

Network Working Group
Internet-Draft
Updates: [8296](#) (if approved)
Intended status: Standards Track
Expires: August 26, 2021

J. Xie
Huawei Technologies
L. Geng
China Mobile
M. McBride
Futurewei
R. Asati
Cisco
S. Dhanaraj
Huawei
Y. Zhu
China Telecom
Z. Qin
China Unicom
M. Shin
LG Uplus
G. Mishra
Verizon Inc.
X. Geng
Huawei
February 22, 2021

Encapsulation for BIER in Non-MPLS IPv6 Networks
draft-xie-bier-ipv6-encapsulation-10

Abstract

This document proposes a BIER IPv6 (BIERv6) encapsulation for Non-MPLS IPv6 Networks using the IPv6 Destination Option extension header. This document updates [RFC 8296](#).

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 August 26, 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	BIER IPv6 Encapsulation	4
3.1.	BIER Option in IPv6 Destination Options Header	4
3.2.	Destination Address in BIERv6 Encapsulation	6
3.3.	BIERv6 Packet Format	8
4.	BIERv6 Packet Processing	9
5.	Security Considerations	11
5.1.	Intra Domain Deployment	12
5.2.	ICMP Error Processing	13
5.3.	Security caused by BIER option	13
5.4.	Applicability of IPsec	14
6.	IANA Considerations	15
6.1.	BIER Option Type	15
6.2.	End.BIER Function	15
7.	Implementation Status	15
8.	Acknowledgements	16
9.	Contributors	17
10.	References	17
10.1.	Normative References	17
10.2.	Informative References	18
Appendix A.	Relationship to BIER Core Standards	19
Appendix B.	Extensions to BIER Control-plane Standards	20
Appendix C.	Considerations of Using Unicast Address	20

Authors' Addresses [21](#)

[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](#)].

The relationship of this document to BIER core standards is listed in [Appendix A](#).

The relevant extensions to BIER Control-plane Standards are listed in [Appendix B](#).

[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:

- o BIERv6 - Bit indexed explicit replication using IPv6 data plane.
- o BIERv6 Domain - A limited-domain using BIERv6 encapsulation as specified in this document for transporting customer multicast packets from one router to multiple destination routers. It is usually managed by a single administrative entity, e.g., a service-provider. It could be a single AS network or a large-scale network that includes multiple ASes. BIER Domain is also used for the same meaning as BIERv6 domain in this document.

- o BIERv6 Option - An Option type carried in IPv6 Destination Options Header (DO header, DOH) which includes the standard Non-MPLS BIER Header. It is in type-length-value (TLV) format. The value portion of the BIERv6 Option TLV, or the BIERv6 Option Data, is in the format of the standard Non-MPLS BIER header. BIER option is also used for the same meaning as BIERv6 option in this document.
- o BIERv6 Header - An IPv6 Header with BIER Option.
- o BIERv6 Packet - An IPv6 packet with BIERv6 Header. An IP/IPv6/Ethernet multicast packet is encapsulated with an outside BIERv6 header and transformed to a BIERv6 packet on the ingress PE (BFIR). BIERv6 packet is transported by the transit routers (BFRs) through a BIERv6 domain towards egress PEs (BFERs). BIERv6 packet is decapsulated by the BFERs, with the original IP/IPv6/Ethernet multicast packet being obtained and forwarded towards the multicast receivers .

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.

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

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	2	3	4	5	6	7	8	9
Next Header										Hdr Ext Len										Option Type										Option Length									
~ BIERv6 Option Data (Non-MPLS BIER Header defined in RFC8296) ~																																							

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.

BIERv6 Option Data The BIERv6 Option Data contains 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 reception. 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 a value larger than 0 upon encapsulation, and SHOULD decrease by 1 by a BFR when forwarding a BIERv6 packet to a BFR adjacency. If the incoming TTL is 0, the packet is considered to be "expired". See [Section 2.1.1.2 of RFC 8296](#).

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 BIERv6 encapsulation, uses Traffic Class field of IPv6 header instead.

