

BIER WG
Internet-Draft
Intended status: Standards Track
Expires: January 19, 2018

Quan. Xiong
Fangwei. Hu
Zheng. Zhang
ZTE Corporation
July 18, 2017

BIER-TE Encapsulation
draft-xiong-bier-te-encapsulation-00.txt

Abstract

Traffic Engineering for Bit Index Explicit Replication (BIER-TE) shares part of architecture, definition and packet format with Bit Index Explicit Replication (BIER) according to the introduce in [[I-D.eckert-bier-te-arch](#)]. BIER-TE supports the traffic engineering by explicit hop-by-hop forwarding and loose hop forwarding of packets. [[I-D.ietf-bier-mpls-encapsulation](#)] specifies a BIER encapsulation that BIER header contains a bitstring in which each bit represents exactly one egress router in the domain.

This document proposes a set of extensions to BIER encapsulation for BIER-TE. The extensions define the BIER-TE header which contains several bitstrings and each bit in each bitstring represents one or more adjacencies in BIER-TE domain. The encapsulation can be used both in an MPLS network and a non-MPLS 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 January 19, 2018.

Copyright Notice

Copyright (c) 2017 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	2
2.	Requirements Language	3
2.1.	Terminology	3
3.	BIER-TE Encapsulation	3
3.1.	BitString Sub-TLV	4
4.	Processing Rules with the BIER-TE Encapsulation	5
5.	Security Considerations	6
6.	IANA Considerations	6
7.	Acknowledgements	6
8.	Normative References	6
	Authors' Addresses	7

[1.](#) Introduction

Traffic Engineering for Bit Index Explicit Replication (BIER-TE) shares part of architecture, definition and packet format with Bit Index Explicit Replication (BIER) according to the introductions in [[I-D.eckert-bier-te-arch](#)]. But BIER-TE supports the traffic engineering by explicit hop-by-hop forwarding and loose hop forwarding of packets. The BIER-TE controller host determines and assigns the BitPositions to the adjacencies which explicit paths passing through.

[[I-D.ietf-bier-mpls-encapsulation](#)] specifies a BIER encapsulation that BIER header contains a bitstring in which each bit represents exactly one egress router in the domain. But in BIER-TE every BitPosition of the BitString of a BIER-TE packet indicates one or more adjacencies instead of an egress router as in BIER. That MUST be a huge number of adjacencies from BFIR to all BFERs and the BitString in BIER encapsulation is related to SD,BSL and SI combination. For these distinct SD,BSL and SI combinations, there

must be more than one BitStrings and the BFR must make many copies of multicast data packet. Even more, BitPositions of all adjacencies passing through BFIR to a BFER MAY be carried in different BitStrings and within the different packets.

Based on the discussion above, this document proposes a set of extensions to BIER encapsulation for BIER-TE. The extensions define the BIER-TE header which contains one or more bitstrings and each bit in each bitstring represents one or more adjacencies in BIER-TE domain. The encapsulation can be used both in an MPLS network and a non-MPLS network.

2. 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 [\[RFC2119\]](#).

2.1. Terminology

The terminology is defined as [\[I-D.ietf-bier-architecture\]](#), [\[I-D.eckert-bier-te-arch\]](#) and [\[I-D.ietf-bier-mpls-encapsulation\]](#).

3. BIER-TE Encapsulation

The BIER-TE header is shown in Figure 1. It extends the BIER encapsulation and adds one or more BitString Sub-TLVs.

