

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 3, 2019

J. Xie
Huawei Technologies
L. Geng
L. Wang
China Mobile
G. Yan
M. McBride
Y. Xia
Huawei
July 2, 2018

Encapsulation for BIER in Non-MPLS IPv6 Networks
draft-xie-bier-6man-encapsulation-01

Abstract

Bit Index Explicit Replication (BIER) introduces a new multicast-specific BIER Header. Currently BIER has two types of encapsulation formats: one is MPLS encapsulation, the other is Ethernet encapsulation. This document proposes a BIER IPv6 encapsulation for Non-MPLS IPv6 Networks using an 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].

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 3, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 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	2
2. Terminology	3
3. Problem Statement and Requirements	3
3.1. Problem Statement	3
3.2. Requirements	4
4. IPv6 BIER Encapsulation	4
4.1. Considerations	4
4.2. IPv6 BIER Destination Option	4
4.3. The whole IPv6 header for BIER packets	5
5. BIER Forwarding in Non-MPLS IPv6 Networks	7
6. Security Considerations	8
7. IANA Considerations	8
8. Acknowledgements	8
9. References	8
9.1. Normative References	8
9.2. Informative References	9
Authors' Addresses	9

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 two types of BIER encapsulation formats: one is MPLS encapsulation, the other is non-MPLS encapsulation. The Non-MPLS encapsulation defined in [RFC8296] is in fact an Ethernet encapsulation with an ethertype 0xAB37, and an 'Ethernet encapsulation' will be used to refer to such an encapsulation in the following text. This document proposes a BIER IPv6 encapsulation for Non-MPLS IPv6 Networks using an IPv6 Destination Option extension header.

2. Terminology

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

3. Problem Statement and Requirements

3.1. Problem Statement

MPLS is a very popular and successful encapsulation. One of the benefits of MPLS is its ability to easily stack a label onto another, thus forming a label stack. This same label stacking benefit is also available for BIER by using an MPLS encapsulation. For example, an MPLS-encapsulated BIER packet can easily run over an MPLS tunnel, either a legacy RSVP-TE/LDP LSP, or an MPLS Segment Routing tunnel. Such a mechanism is the key to obtain the capability of "fast reroute" or "bypass a Non-capable router". To quote [RFC8279]:

- o In the event that unicast traffic to the BFR-NBR is being sent via a "bypass tunnel" of some sort, the BIER-encapsulated multicast traffic sent to the BFR-NBR SHOULD also be sent via that tunnel. This allows any existing "fast reroute" schemes to be applied to multicast traffic as well as to unicast traffic.
- o Unicast tunnels are used to bypass non-BFRs.

Some other scenarios also need BIER to run on a tunnel, such as transferring a BIER packet through a whole Non-BIER network or domain.

The capability to run BIER on a tunnel, especially the widely deployed mpls tunnel, can be obtained by using a BIER MPLS encapsulation, but cannot be obtained by using a BIER Ethernet encapsulation. It is not possible either, to run BIER on other links such as POS, by using BIER Ethernet encapsulation.

The capability of running BIER on various kinds of links and tunnels, by using an MPLS encapsulation, is beneficial to BIER deployments. In an IPv6 network, however, there are considerations of using a non-MPLS encapsulation for unicast as the data-plane, such as SRH defined in [I-D.ietf-6man-segment-routing-header], where the function of a bypass tunnel uses an SRH header, with one or many Segments (or SIDs), instead of MPLS Labels.

3.2. Requirements

This chapter lists the BIER IPv6 encapsulation requirements needed to make the deployment of BIER on IPv6 network with SRH data-plane the same as on IPv4/IPv6 network with MPLS data-plane. These BIER IPv6 encapsulation requirements should provide similar benefits to MPLS encapsulation such as "fast reroute" or "run on any link or interface".

1. The listed requirements MUST be supported with any L1/L2 over which BIER layer can be realized.
2. It SHOULD support a hop-by-hop replication to multiple destinations in a BIER Domain.
3. It SHOULD support BIER on an "SRH tunnel".
4. It SHOULD align with the recommendations of the 6MAN working group.

4. IPv6 BIER Encapsulation

4.1. Considerations

BIER is generally a hop-by-hop and one-to-many architecture, while Segment Routing is a source-routing and one-to-one architecture. One of the challenges of an BIER IPv6 Encapsulation is how to allow BIER to run over a Segment Routing tunnel. A suitable method for such a combination is to use a Multicast Address as the Last Segment (or SID). After all the source-routing hops have been processed, the remaining Multicast Address becomes the IPv6 Destination Address. A hop-by-hop replicating diagram begins by using the Destination Multicast Address.

