

BIER WG
Internet-Draft
Intended status: Standards Track
Expires: September 16, 2016

C. Wang
Z. Zhang
ZTE Corporation
A. Qu
March 15, 2016

BIER Ethernet
draft-wang-bier-ethernet-01

Abstract

Bit Index Explicit Replication (BIER) [I-D.ietf-bier-architecture] is an architecture that provides optimal multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain any multicast related per-flow state. BIER also does not require any explicit tree-building protocol for its operation. When a multicast data packet enters the BIER domain, the BFIR determines the BFERs to which the packet needs to be sent. Then the BFIR encapsulates the packet in a BIER header and forwards the packet according to the BIFTs. Currently, there is a BIER-MPLS solution to transmit multicast traffic using MPLS label indication. Alternatively, this document tries to propose a solution named BIER Ethernet to support BIER forwarding in Ethernet network.

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 <http://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 September 16, 2016.

Copyright Notice

Copyright (c) 2016 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
(<http://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. Convention and Terminology	4
3. BIER Ethernet Header	5
4. Imposing and processing the BIER Ethernet header	8
5. Control Plane Considerations about BIER Ethernet	9
6. BIER Ethernet Considerations	11
6.1. BIER Ethernet for Traffic Engineering	11
6.2. BIER Ethernet for Multicast VPN	11
7. Assignment Considerations	12
7.1. IEEE Registration Authority Considerations	12
7.2. IANA Considerations	12
8. Acknowledgements	13
9. References	14
9.1. Normative References	14
9.2. Informative References	14
Authors' Addresses	16

1. Introduction

Bit Index Explicit Replication (BIER) [I-D.ietf-bier-architecture] is an architecture that provides optimal multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain any multicast related per-flow state. BIER also does not require any explicit tree-building protocol for its operation. A multicast data packet enters a BIER domain at a "Bit-Forwarding Ingress Router" (BFIR), and leaves the BIER domain at one or more "Bit-Forwarding Egress Routers" (BFERs). The BFIR encapsulates a BIER header to the packet. The BIER header contains a BitString in which each bit represents exactly one BFER to forward the packet to.

Specifically, after encapsulating a BIER header to the original multicast data packet, the BFIR and the intermediate BFRs as well as the BFERs require to carry out the BIER forwarding procedures to the BIER-encapsulated packet according to the information in the BIER header. As described in [I-D.ietf-bier-architecture], each BFR firstly should determine the packet's Subdomain-ID, BitStringLength and Set ID information to locate the exact "Bit Index Forwarding Table" (BIFT), and then do the subsequent procedures in terms of BitString and the found BIFT.

However the existing draft requires MPLS label preceding the BIER header and using Bottom label to carry BIER forwarding information is not clean cut design.

Hence we should design BIER header that holds all BIER forwarding related information, and just let MPLS as an independent layer protocol to help BIER forwarding as it does for IPv4/IPv6/IPmcast traffic.

Additionally, the BIER forwarding capability will be also introduced in enterprise/data center, such feature may be newly implemented in switch ASICs, with clean cut design using BIER-ethernet draft, the implementation will be more clean as well.

So this document tries to propose this kind of BIER header which contains significant BIER information directly such as Subdomain-ID, BitStringLength and Set ID as well as BitString. It is applicable when a given BIER domain is an Ethernet network.

2. Convention and Terminology

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].

The terms about BIER are defined in [I-D.ietf-bier-architecture].

3. BIER Ethernet Header

The BIER Ethernet header is shown in Figure 1.