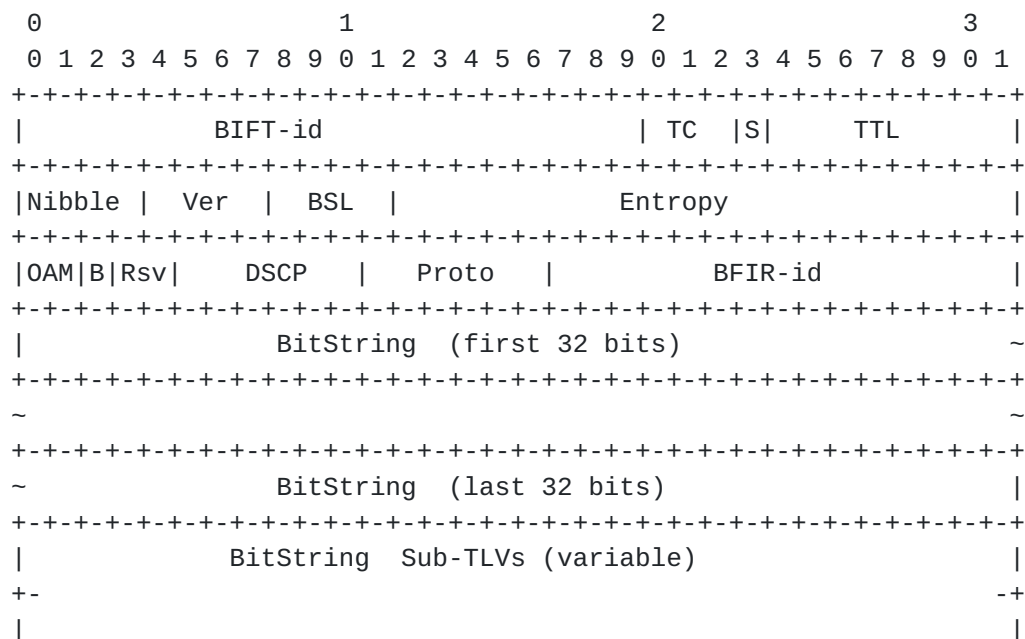


Figure 1:BIER-TE Header Format

B bit : 1bit, indicates BIER-TE packet when it is set. As the [I-D.eckert-bier-te-arch] mentioned, when a BFR receives a packet, it needs to interpret the BitString of a BIER-TE packet differently from a BIER packet and it is necessary to distinguish BIER from BIER-TE packets.

BitString Sub-TLV: identifies BitString related information and each BitString Sub-TLV corresponds to a particular combination of SD, BSL, SI and bitstring. The format details see [section 3.1](#).

The definition of other fields is the same with [I-D.ietf-bier-mpls-encapsulation].

3.1. BitString Sub-TLV

This document proposes BitString Sub-TLV for BIER-TE header. The TLV is optional. The format of the new sub-TLV is shown in Figure 2 and 3.

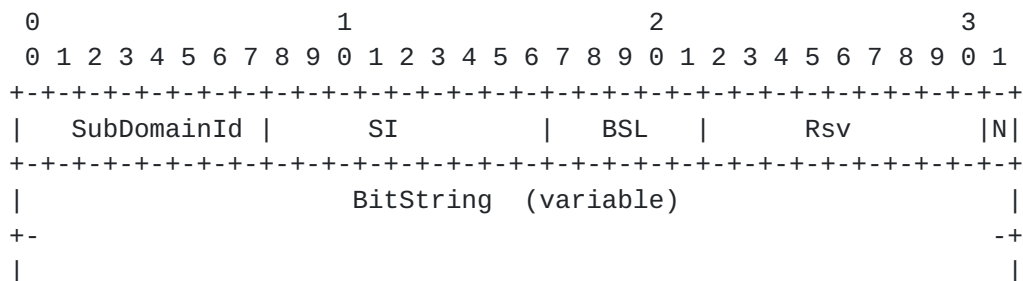


Figure 2: BitString Sub-TLV Format 1

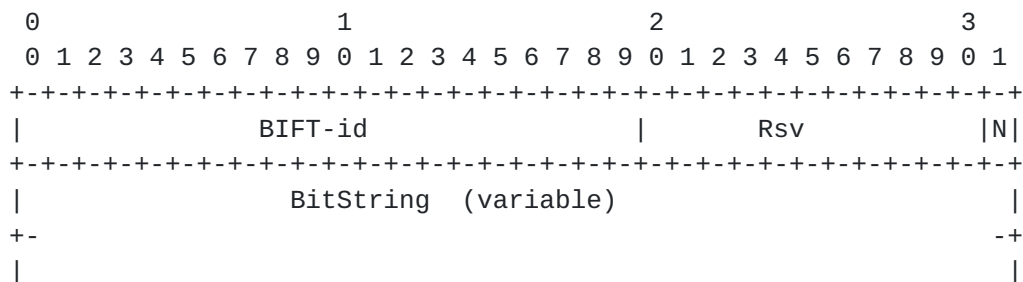


Figure 3: BitString Sub-TLV Format 2

The fields in the format are:

SD : 8bits, indicates the Sub-DomainId of the BitString in the range [0,255].