Proto: SHOULD be set to 0 upon transmission and be ignored upon reception. In BIERv6 encapsulation, the functionality of this 6-bit Proto field is replaced by the Next Header field in Destination Options header or the last IPv6 extension header to indicate the type of the payload. This updates [section 2.1.2 of \[RFC8296\]](#) about Proto definition. Next Header value in BIERv6 encapsulation for common usage includes:

Value 4 for IPv4 packet as BIERv6 payload.

Value 41 for IPv6 packet as BIERv6 payload.

Value 143 for Ethernet packet as BIERv6 payload.

Multicast VPN (MVPN) service is considered as part of the BIER layering mode defined in [\[RFC8279\]](#), and should be supported by BIERv6 encapsulation. [\[I-D.xie-bier-ipv6-mvpn\]](#) illustrates how MVPN is supported in BIERv6 encapsulation without using this Proto field.

BIER-PING [\[I-D.ietf-bier-ping\]](#) is considered a useful function of the BIER architecture, and should be supported by BIERv6 encapsulation. How BIER-PING is supported in BIERv6 encapsulation without using this Proto field is outside the scope of this document.

BFIR-id: See [Section 2.1.2 of RFC 8296](#).

BitString: See [Section 2.1.2 of RFC 8296](#).

3.2. Destination Address in BIERv6 Encapsulation

When a BIERv6 packet is replicated to a next hop BFR, an unicast address of the next hop BFR is used as the destination address of the BIERv6 packet. Considerations of using unicast (or multicast) address is listed in [Appendix C](#).

The unicast address used in BIERv6 packet targeting a BFR SHOULD be advertised as part of the BIER IPv6 Encapsulation. When a BFR advertises the BIER information with BIERv6 encapsulation capability, an IPv6 unicast address 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.

If a BFR belongs to more than one sub-domain, it may (though it need not) have a different End.BIER in each sub-domain. If different End.BIER is used for each sub-domain, implementation SHOULD support

verifying the DA of a BIERv6 packet is the End.BIER address bound by the sub-domain of the packet.

For security deployment of BIERv6, the End.BIER address(es) is required to be allocated from an IPv6 address block, and the IPv6 address block is used for domain boundary security policy. See [section 5.1](#) of this document for such security policy. Such kind of security policy using IPv6 address block follows the paradigm settled by the [\[RFC8754\] section 5](#).

Deployment of BIERv6 in SRv6 network is allowed. In this case, the BIERv6 domain is the same as SRv6 domain, and the End.BIER address is allocated from the locator of SRv6.

To better understand the configuration mode of End.BIER address in BIERv6, [\[I-D.geng-bier-bierv6-yang\]](#) could be referenced.

For the convenience of such co-existence of BIERv6 and SRv6, the indication of End.BIER or "BIER specific handling" in FIB shares the same space as SRv6 Endpoints Behaviors defined in [\[I-D.ietf-spring-srv6-network-programming\]](#).

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                                                        ;;Ref3
6.     Drop the packet and end the process.
7. ELSE IF NH=ICMPv6 or (NH=60 and Dest_NH=ICMPv6)             ;;Ref4
8.   Send to CPU.
9. ELSE                                                          ;;Ref5
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 TLV present in Destination options header.

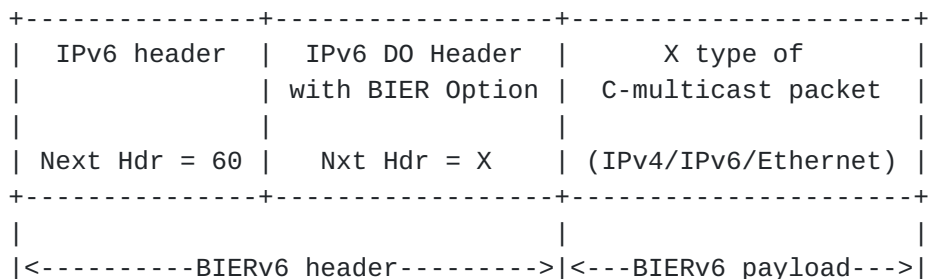
Ref3/Ref5: Undesired packet is dropped because the destination address is the BIER specific IPv6 address (End.BIER function).