We then need to decide where to place the BIER header. According to [RFC8200], [RFC6564], and [RFC7045], a suitable place for a well-known BIER header is an IPv6 Destination Option extension header. Such a Destination Option carrying BIER header is only used for a hop-by-hop Multicast Address destination, but not for the transit router along the source-routing path.

4.2. IPv6 BIER Destination Option

The IPv6 BIER Destination Option is carried by the IPv6 Destination Option Header (indicated by a Next Header value 60). It is used in a packet sent by an IPv6 BFIR router to inform the routers in an IPv6 BIER domain to replicate to destination BFER routers.

The IPv6 BIER Destination Option is encoded in type-length-value (TLV) format as follows:

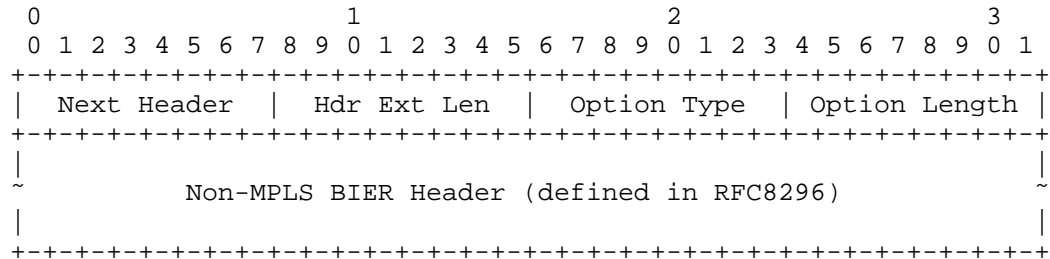


Figure 1: IPv6 BIER Destination Option

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 TBD. Need to be allocated by IANA.

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, including the BIFT-id.

4.3. The whole IPv6 header for BIER packets

[RFC8200] specifies that the Destination Option Header can be located either before the Routing Header or after the Routing Header. However, this document requires that the Destination Option Header with a BIER Destination Option TLV is always located after the Routing Header if the Routing Header is present.

This is because the BIER header is always handled after the tunnels (or bypass tunnels) have been handled. BIER MPLS encapsulation has the same behavior. To quote [RFC8296]:

- o It is crucial to understand that in an MPLS network the first four octets of the BIER encapsulation header are also the last four octets of the MPLS header. Therefore, any prior MPLS label stack entries MUST have the S bit (see [RFC3032]) clear (i.e., the S bit must be 0).

Other IPv6 extension headers are not commonly used in the current Internet. For Example, [RFC6744] says that "IPv6 Destination Options headers, and the options carried by such headers, are extremely uncommon in the deployed Internet". [RFC6564] says that "Extension headers, with the exception of the Hop-by-Hop Options header, are not usually processed on intermediate nodes", and that "Reports from the field indicate that some IP routers deployed within the global Internet are configured either to ignore the presence of headers with hop-by-hop behavior or to drop packets containing headers with hop-by-hop behavior."

Such IPv6 extension headers will even be more uncommon when a BIER encapsulation is used in data-plane forwarding. The entire IPv6 header, with BIER encapsulation and Routing Header, is expected to look like this:

IPv6 header

Hop-by-Hop Options header [Not Used]

Destination Options header [Not Used]

Routing header [SRH Header with Multicast Address as last SID]

Fragment header [Not Used]

Authentication header [Not Used]

Encapsulating Security Payload header [Not Used]

Destination Options header [BIER header in BIER Option TLV]

Upper-layer header [Data-plane Data]

Once a packet is encapsulated with a BIER Destination Option, it is basically assumed to be a data-plane multicast packet, so the 'OAM' or similar functions in the SRH Header Optional TLV Objects field should not exist.

The last Segment (SID) in the SRH header, or Segment List[0], should be a Multicast Address to indicate a hop-by-hop behavior. Such a Multicast Address can be reserved or unreserved as the Destination Option Header can inform the routers to do the address check. A reserved multicast address should be indicating a 'BIER specific' address.

BIER header has a 'proto' field to identify the type of BIER packet payload, and the IANA has created a registry called "BIER Next

Protocol Identifiers" to assign the value. That means the 'Upper-layer header' of a BIER packet have already been identified by the 'proto' field of the BIER header in the Destination Option Header. Thus the 'Next Header' in the Destination Option Header is not need to identify the 'Upper-layer header' any more, and is recommended to be set to 'No Next Header (value 59)'.