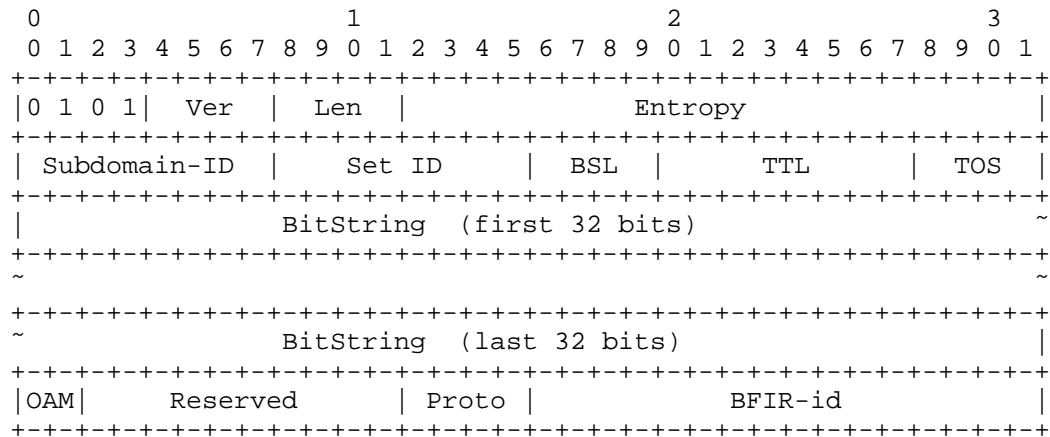


Figure 1: BIER Ethernet Header

First nibble: The first 4 bits of the header are set to 0101; this ensures that the BIER header will not be confused with an IP header or with the header of a pseudowire packet.

Ver: this 4-bit field identifies the version of the BIER header.

Len: This 4-bit field encodes the length of BIER Ethernet header.

Entropy: This 20-bit field specifies an "entropy" value that can be used for load balancing purposes. The BIER forwarding process may do equal cost load balancing, but the load balancing procedure **MUST** choose the same path for any two packets have the same entropy value.

Subdomain-ID: unique value identifying the BIER Subdomain within the BIER domain, as described in section 1 of [I-D.ietf-bier-architecture].

Set ID: indicates the packet's Set Identifier.

BSL: indicates the packet's BitStringLength.

TTL: Time to Live

TOS: Type of Service. It can be used to differentiate services to

different BIER packets.

BitString: together with the packet's Set ID, identifies the destination BFERs for this packet.

OAM: These two bits are used for the passive performance measurement marking method.

Reserved: These 10 bits are currently unused. They SHOULD be set to zero upon transmission, and MUST be ignored upon reception.

Proto: This 4-bit field identifies the type of the payload. (The "payload" is the packet or frame immediately following the BIER header.)

BFIR-id: By default, this is the BFR-id of the BFIR, in the Subdomain to which the packet has been assigned. The BFR-id is encoded in the 16-bit field as an unsigned integer in the range [1,65535].

Furthermore, BIER Ethernet encapsulated packet has the following format. The original multicast data packet is encapsulated with two headers (starting from the outermost header): Outer Ethernet Header + BIER Ethernet Header. Figure 2 is an example of an outer Ethernet Header. The outer VLAN tag is optional. In some situations, there may be some other encapsulation headers before the multicast data packet and after the BIER Ethernet header.

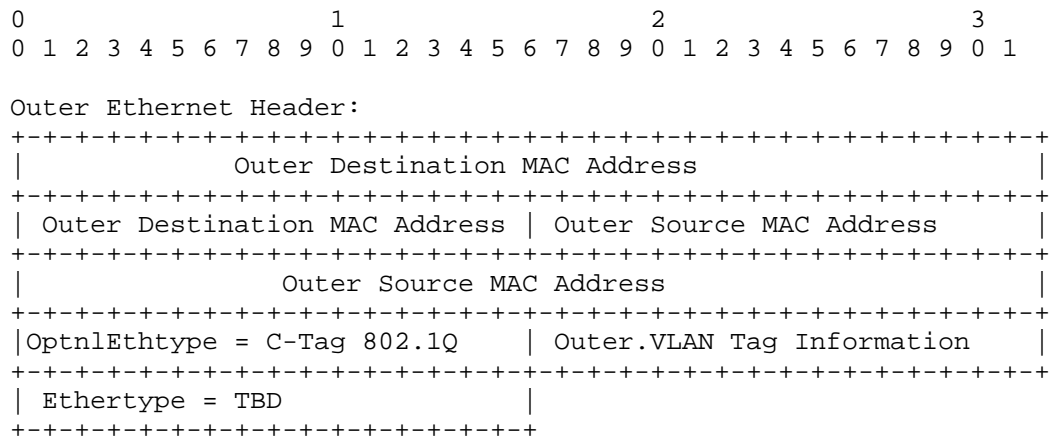


Figure 2: Outer Ethernet Header

Ethertype: requires a new Ethertype for BIER Ethernet header.