Ref4: 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 by the Ingress BFR (BFIR).

Typically a BIERv6 header would contain the Destination Options Header as the only Extensions Header besides IPv6 Header, as 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:

- o After the routing header, if that header is present
- o Before the Fragment Header, if that header is present
- o 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 BIERv6 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 BitString field in the BIERv6 Option Data 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 BIERv6 Option Data. The source address in the IPv6 header MUST NOT be modified by the BFRs.

4. BIERv6 Packet Processing

When a multicast packet enters the BIER domain, 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 a BIERv6 Option in IPv6 Destination Options 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 the End.BIER address of the next-hop BFR the BIERv6 packet replicating to, no matter next-hop BFR is directly connected (one-hop) or not directly connected (multi-hop).

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 an End.BIER IPv6 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 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 is supposed to be a wrong packet or an ICMPv6 packet, and the process can be referred to the example in [section 3.2](#).

A packet not having a 'BIER Specific Handling' indication but having a BIER option SHOULD be processed normally as 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 End.BIER Unicast address 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 on a Non-BIER router, or decrease by 1 on a BFR.

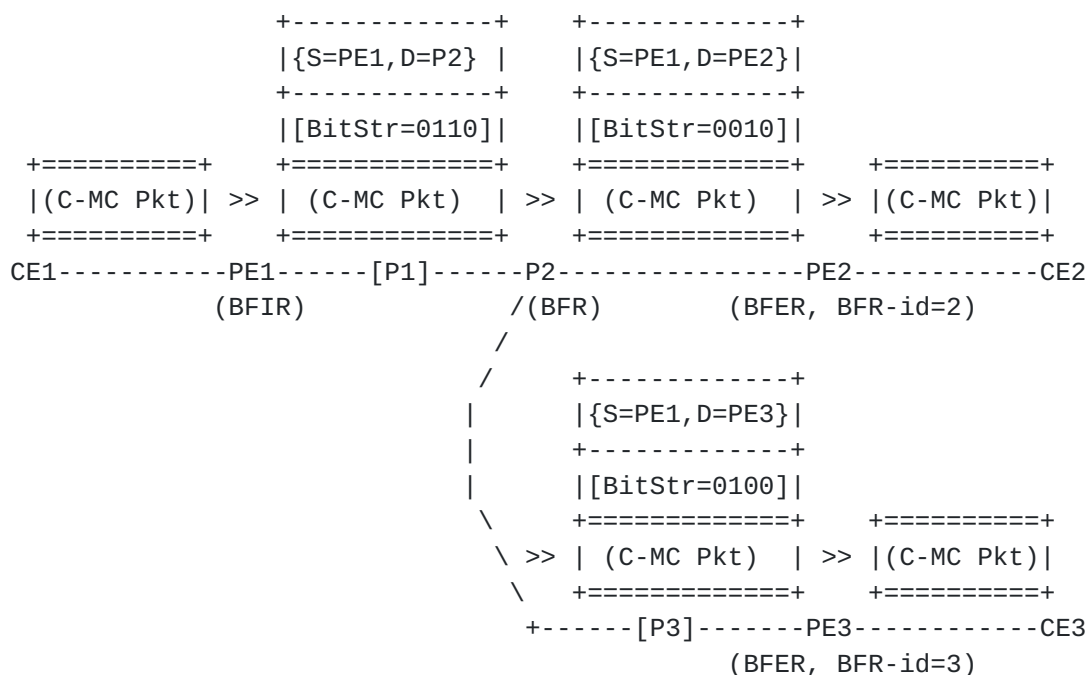
The BitString in the BIER header in the Destination Options Header may change when sending packets to a neighbor. Such change of BitString MUST be aligned with the procedure defined in [RFC8279](#).

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 BitString field of the BIERv6 Option Data 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 further lookup on the IPv6 source address instead of the upstream-assigned MPLS Label as described in [\[RFC8556\]](#).

