

Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 30, 2019

Ran. Chen
Greg. Mirsky
Shaofu. Peng
ZTE Corporation
February 26, 2019

BIER-TE Ping and Trace
draft-chen-bier-te-ping-04

Abstract

Bit Index Explicit Replication (BIER)-TE shares architecture and packet formats with BIER-TE forwards and replicates packets based on a BitString in the packet header, but every BitPosition of the BitString of a BIER-TE packet indicates one or more adjacencies.

This document describes the mechanism and basic BIER-TE OAM packet format that can be used to perform Ping and Traceroute on the BIER-TE 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 <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 30, 2019.

Copyright Notice

Copyright (c) 2019 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. Conventions used in this document](#) [3](#)
- [3. BIER-TE OAM Packet format](#) [3](#)
 - [3.1. Target FEC Stack](#) [3](#)
 - [3.1.1. BIER-TE forward-connected TLV](#) [4](#)
 - [3.1.2. BIER-TE local-decap sub-TLV](#) [5](#)
 - [3.1.3. BIER-TE forward-routed TLV](#) [6](#)
 - [3.2. Downstream Mapping TLV](#) [6](#)
 - [3.2.1. FEC Stack Change Sub-TLV](#) [6](#)
 - [3.3. Reply-To TLV](#) [7](#)
- [4. BIER-TE OAM Processing](#) [8](#)
 - [4.1. Sending BIER Echo Request](#) [8](#)
 - [4.2. Receiving BIER Echo Request](#) [9](#)
 - [4.3. Sending Echo Reply](#) [10](#)
 - [4.4. Receiving Echo Reply](#) [11](#)
- [5. Security Considerations](#) [11](#)
- [6. IANA Considerations](#) [12](#)
 - [6.1. TLVs](#) [12](#)
 - [6.2. Target FEC Stack](#) [12](#)
 - [6.3. Downstream Detailed Mapping Sub-TLVs](#) [12](#)
- [7. Normative references](#) [12](#)
- Authors' Addresses [14](#)

1. Introduction

[I-D.ietf-bier-te-arch] introduces and explains BIER-TE architecture that provides policy-based multicast forwarding through a "BIER-TE domain" without requiring intermediate routers to maintain any multicast related per-flow state. BIER-TE forwards and replicates packets based on a BitString in the packet header, but every BitPosition of the BitString of a BIER-TE packet indicates one or more adjacencies.

This document describes the mechanism and the basic BIER-TE OAM packet format that can be used to perform Ping and Traceroute on the BIER-TE network.

This document enhances the BIER Ping and Traceroute, as defined in [I-D.ietf-bier-ping]. [RFC8296] defines a 4-bit field as "Proto" to identify the payload following BIER header. When the payload is

BIER-TE OAM, the "Proto" field is the same with the BIER OAM "Proto" field.

2. Conventions used in this document

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](#).

3. BIER-TE OAM Packet format

The BIER-TE OAM packet header format and the fields are the same as the BIER OAM header [[I-D.ietf-bier-ping](#)]. This document defines two new return codes and the new TLVs and Sub-TLVs.

The new Return codes are follows:

TBA1 Replying BFR is not in the path to any target BFER

TBA2 Mapping for this FEC is not the given BitPosition in BitString

The TLVs and Sub-TLVs requested by this document for IANA consideration are the following:

Type	value field
9	Target FEC Stack
10	Reply-To TLV

3.1. Target FEC Stack

A BIER-TE echo request MAY include the Target FEC Stack TLV that describes the FEC Stack being tested. If there aren't adjacency keyword information in BFIR, the FEC Stack MUST not be tested, and the Nil FEC MUST be used.

