Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2020

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

March 9, 2020

Encapsulation for BIER in Non-MPLS IPv6 Networks
draft-xie-bier-ipv6-encapsulation-06

Abstract

This document proposes a BIER IPv6 (BIERv6) encapsulation for Non-MPLS IPv6 Networks using the IPv6 Destination Option extension header. This document updates [RFC8296].

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

Copyright Notice

Table of Contents

1. Introduction   2
2. Terminology   3
3. BIER IPv6 Encapsulation   3
  3.1. BIER Option in IPv6 Destination Options Header   3
  3.2. Multicast and Unicast Destination Address   6
  3.3. BIERv6 Packet Format   8
4. BIERv6 Packet Processing   9
5. Security Considerations 10
  5.1. Intra Domain Deployment 11
  5.2. ICMP Error Processing 12
  5.3. Security caused by BIER option 13
  5.4. Applicability of IPsec 13
6. IANA Considerations 14
  6.1. BIER Option Type 14
  6.2. End.BIER Function 15
  6.3. BIER Next Protocol Identifiers 15
7. Acknowledgements 15
8. Contributors 15
9. References 16
  9.1. Normative References 16
  9.2. Informative References 17
Authors' Addresses 17

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.

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:

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Next Header  |  Hdr Ext Len  |  Option Type  | Option Length |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
~          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.

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 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 IPv6 BIER encapsulation, uses highest 6-bit of Traffic Class field of IPv6 header to hold a Differentiated Services Codepoint [RFC2474].

Proto: This fileld is used for two functions. The first is to identify the BIER Payload following the BIER header, and the second is to deliver the packet to a proper overlay module by BIER layer, with VRF lookup in case of BIER data process, or without VRF lookup in case of BIER OAM process. In BIER IPv6 encapsulation, the first function of Proto is compromised by a proper Next Header value combination, and the second function of Proto should be kept with the semantic unique and consistent for BIER demultiplexing. This updates section 2.1.2 of [RFC8296] about Proto defination. This document support the following combination of BIER Proto and Next Header:

Use Next Header value 4, 41 and 143 for IPv4 packet, IPv6 packet and Ethernet packet respectively, and use Proto value TBD1 indicating "Delivering the data packet with VRF lookup in IPv6 source address".

Use Next Header value 59 for private packet format, and use Proto value 5 indicating "Delivering the BIER OAM packet without VRF lookup". The BIER-PING [I-D.ietf-bier-ping] works equally in BIER IPv6 encapsulation as well as in BIER MPLS or BIER Ethernet encapsulation.

Allocation of BIER Proto value TBD1 is listed in the IANA considerations section of this document.

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 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. 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.

Some of the above functions are assumed very basic requirements and part of BIER architecture as described in [RFC8279]. This document uses unicast address for both one-hop replication and multi-hop replication.

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. See section 5.2 for such use case.

The following is an example of configuring a sub-domain using BIER IPv6 encapsualation:

```
# Config an IPv6 block for End.BIER IPv6 address allocation
ipv6-block blk1 2001:DB8:A1:: 96 static 32

# Config BIER Sub-domain using End.BIER allocated from blk1
bier sub-domain 6 ipv6-underlay
    bfr-prefix interface loopback0
    end-bier ipv6-block blk1 opcode ::1
    encapsulation ipv6 bsl 256 max-si 0
```

The co-existance of BIERv6 and SRv6 is allowed. The following is an example of configuring a sub-domain using BIERv6 when SRv6 is already deployed with a locator 'loc1' configured:

```
# Config BIER Sub-domain using End.BIER allocated from loc1
bier sub-domain 6 ipv6-underlay
    bfr-prefix interface loopback0
    end-bier locator loc1 opcode ::1
    encapsulation ipv6 bsl 256 max-si 0
```

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]. Apart from this sharing of code space, there is nothing dependent on SRv6.

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 droped 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.

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.

```
+---------------+--------------+------------
| IPv6 header   | Dest Options | X type of
|               | Header with  | multicast
|               | BIER Option  | packet
|               |              |
| Next Hdr = 60 |  Nxt Hdr = X |
+---------------+--------------+------------
```

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 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 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).

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 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 refered 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 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 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.

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 a BIER packet is altered, either the BIER header, or the multicast data packet, by an intermediate router, packets may be lost, stolen, or otherwise misdelivered. BIER IPv6 encapsulation provides the ability of IPsec to ensure the confidentiality or integrity. See section 5.4 for this security problem.

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 configs 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

ICMP error packets generated within the BIER Domain are sent 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 BIER IPv6 packet is filled with wrong TTL, 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.

OAM functions of PING and TRACE for BIER IPv6 encapsulation may also need ICMP based on BIER IPv6 encapsulation and cause ICMP response message containing BIER option. The ability of seperating such OAM ICMP packets from error ICMP packets caused by error is necessary for the availability of OAM, otherwise the OAM function may fail due to the rate-limiting of ICMP error packets.

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) [RFC5996] 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. The value 0x70 is suggested.

Hex Value	act	chg	rest	Description	Reference
0x70	01	1	10000	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

6.3. BIER Next Protocol Identifiers

Allocation is expected from IANA for a BIER Proto codepoint from the "BIER Next Protocol Identifiers" sub-registry.

TBD1: Delivering the packet with VRF lookup in IPv6 source address

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.

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.

8. Contributors

Gang Yan

Huawei Technologies

China

Email: yangang@huawei.com

Yang(Yolanda) Xia

Huawei Technologies

China

Email: yolanda.xia@huawei.com

9. References

9.1. Normative References

[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>.

[RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, DOI 10.17487/RFC5996, September 2010, <https://www.rfc-editor.org/info/rfc5996>.

[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>.

[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>.

[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>.

9.2. Informative 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-ipv6-requirements] McBride, M., Xie, J., Dhanaraj, S., Asati, R., and Y. Zhu, "BIER IPv6 Requirements", draft-ietf-bier-ipv6-requirements-04 (work in progress), January 2020.

[I-D.ietf-bier-ping] Kumar, N., Pignataro, C., Akiya, N., Zheng, L., Chen, M., and G. Mirsky, "BIER Ping and Trace", draft-ietf-bier-ping-06 (work in progress), October 2019.

[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-10 (work in progress), February 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>.

[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