The Fragment Header, AH Header or ESP Header, if exists after the BIER options header, can be processed on BFER only as part of the multicast flow overlay process.

The following diagram shows the whole progression of the multicast packet as it enters the BIERv6 domain on PE1, and leaves the BIERv6 domain on PE2 and PE3.



{S=PE1,D=PE2}: Source address and Destination address in IPv6 header.

[BitStr=0110]: BitString value in IPv6 DO Header.

(C-MC Pkt): Customer MultiCast packet.

- o PE1 is Provider Edge router, acting as BFIR.

- o P2 is Provider Core router, acting as BFR.
- o P1 and P3 are IPv6 routers, acting as Non-BFR.
- o PE2 and PE3 are Provider Edge routers, acting as BFER.
- o CE1 and CE2 are Customer Edge routers.

5. Security Considerations

BIER IPv6 encapsulation provides a new encapsulation based on IPv6 and BIER to transport multicast data packet in a BIER domain. The BIER domain can be a single IGP area, an anonymous system (AS) with multiple IGP areas, or multiple anonymous systems (ASes) operated by a network operator. A single BIER Sub-domain may be deployed through the whole BIER Domain, as illustrated in [\[I-D.geng-bier-ipv6-inter-domain\]](#).

This section reviews security considerations related to the BIER IPv6 encapsulation, based on security considerations of [\[RFC8279\]](#), [\[RFC8296\]](#), and other documents related to IPv6 extension.

It is expected that all nodes in a BIER IPv6 domain are managed by the same administrative entity. BIER-encapsulated packets should generally not be accepted from untrusted interfaces or tunnels. For example, an operator may wish to have a policy of accepting BIER-encapsulated packets only from interfaces to trusted routers, and not from customer-facing interfaces. See [section 5.1](#) for normal Intra domain deployment.

For applications that require a BFR to accept a BIER-encapsulated packet from an interface to a system that is not controlled by the network operator, the security considerations of [\[RFC8296\]](#) apply.

BIER IPv6 encapsulation may cause ICMP packet sent to BFIR and cause security problems. See [section 5.2](#) for ICMP related problems.

This document introduces a new option used in IPv6 Destination Options Header, together with the special-use IPv6 address called End.BIER in IPv6 destination address for BIER IPv6 forwarding. However the option newly introduced may be wrongly used with normal IPv6 destination address. See [section 5.3](#) for problems introduced by the new IPv6 option with normal IPv6 destination address.

If the multicast data packet of a BIERv6 packet is altered by an intermediate router, contents of the multicast data packet will be damaged. BIER IPv6 encapsulation provides the ability of IPsec to

ensure the confidentiality or integrity for multicast data packet. See [section 5.4](#) for this security problem.

If the BIERv6 encapsulation of a particular packet specifies a BitString (together with SI) other than the one intended by the BFIR, the packet is likely to be misdelivered. Some modifications of the BIER encapsulation, e.g., setting every bit in the BitString, may result in denial-of-service (DoS) attacks. This kind of DoS attack is a challenge not only in BIERv6 but also in BIER as specified in [\[RFC8279\]](#) and [\[RFC8296\]](#), as the BitString is required to change on BFR per the BIER forwarding procedures. This document does not provide new mechanisms to improve this kind of weakness.

A BIER router accepts and uses the End.BIER IPv6 address to construct BIFT only when the IPv6 address is configured explicitly, or received from a router via control-plane protocols. The received information is validated using existing authentication and security mechanisms of the control-plane protocols. BIER IPv6 encapsulation does not define any additional security mechanism in existing control-plane protocols, and it inherits any security considerations that apply to the control-plane protocols.

[5.1.](#) Intra Domain Deployment

Generally nodes outside the BIER Domain are not trusted: they cannot directly use the End.BIER of the domain. This is enforced by two levels of access control lists:

1. Any packet entering the BIER Domain and destined to an End.BIER IPv6 Address within the BIER Domain is dropped. This may be realized with the following logic. Other methods with equivalent outcome are considered compliant:

- * allocate all the End.BIER IPv6 Address from a block S/s

- * configure each external interface of each edge node of the domain with an inbound infrastructure access list (IACL) which drops any incoming packet with a destination address in S/s

- * Failure to implement this method of ingress filtering exposes the BIER Domain to BIER attacks as described and referenced in [\[RFC8296\]](#).