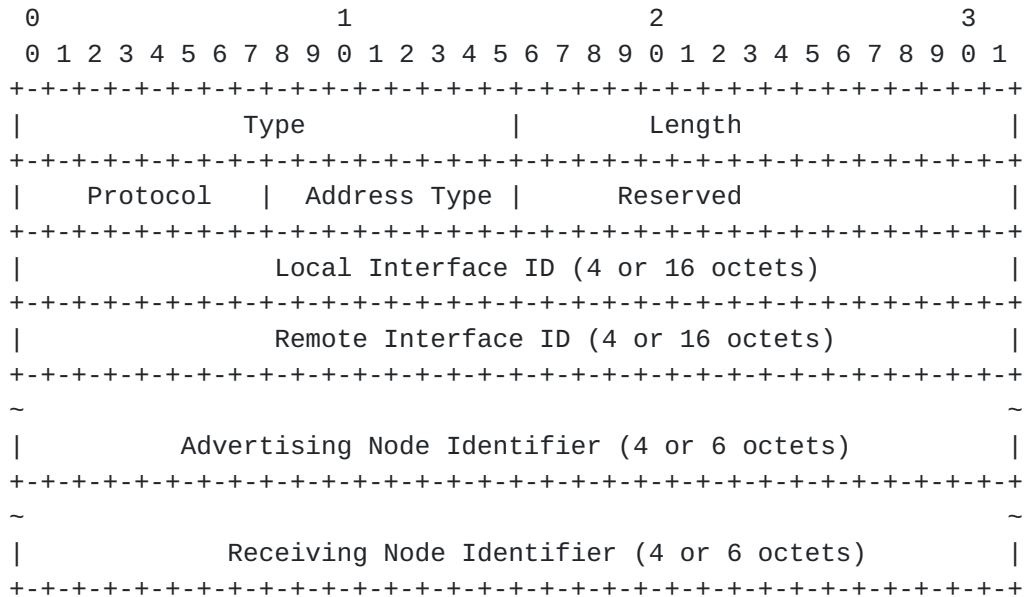
We define three new FEC Stack types. The Target FEC Stack is a list of sub-TLVs.

Sub-Type	Length	Value Field
29	20 or 48 octets	BIER-TE forward-connected
30	8 or 10 octets	BIER-TE local-decap
31	4 or 6 octets	BIER-TE forward-routed

Other FEC Types will be defined as needed.

3.1.1. BIER-TE forward-connected TLV

The format is as below:



Address Type

The Address Type indicates the address type and length of the IP address for the downstream interface. The Address type MAY be set to one of the values listed in Table below:

Type	Addr. Type	DA Length	DIA Length
1	IPv4 Numbered	4	4
2	IPv4 Unnumbered	4	4
3	IPv6 Numbered	16	16
4	IPv6 Unnumbered	16	4

DA Length - Downstream Address field Length.

DIA Length - Downstream Interface Address field Length.

Local Interface ID

Local BFR assigns Local Interface ID for a link to which Adjacency ID is bound. This field is set to local link address (IPv4 or IPv6).

Remote Interface ID

Remote BFR assigns Remote Interface ID for a link to which Adjacency ID is bound. This field is set to remote link address (IPv4 or IPv6).

Advertising Node Identifier

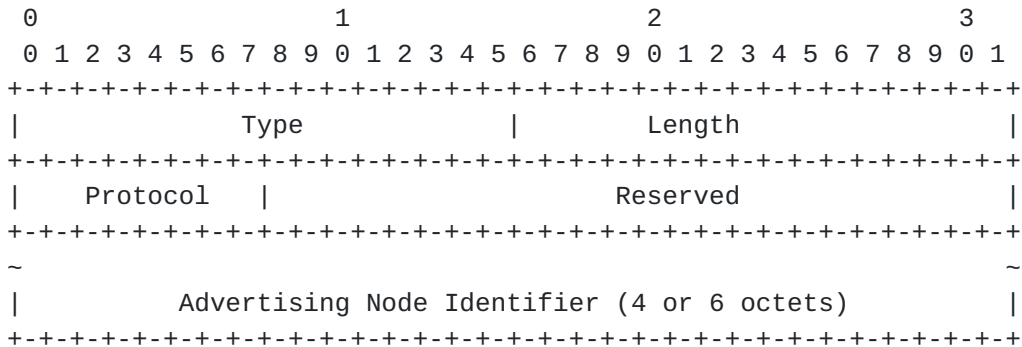
Advertising Node Identifier is the advertising node identifier. When Protocol is set to 1, then the 32 rightmost bits represent OSPF Router ID, and if Protocol is set to 2, this field carries 48 bit ISIS System ID.

Receiving Node Identifier

Receiving Node Identifier is the downstream node identifier. When Protocol is set to 1, then the 32 rightmost bits represent OSPF Router ID, and if Protocol is set to 2, this field carries 48 bit ISIS System ID.