5. BIER Forwarding in Non-MPLS IPv6 Networks

In a Non-MPLS IPv6 Network, BIER may be deployed in a hop-by-hop manner, or possibly be deployed through an SRH tunnel either for "bypassing Non-capable BIER routers" or "fast rerouting". Here is an example where a packet is firstly forwarded through an SRH tunnel and then through a hop-by-hop BIER domain.

When a router along the Segment Routing path receives an IPv6 BIER packet with an SRH header, and if the IPv6 destination address is not one of the router's address, then the packet is forwarded by an IPv6 FIB lookup of the destination address and none of the IPv6 extension headers will be checked. If the IPv6 Destination Address is one of the router's address, and also one of the router's Segment (or SID) of some type, then the router will do a specific function indicated by the Segment, as defined in

[I-D.filsfils-spring-srv6-network-programming]. If the IPv6 Destination Address is a specific type of Segment, called BIER Segment or BIER SID, then the according function is called Endpoint BIER function or 'End.BF' function for short.

When router receives a packet destined to X and X is a local End.BF SID, the router does:

1. IF SL > 0
2. decrement SL
3. update IPv6 DA with SRH[SL]
4. IF SL = 0 & STATE(SRH[0]) = BIER
5. update IPv6 header NH with SRH NH
6. pop the SRH
7. forward the updated packet
8. ELSE
9. drop the packet
10. ELSE
11. drop the packet

Figure 2: End.BF Function

The End.BF function is used for the SRH tunnel destination router to terminate the source-routing SRH forwarding and begin the hop-by-hop BIER IPv6 forwarding. After the SRH header is popped, the multicast

address in the updated IPv6 Destination Address indicates the BIER information of this 'host', and the packet will be forwarded according to the BIER Header in the BIER Destination Option TLV in the IPv6 Destination Option extension header of this 'host'.

In the following hop-by-hop forwarding procedure, the IPv6 Destination Address in an incoming packet indicates the BIER information of this 'host', and the packet will be forwarded according to the BIER Header in the BIER Destination Option TLV in the IPv6 Destination Option extension header. A router is required to ignore the IPv6 BIER Destination Option if the IPv6 Destination Address of a packet is not a multicast address, or is a multicast address without indicating the BIER information of this 'host'.

6. Security Considerations

An IPv6 BIER Destination Option with Multicast Address Destination would be used only when an IPv6 BIER state with the specific Multicast Address Destination has been built by the control-plane. Otherwise the packet with an IPv6 BIER Destination Option will be discarded.

7. IANA Considerations

Allocation is expected from IANA for a Destination Option Type codepoint from the "Destination Options and Hop-by-Hop Options" sub-registry of the "Internet Protocol Version 6 (IPv6) Parameters" registry [RFC2780] at <<https://www.iana.org/assignments/ipv6-parameters/>>.

8. Acknowledgements

TBD.

9. References

9.1. Normative References

[I-D.filsfils-spring-srv6-network-programming]
 Filsfils, C., Li, Z., Leddy, J., daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d., Steinberg, D., Raszuk, R., Matsushima, S., Lebrun, D., Decraene, B., Peirens, B., Salsano, S., Naik, G., Elmalky, H., Jonnalagadda, P., and M. Sharif, "SRv6 Network Programming", draft-filsfils-spring-srv6-network-programming-04 (work in progress), March 2018.

- [I-D.ietf-6man-segment-routing-header]
Previdi, S., Filsfils, C., Leddy, J., Matsushima, S., and
d. daniel.voyer@bell.ca, "IPv6 Segment Routing Header
(SRH)", draft-ietf-6man-segment-routing-header-13 (work in
progress), May 2018.
- [RFC6564] Krishnan, S., Woodyatt, J., Kline, E., Hoagland, J., and
M. Bhatia, "A Uniform Format for IPv6 Extension Headers",
RFC 6564, DOI 10.17487/RFC6564, April 2012,
<<https://www.rfc-editor.org/info/rfc6564>>.
- [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing
of IPv6 Extension Headers", RFC 7045,
DOI 10.17487/RFC7045, December 2013,
<<https://www.rfc-editor.org/info/rfc7045>>.
- [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>>.

9.2. Informative 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>>.

Authors' Addresses

Jingrong Xie
Huawei Technologies

Email: xiejingrong@huawei.com

Liang Geng
China Mobile
Beijing 10053

Email: gengliang@chinamobile.com

Lei Wang
China Mobile
Beijing 10053

Email: wangleiyjy@chinamobile.com

Gang Yan
Huawei

Email: yangang@huawei.com

Mike McBride
Huawei

Email: mmcbride7@gmail.com

Yang Xia
Huawei

Email: yolanda.xia@huawei.com