2. The distributed protection in #1 is complemented with per node protection, dropping packets to End.BIER IPv6 Address from source addresses outside the BIER Domain. This may be realized with the following logic. Other methods with equivalent outcome are considered compliant:

- * assign all interface addresses from prefix A/a
- * assign all the IPv6 addresses used as source address of BIER IPv6 packets from a block B/b
- * at node k, all End.BIER IPv6 addresses local to k are assigned from prefix Sk/sk
- * configure each internal interface of each BIER node k in the BIER Domain with an inbound IACL which drops any incoming packet with a destination address in Sk/sk if the source address is not in A/a or B/b.

For simplicity of deployment, a configuration of IACL effective for all interfaces can be provided by a router. Such IACL can be referred to as global IACL(GIACL). Each BIER node k then simply config a GIACL which drops any incoming packet with a destination address in Sk/sk if the source address is not in A/a or B/b for the intra-domain deployment mode.

5.2. ICMP Error Processing

The BIERv6 BFR does not send ICMP error messages to the source address of a BIERv6 packet, there is still chance that Non-BFR routers send ICMP error messages to source nodes within the BIER Domain.

A large number of ICMP may be elicited and sent to a BFIR router, in case when a BIERv6 packet is filled with wrong Hop Limit, either error or malfeasance. A rate-limiting of ICMP packet should be implemented on each BFR.

The ingress node can take note of the fact that it is getting, in response to BIER IPv6 packet, one or more ICMP error packets. By default, the reception of such a packets MUST be countered and logged. However, it is possible for such log entries to be "false positives" that generate a lot of "noise" in the log; therefore, implementations SHOULD have a knob to disable this logging.

5.3. Security caused by BIER option

This document introduces a new option used in IPv6 Destination Options Header. An IPv6 packet with a normal IPv6 address of a router (e.g. loopback IPv6 address of the router) as destination address will possibly carry a BIER option.

For a router incapable of BIERv6, such BIERv6 packet will not be processed by the procedure described in this document, but be

processed as normal IPv6 packet with unknown option, and the existing security considerations for handling IPv6 options apply. Possible way of handling IPv6 packets with BIER option may be send to CPU for slow path processing, with rate-limiting, or be discarded according to the local policy.

For a router capable of BIERv6, such BIERv6 packet MUST NOT be forwarded, but should be processed as a normal IPv6 packet with unknown option, or additionally and optionally be countered and logged if the router is capable of doing so.

5.4. Applicability of IPsec

IPsec [[RFC4301](#)] uses two protocols to provide traffic security services -- Authentication Header (AH) [[RFC4302](#)] and Encapsulating Security Payload (ESP) [[RFC4303](#)]. Each protocol supports two modes of use: transport mode and tunnel mode. IPsec support both unicast and multicast. IPsec implementations MUST support ESP and MAY support AH.

This document assume IPsec working in tunnel mode with inner IPv4 or IPv6 multicast packet encapsulated in outer BIERv6 header and IPsec header(s).

IPsec used with BIER IPv6 encapsulation to ensure that a BIER payload is not altered while in transit between BFIR and BFERs. If a BFR in between BFIR and BFERs is compromised, there is no way to prevent the compromised BFR from making illegitimate modifications to the BIER payload or to prevent it from misforwarding or misdelivering the BIER-encapsulated packet, but the BFERs will detect the illegitimate modifications to the BIER Payload (or the inner multicast data packet). This could provide cryptographic integrity protection for multicast data transport. This capability of IPsec comes from the design that, the destination options header carrying the BIER header is located before the AH or ESP and the BFR routers in between BFIR and BFERs can process the BIER header without aware of AH or ESP.