4. Imposing and processing the BIER Ethernet header

When a BFIR receives a multicast packet from outside the BIER domain, the BFIR carries out the following procedure:

1. By consulting the "multicast flow overlay", the BFIR determines the set of BFERs that must receive the packet.

2. By consulting the "BIER layer", the BFIR determines the packet's Subdomain, BitStringLength, Set Identifier and BitString information. The BFERs have the same Set Identifier can be encoded into the same BitString.

3. Using information provided by the routing underlay associated with the packet's BIER information, the BFIR determines the next hop for each (Set Identifier, the BitString) combination, and copies packet to each Set Identifier.

4. Before transmitting the packet to the next hop, the BFIR updates the BitString information and encapsulates the BIER Ethernet header to the multicast packet.

When an intermediate BFR receives a BIER Ethernet encapsulated packet, it acquires Subdomain-ID, BitStringLength as well as Set Identifier information directly from the BIER Ethernet header to determine the BIFT, and then forwards the received BIER packet according to the procedures described in [I-D.ietf-bier-architecture].

When a BFR receives a BIER Ethernet encapsulated packet whose Subdomain ID, Set Identifier and BitString identify the BFR itself, then the BFR is also a BFER for that packet. As a BFER, it must decapsulate the BIER Ethernet header, and pass the original multicast packet out.

5. Control Plane Considerations about BIER Ethernet

As described in the BIER OSPF extensions [I-D.ietf-bier-ospf-bier-extensions] and BIER ISIS extensions [I-D.ietf-bier-isis-extensions], they already define BIER Info Sub-TLV as the following format in Figure 3 (take the ospf extensions for example).

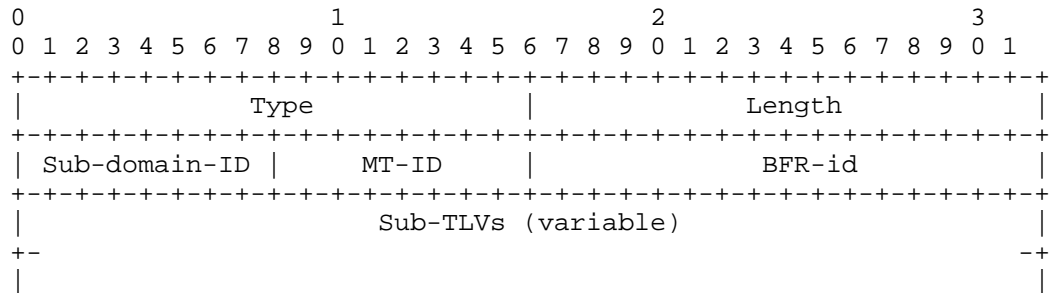


Figure 3: BIER Info Sub-TLV

To implement BIER Ethernet, the control plane, except the advertisements of BIER Info Sub-TLV, should have advertisements about BitStringLengths information the sending BFR supports. A reference format of BSL Sub-sub-TLV is illustrated in Figure 4 (take the ospf extensions for example as well).




Figure 4: BSL Sub-sub-TLV

Type: value of 1 indicating BSL Sub-sub-TLV for BIER Ethernet encapsulation.

Len: This 8-bit field encodes the length of this sub-sub-TLV.

BSL Identifier: indicating the BSL the sending BFR supporting.