SI : 10bits, indicates the Set Identification of the BitString in the range [0,1023].

BSL : 4bits, indicates the length in bits of the BitString. If k is the length of the BitString, the value of this field is $\log_2(k)-5$. The values are supported as follows:

1: 64 bits

2: 128 bits

3: 256 bits

4: 512 bits

5: 1024 bits

6: 2048 bits

7: 4096 bits

N : 1bit, indicates that there are one or more BitString Sub-TLVs immediately preceding the TLV when it is set.

BitString: variable, identifies the collection of the adjacencies from BFIR to all BFERS and each BitString is related to SD, BSL and SI combination or BIFT-id of the packet.

BIFT-id: 22bits, The BIFT-id represents a particular Bit Index Forwarding Table (BIFT); see Section 6.4 of [\[I-D.ietf-bier-architecture\]](#). Each BIFT corresponds to a particular combination of SD, BSL and SI.

4. Processing Rules with the BIER-TE Encapsulation

As defined in [\[I-D.eckert-bier-te-arch\]](#), the BIER-TE operations consists of four layers: the "Multicast Flow Overlay", the "BIER-TE Controller Host", the "Routing Underlay" and the "BIER-TE forwarding layer". The BIER-TE Multicast flow processing with BIER-TE encapsulation is as follows:

1. The BIER-TE Controller assigns the BitPositions for adjacencies based on the operator policy and populates the BitPositions to the BIFT of each BFR as mentioned in [\[I-D.eckert-bier-te-arch\]](#).
2. The Multicast Flow Overlay determines the BFIR and a set of BFERS and sends this information to the BIER-TE controller.

3. The BIER-TE controller calculates the explicit paths based on algorithms from BFIR to all BFERs.
4. The BIER-TE controller gets all adjacencies which the paths passing through and determines the list of bitstrings based on the SD,BSL and SI combination and BitPositions/adjacencies assignments. Each bit/BitPosition in each bitstring represents one or more adjacencies in BIER-TE domain. The BitPositions of the adjacencies that have the same SD,BSL and SI combination can be encoded into the same BitString. It then pushes the BitStrings into the BFIR.
5. When a BFIR receives a multicast packet from outside the BIER-TE domain, the BFIR carries out the following procedure:
 - a. The BFIR makes a copy of the multicast data packet and encapsulates the copy in a BIER-TE header as this document proposes(see [Section 3](#)). The BitStrings which received from the BIER-TE controller are mapped to the field of BitString Sub-TLVs.
 - b. The BFIR checks the BIER-TE header and get the BitString Sub-TLVs information. Then traverses the Sub-TLVs and related local BIFT which has the same SD,BSL and SI combination. The packet may then be transmitted to adjacencies/neighboring BFRs and applies to that copy with the forwarding procedure of [[I-D.eckert-bier-te-arch](#)].

5. Security Considerations

TBD.

6. IANA Considerations

TBD.

7. Acknowledgements

TBD.

8. Normative References

[[I-D.eckert-bier-te-arch](#)]

Eckert, T., Cauchie, G., Braun, W., and M. Menth, "Traffic Engineering for Bit Index Explicit Replication BIER-TE", [draft-eckert-bier-te-arch-05](#) (work in progress), June 2017.

[I-D.ietf-bier-architecture]

Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", [draft-ietf-bier-architecture-07](#) (work in progress), June 2017.

[I-D.ietf-bier-mpls-encapsulation]

Wijnands, I., Rosen, E., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication in MPLS and non-MPLS Networks", [draft-ietf-bier-mpls-encapsulation-07](#) (work in progress), June 2017.

[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>>.

Authors' Addresses

Quan Xiong
ZTE Corporation
No.6 Huashi Park Rd
Wuhan, Hubei 430223
China

Phone: +86 27 83531060
Email: xiong.quan@zte.com.cn

Fangwei Hu
ZTE Corporation
No.889 Bibo Rd
Shanghai 201203
China

Phone: +86 21 68896273
Email: hu.fangwei@zte.com.cn

Zheng Zhang
ZTE Corporation
No.50 Software Avenue, Yuhuatai District
Nanjing, Jiangsu Province 210012
China

Phone: +86 025 88016576
Email: zhang.zheng@zte.com.cn