For ESP, the Integrity Check Value (ICV) is computed over the ESP header, Payload, and ESP trailer fields. It doesn't require the IP or extension header for ICV calculating, and thus the change of DA and BIER option data does not affect the function of ESP.

For AH, the Integrity Check Value (ICV) is computed over the IP or extension header fields before the AH header, the AH header, and the Payload. The IPv6 DA is immutable for unicast traffic in AH, and the change of DA in BIER IPv6 forwarding for multicast traffic is incompatible to this rule. How AH is extended to support multicast

traffic transporting through BIER IPv6 encapsulation is outside the scope of this document.

The detailed control-plane for BIER IPv6 encapsulation IPsec function is outside the scope of the document. Internet Key Exchange Protocol Version 2 (IKEv2) [[RFC7296](#)] and Group Security Association (GSA) [[RFC5374](#)] can be referred to for further studying.

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.

Hex Value	act	chg	rest	Description	Reference
TBD	01	1	TBD	BIER Option	This draft

6.2. 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

7. Implementation Status

Implementation of BIERv6 as described in this document and related drafts defined in [appendix B](#) has been operational.

In February 2021, China Unicom successfully validated this BIERv6 solution over its Beijing Metro network.

The BIERv6 test network contains the following network devices:

STB, ONT, OLT, BFER, non-BFR, BFR, BFIR, Switch, CR (Core Router, video flow input node)

where BIERv6 capable nodes include:

- o BFER: Huawei ATN980C
- o BFR&BFIR: Huawei NE40E

and IPv6 capable nodes include:

- o non-BFR: Cisco ASR 9006

BIERV6 test is composed of the following scenarios:

- o Intra-AS BIERV6: mechanisms defined in this document and [[I-D.xie-bier-ipv6-isis-extension](#)] are validated.
- o Inter-AS BIERV6: mechanisms defined in [[I-D.geng-bier-ipv6-inter-domain](#)] are validated.
- o IPv4 multicast traffic over BIERV6 MVPN. Mechanisms defined in [[I-D.xie-bier-ipv6-mvpn](#)] are validated.
- o Deployment of Intra-AS and Inter-AS BIERV6 with non-BIERV6 capable intermediate nodes, where these nodes only need to be IPv6 capable.
- o BIERV6 Ping
- o Deterministic convergence as indicated in [[RFC8279](#)], where intermediate link or BFR node fails.
- o Service reliability where BIERV6 source side link or BFIR node fails.
- o Service reliability where BIERV6 receiver side link or BFER node fails.

8. Acknowledgements

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

Thanks Aijun Wang for comments about BIER OAM function in BIER IPv6 encapsulation.

Thanks Mach Chen for review and suggestions about BIER-PING function in BIER IPv6 encapsulation.

9. Contributors

Gang Yan

Huawei Technologies

China

Email: yangang@huawei.com

Yang(Yolanda) Xia

Huawei Technologies

China

Email: yolanda.xia@huawei.com

10. References

10.1. 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>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC4302] Kent, S., "IP Authentication Header", [RFC 4302](#), DOI 10.17487/RFC4302, December 2005, <<https://www.rfc-editor.org/info/rfc4302>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), DOI 10.17487/RFC4303, December 2005, <<https://www.rfc-editor.org/info/rfc4303>>.
- [RFC5374] Weis, B., Gross, G., and D. Ignjatic, "Multicast Extensions to the Security Architecture for the Internet Protocol", [RFC 5374](#), DOI 10.17487/RFC5374, November 2008, <<https://www.rfc-editor.org/info/rfc5374>>.
- [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>>.

- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, [RFC 7296](#), DOI 10.17487/RFC7296, October 2014, <<https://www.rfc-editor.org/info/rfc7296>>.
- [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>>.
- [RFC8556] Rosen, E., Ed., Sivakumar, M., Przygienda, T., Aldrin, S., and A. Dolganow, "Multicast VPN Using Bit Index Explicit Replication (BIER)", [RFC 8556](#), DOI 10.17487/RFC8556, April 2019, <<https://www.rfc-editor.org/info/rfc8556>>.
- [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>>.