The sending BFR may support one or several BSLs, as following:

00000001: represents BSL 64 bits;

00000010: represents BSL 128 bits;

00000100: represents BSL 256 bits;

00001000: represents BSL 512 bits;

00010000: represents BSL 1024 bits;

00100000: represents BSL 2048 bits;

01000000: represents BSL 4096 bits;

Each bit represents one BSL. When there are two or more bits set, that means the sending BFR supports more than one BSL. For example, if the BSL Identifier is 00010101, it means the sending BFR supports 1024 bits, 256 bits and 64 bits.

Additionally, the similar BSL Sub-sub-TLV extension can be introduced in ISIS extension and IDR extension for implementing BIER Ethernet.

6. BIER Ethernet Considerations

6.1. BIER Ethernet for Traffic Engineering

Specifically, BIER-TE encapsulation format may be the same as BIER encapsulation. However, how to interpret the BitString is totally different. Hence, BIER-Ethernet encapsulation MUST need one identifier to be assigned to identify the BIER header is for BIER forwarding or BIER-TE forwarding. For example, one bit in Reserved field can be reserved for this purpose.

6.2. BIER Ethernet for Multicast VPN

In MVPN, the P-tunnels are used for carrying multicast traffic across backbone. BIER tunnel Type is newly defined in [I-D.ietf-bier-mvpn]. The BIER Encapsulation used for multicast tunnel is independent of the (upstream assigned) MPLS label. Hence, BIER-Ethernet can also be used as P-Tunnel. In other words, there may need a new Tunnel Type to identify BIER-Ethernet Tunnel type, or a new flag to distinguish BIER-MPLS tunnel and BIER-Ethernet Tunnel.

7. Assignment Considerations

7.1. IEEE Registration Authority Considerations

This document requests the IEEE Registration Authority to assign a new Ethertype for BIER Ethernet Header.

7.2. IANA Considerations

This document requires new IANA allocation for BSL Sub-sub-TLV extension in different routing protocol.

8. Acknowledgements

The authors would like to thank IJsbrand Wijnands, Tony Przygienda and Andrew Qu for their ideas and contribution to this document.

9. References

9.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007, <<http://www.rfc-editor.org/info/rfc4915>>.

9.2. Informative References

- [I-D.ietf-bier-architecture]
Wijnands, I., Rosen, E., Dolganow, A., P, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", draft-ietf-bier-architecture-03 (work in progress), January 2016.
- [I-D.ietf-bier-isis-extensions]
Ginsberg, L., P, T., Aldrin, S., and J. Zhang, "BIER support via ISIS", draft-ietf-bier-isis-extensions-01 (work in progress), October 2015.
- [I-D.ietf-bier-mpls-encapsulation]
Wijnands, I., Rosen, E., Dolganow, A., Tantsura, J., and S. Aldrin, "Encapsulation for Bit Index Explicit Replication in MPLS Networks", draft-ietf-bier-mpls-encapsulation-03 (work in progress), February 2016.
- [I-D.ietf-bier-mvpn]
Rosen, E., Sivakumar, M., Aldrin, S., Dolganow, A., and T. P, "Multicast VPN Using BIER", draft-ietf-bier-mvpn-02 (work in progress), December 2015.
- [I-D.ietf-bier-ospf-bier-extensions]
Psenak, P., Kumar, N., Wijnands, I., Dolganow, A., P, T., Zhang, J., and S. Aldrin, "OSPF Extensions For BIER", draft-ietf-bier-ospf-bier-extensions-01 (work in progress), October 2015.
- [I-D.ietf-ospf-prefix-link-attr]
Psenak, P., Gredler, H., rjs@rob.sh, r., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute

Advertisement", draft-ietf-ospf-prefix-link-attr-13 (work in progress), August 2015.

Authors' Addresses

Cui(Linda) Wang
ZTE Corporation
No.50 Software Avenue, Yuhuatai District
Nanjing
China

Email: wang.cuil@zte.com.cn

Zheng(Sandy) Zhang
ZTE Corporation
No.50 Software Avenue, Yuhuatai District
Nanjing
China

Email: zhang.zheng@zte.com.cn

Andrew Qu
2860 Junction Ave
San JoseGBP[not] CA 95134

Email: laodulaodu@gmail.com

