MUST follow the procedures defined in section 3.1 as they modify the other fields in the Non-MPLS BIER header. The source address in the IPv6 header MUST NOT be modified by the BFRs.

4. BIERv6 Packet Processing

There is no BIER-specific processing, and all the 8 steps in section 6.5 of RFC8279 apply to BIERv6 packet processing. However, there are some IPv6-specific processing procedures due to the base and general procedures of IPv6.

On the overlay layer, when a multicast packet enters the BIER domain in a Non-MPLS IPv6 network, the Ingress BFR (BFIR) encapsulates the multicast packet with a BIERv6 Header, transforming it to a BIERv6 packet. The BIERv6 header includes an IPv6 header and IPv6 Destination Options Header within a standard Non-MPLS BIER header. Source Address field in the IPv6 header MUST be set to a routable IPv6 unicast address of the BFIR. Destination Address field in the IPv6 header is set to a BIER multicast address, FF0X::AB37, if the next-hop BFR is directly connected, or MAY be set to a unicast address in case of the scenarios discussed in section 3.2.

On the BIER layer, upon receiving an BIERv6 packet, the BFR processes the IPv6 header first. This is the general procedure of IPv6.

If the IPv6 Destination address is the BIER multicast address, a 'BIER Specific Handling' indication will be obtained by the preceding

Multicast DA lookup (MFIB lookup). The BIER option, if exists, will be checked to decide which neighbor(s) to replicate the BIERv6 packet to.

If the IPv6 Destination address is an IPv6 BFR-Prefix unicast address of this BFR, a 'BIER Specific Handling' indication will be obtained by the preceding Unicast DA lookup (FIB lookup). The BIER option, if exists, will be checked to decide which neighbor(s) to replicate the BIERv6 packet to.

It is a local behavior to handle the combination of extension headers, options and the BIER option(s) in destination options header when a 'BIER Specific Handling' indication is got by the preceding MFIB or FIB lookup. Early deployment of BIERv6 may require there is only one BIER option TLV in the destination options header followed the IPv6 header. How other extension headers or more BIER option TLVs in a BIERv6 packet is handled is outside the scope of this document.

A packet having a 'BIER Specific Handling' indication but not having a BIER option MAY be processed normally as normal multicast or unicast forwarding procedures do, or MAY be dropped.

A packet not having a 'BIER Specific Handling' indication but having a BIER option SHOULD be processed normally as normal multicast or unicast forwarding procedures, which may be a behavior of drop, or send to CPU, or other behaviors in existing implementations.

The Destination Address field in the IPv6 Header MUST change to the nexthop BFR's BFR Prefix if Unicast address is used in BIERv6.

The Hop Limit field of IPv6 header MUST decrease by 1 when sending packets to a BFR neighbor, while the TTL in the BIER header MUST be unchanged.

The BitString in the BIER header in the Destination Options Header may change when sending packets to a neighbor. Such change of BitString MUST be aligned with the procedure defined in RFC8279. Because of the requirement to change the content of the option when forwarding BIERv6 packet, the BIER option type should have chg flag 1 per section 4.2 of RFC8200.

The procedures applies normally if a bit corresponding to the self bfr-id is set in the bit-string field of the Non-MPLS BIER header of the BIERv6 packet. The node is considered to be an Egress BFR (BFER) in this case. The BFER removes the BIERv6 header, including the IPv6 header and the Destination Options header, and copies the packet to the multicast flow overlay. The egress VRF of a packet may be

determined by a single MFIB lookup on the BFER using both the IPv6 SA and IPv6 DA.

5. Security Considerations

A BIERv6 packet with a special IPv6 Destination Address, either multicast or unicast, would be processed by BIER forwarding procedure only when the 'BIER valid' flag has been obtained ahead of time in the normal MFIB or FIB lookup of the IPv6 header. Otherwise the packet with an IPv6 BIER Option will be dropped, as if the Option is not recognize by the node.

An IPv6 packet with BIER multicast address FF0X::AB37 as destination address, but does not carry IPv6 BIER Option will be dropped.

6. IANA Considerations

6.1. BIER Option Type

Allocation is expected from IANA for a BIER Option Type codepoint from the "Destination Options and Hop-by-Hop Options" sub-registry of the "Internet Protocol Version 6 (IPv6) Parameters" registry. The value 0x70 is suggested.

Hex Value	act	chg	rest	Description	Reference
0x70	01	1	10000	BIER Option	This draft

Figure 1: IPv6 Option Type Suggested

6.2. BIER Multicast Address

Allocation is expected from IANA for a BIER Multicast Address from the "Variable Scope Multicast Addresses" sub-registry of the "IPv6 Multicast Address Space Registry" registry. The address 'FF0X::AB37' is suggested.

Address(es)	Description	Reference
FF0X:0:0:0:0:0:0:AB37	ALL_BIER_FORWARDERS	This draft

Figure 2: Multicast Address Suggested

6.3. End.BIER Function

Allocation is expected from IANA for an End.BIER function codepoint from the "SRv6 Endpoint Behaviors" sub-registry. The value 60 is suggested.

Value	Hex	Endpoint function	Reference
TBD	TBD	End.BIER	This draft

Figure 3: End.BIER Function

7. Acknowledgements

The authors would like to thank Stig Venaas for his valuable comments. Thanks IJsbrand Wijnands, Greg Shepherd, Tony Przygienda, Toerless Eckert, Jeffrey Zhang for the helpful comments to improve this document.

8. References

8.1. Normative References

[RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <https://www.rfc-editor.org/info/rfc6275>.

[RFC7346] Droms, R., "IPv6 Multicast Address Scopes", RFC 7346, DOI 10.17487/RFC7346, August 2014, <https://www.rfc-editor.org/info/rfc7346>.

[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>.

8.2. Informative References

[I-D.ietf-bier-ipv6-requirements]
McBride, M., Xie, J., Dhanaraj, S., and R. Asati, "BIER IPv6 Requirements", draft-ietf-bier-ipv6-requirements-00 (work in progress), May 2019.

[I-D.ietf-spring-srv6-network-programming]
Filsfils, C., Camarillo, P., Leddy, J., daniel.voyer@bell.ca, d., Matsushima, S., and Z. Li, "SRv6 Network Programming", draft-ietf-spring-srv6-network-programming-00 (work in progress), April 2019.

[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>.

[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>.

Authors' Addresses

Jingrong Xie
Huawei Technologies

Email: xiejingrong@huawei.com

Liang Geng
China Mobile
Beijing 10053

Email: gengliang@chinamobile.com

Mike McBride
Futurewei

Email: mmcbride7@gmail.com

Rajiv Asati
Cisco

Email: rajiva@cisco.com

Senthil Dhanaraj
Huawei

Email: senthil.dhanaraj@huawei.com

Gang Yan
Huawei

Email: yangang@huawei.com

Yang Xia
Huawei

Email: yolanda.xia@huawei.com