3.1.2. BIER-TE local-decap sub-TLV

The format is as below:




Advertising Node Identifier

Advertising Node Identifier is the advertising node identifier. When Protocol is set to 1, then the 32 rightmost bits represent OSPF Router ID and if protocol is set to 2, this field carries 48 bit ISIS System ID.


3.1.3. BIER-TE forward-routed TLV

The ipv4 format is as below:



IPv4 Prefix: This field carries the IPv4 prefix.

The ipv6 format is as below:



IPv6 Prefix: This field carries the IPv6 prefix.

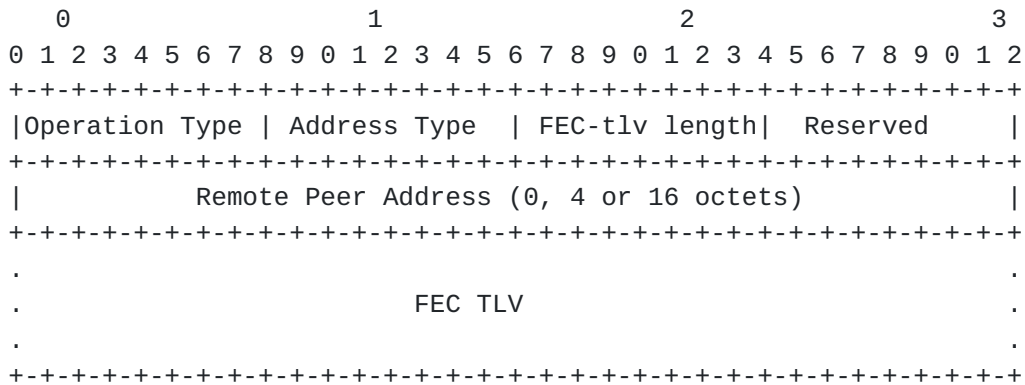
3.2. Downstream Mapping TLV

This TLV format is the same with the BIER OAM [I-D.ietf-bier-ping] and we a new Sub-TLV: FEC stack change Sub-TLV.

Sub-TLV Type	Value
3	FEC stack change

3.2.1. FEC Stack Change Sub-TLV

The format and the usage as defined in [RFC6424].



Operation Type

The operation type specifies the action associated with the FEC Stack Change. A new operation type is defined:

Type	Operation
3	Remove

Address type: 0.

FEC TLV Length: Length in bytes of the FEC TLV.

Reserved: This field is reserved for future use and MUST be set to zero.

Remote Peer Address: 0.

FEC TLV

The FEC TLV is present only when the FEC-tlv length field is nonzero. The FEC TLV specifies the FEC associated with the FEC stack change operation. The FEC type is defined in [section 3.1](#).

3.3. Reply-To TLV

The Initiator BFR MAY include the Reply-To TLV in Echo Request message. The Reply-to TLV is used by a transit BFR or BFER when the reply mode is "Reply via IPv4/IPv6 UDP packet".

The IP address will be used as the destination IP address for the Echo Reply. The format and usage are the same as defined for the BIER OAM header [[I-D.ietf-bier-ping](#)].

4. BIER-TE OAM Processing

BIER-TE OAM packet MUST be sent to BIER control plane for OAM processing if one of the following conditions is true:

- o The receiving BFR is a BFER.
- o TTL of BIER-MPLS Label expired.
- o Presence of Router Alert label in the label stack.

4.1. Sending BIER Echo Request

- o Message Type:1.
- o Return Code:0.
- o Proto:0.
- o Sender's Handle and Sequence number:The local matter to Initiator and SHOULD increment the Sequence number by 1 for every subsequent Echo Request.
- o QTF:Initiator's local timestamp format.
- o TimeStamp Sent:the time that the Echo Request is sent.
- o MUST include Original SI-BitString TLV.
- o In Ping mode, the Initiator MAY include Target SI-BitString TLV to control the responding BFER(s) by listing all local-decap Adjacency ID of the BFERs from which the Initiator expects a response. Initiator on receiving a reply with Return code as "Replying BFR is the only BFER in header Bitstring" or "Replying router is one of the BFER in header Bitstring", SHOULD remove the BFER's local-decap ID from Target SI-BitString for any subsequent Echo Request.
- o When the Reply mode is set to 2, Initiator MUST include Reply-To TLV in the Echo Request.
- o The Initiator MAY include Downstream Mapping TLV in the Echo Request to query additional information from transit BFRs and BFERs. In the case of ECMP discovery, Initiator MUST include the Multipath Entropy Data Sub-TLV and SHOULD set the Target SI-BitString TLV carrying a specific BFER's local-decap Adjacency ID.