10.2. Informative References

- [I-D.geng-bier-bierv6-yang]
Geng, X., Qin, Z., and F. Zheng, "YANG Data Model for Bierv6", [draft-geng-bier-bierv6-yang-00](#) (work in progress), June 2020.

[I-D.geng-bier-ipv6-inter-domain]

Geng, L., Xie, J., McBride, M., Yan, G., and X. Geng, "Inter-Domain Multicast Deployment using BIERv6", [draft-geng-bier-ipv6-inter-domain-02](#) (work in progress), October 2020.

[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-ping]

Nainar, 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-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-28](#) (work in progress), December 2020.

[I-D.xie-bier-ipv6-isis-extension]

Xie, J., Wang, A., Yan, G., Dhanaraj, S., and X. Geng, "BIER IPv6 Encapsulation (BIERv6) Support via IS-IS", [draft-xie-bier-ipv6-isis-extension-02](#) (work in progress), October 2020.

[I-D.xie-bier-ipv6-mvpn]

Xie, J., McBride, M., Dhanaraj, S., Geng, L., and G. Mishra, "Use of BIER IPv6 Encapsulation (BIERv6) for Multicast VPN in IPv6 networks", [draft-xie-bier-ipv6-mvpn-03](#) (work in progress), October 2020.

[Appendix A](#). Relationship to BIER Core Standards

The BIER architecture [[RFC8279](#)] is inherited in this BIERv6 proposal, and the layering mode of BIER architecture is fully supported with some necessary extension to the data plane as well as the control plane standards.

The focus of this document is BIERv6 data plane, including the BIERv6 encapsulation and packet forwarding procedures. The common BIER header encoding [[RFC8296](#)] is maximum reused in this BIERv6 proposal.

To better understand the overall BIER IPv6 problem space and requirements, refer to [[I-D.ietf-bier-ipv6-requirements](#)].

Appendix B. Extensions to BIER Control-plane Standards

The relevant control-plane documents that have done or still to be done are listed below.

- o Based on [[RFC8401](#)], IS-IS extension is defined in [[I-D.xie-bier-ipv6-isis-extension](#)] for intra-AS BIERv6 information advertisement and BIRT/BIFT building.
- o OSPFv3 extension for intra-AS BIERv6 information advertisement and BIRT/BIFT building is to be defined.
- o Based on this BIERv6 encapsulation, [[I-D.geng-bier-ipv6-inter-domain](#)] illustrates how inter-AS BIRT/BIFT are built and how inter-AS multicast deployment is supported.
- o BGP extension for inter-AS BIERv6 information advertisement and BIRT/BIFT building is to be defined.
- o Based on [[RFC8556](#)], BGP-MVPN using BIERv6 encapsulation is defined in [[I-D.xie-bier-ipv6-mvpn](#)] for multicast service deployment.

Appendix C. Considerations of Using Unicast 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 way one may think of as an approach for both the two paradigms in BIERv6 encapsulation.

However using a unicast address has the following benefits:

1. Replicating 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 without imposing new requirements of snooping on switches.
4. Replicating a BIERv6 packet through an anonymous system(AS) to BFRs in other ASes, as illustrated in [[I-D.geng-bier-ipv6-inter-domain](#)].

Some of the above scenarios are assumed part of BIER architecture as described in [[RFC8279](#)], and some of them are the scalability aspects for inter-AS stateless multicast this document intends to support. This document intends to fulfil all these requirements (categorized

as multi-hop replication), and proposes to use unicast address for both one-hop replication and multi-hop replication.

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

Yongqing Zhu
China Telecom

Email: zhuyq8@chinatelecom.cn

Zhuangzhuang Qin
China Unicom

Email: qinzhuangzhuang@chinaunicom.cn

MooChang Shin
LG Uplus

Email: himzzang@lguplus.co.kr

Gyan Mishra
Verizon Inc.

Email: gyan.s.mishra@verizon.com

Xuesong Geng
Huawei

Email: gengxuesong@huawei.com