- o The Initiator MUST encapsulate the OAM packet with BIER header and MUST set the Proto as 6 and further encapsulates with BIER-MPLS label. In ping mode, the BIER-MPLS Label TTL MUST be set to 255. In traceroute mode, the BIER-MPLS Label TTL is set successively starting from 1 and MUST stop sending the Echo Request if it receives a reply with Return code as "Replying router is the only BFER in BIER header Bitstring" from all BFER listed in Target SI-BitString TLV.
- o MUST PUSH the corresponding FEC to Target FEC stack, in the same with as the order of the adjacency's bit-position in the BitString.

4.2. Receiving BIER Echo Request

Reply-Flag: This flag is initially set to 1.

Interface-I: The incoming interface on which the Echo Request was received.

BIER-Label-L: The BIER-MPLS Label received as the top label on received Echo Request.

Header-H: The BIER header from the received Echo Request.

Best-return-code: contains the return code for the echo reply packet as currently best known.

If the received Echo Request carries Target SI-BitString TLV, a BFR SHOULD run boolean AND operation between BitString in Header-H and BitString in Target SI-BitString TLV.

If the resulting BitString is all-zero, Set Best-return-code to "Mapping for this FEC is not the given BitPosition in BitString" and Go to [section 4.3](#), Else:

- o If the BIER-Label-L does not correspond to the local label assigned for {sub-domain, BitStringLength, SI} in Original SIBitString TLV, Set the Best-return-code to "Set-Identifier Mismatch" and Go to [section 4.3](#).
- o If any of the TLVs in Echo Request message is not understood. Set the Best-return-code to "One or more of the TLVs was not understood" and Go to [section 4.3](#).
- o If the forwarding lookup defined in [section 6.5 of RFC8279](#) does not match any entry for the received BitString in BIER header.

Set the Best-return-code to "No matching entry in forwarding table" and Go to [section 4.3](#).

- o If any FEC which get from the matched BIFT entry is not consistent with the FEC get from the FEC stack at the same position as entry's BitPosition in Header-H, Set the Best-return-code to "Mapping for this FEC is not the given BitPosition in BitString" and Go to [section 4.3](#).
- o If the DSMAP TLV carries Multipath Entropy Data Sub-TLV and if the BitString in Header-H carries more than one forward routed adjacency and each matches the BIFT entry. Set the Best-return-code to "Invalid Multipath Info Request" and Go to [section 4.3](#). Else, list the ECMP downstream neighbors to reach forward routed adjacency, calculate the Entropy considering the BitString in Header-H and Multipath Entropy Data Sub-TLV from received Echo Request. Set the Best-return-code to 5 (Packet-Forward-Success).
- o For all the forward-connected adjacencies and all the local-decap adjacencies which match the BIFT entry, FEC Change sub-TLV SHOULD be carried in DSMAP TLV and set the operation type filed in the FEC change sub-TLV to remove.
- o For all the forward-routed adjacencies which match the BIFT entry, if the BIFT entry indicates that not the local decapsulation but continue forwarding the OAM packet, FEC change sub-TLV SHOULD NOT be carried in DSMAP TLV. If the BIFT entry indicate that the local decapsulation the OAM packet, FEC change sub-TLV SHOULD be carried in DSMAP TLV, and set the operation type filed in the FEC change sub-TLV to remove.
- o If the responder is BFER which matches the local-decap BIFT, and there are no more bits in BIER header BitString left for forwarding. Set the Best-return-code to "Replying router is the only BFER in BIER header BitString", and go to [section 4.3](#).
- o If the responder is BFER which match the local-decap BIFT, and there are more bits in the BitString left for forwarding. Set the Best-return-code to "Replying router is one of the BFER in BIER header BitString", and go to [section 4.3](#).

[4.3](#). Sending Echo Reply

- o Message Type:2.
- o Return Code:Best-return-code.
- o The Proto :0.

- o When the Best-return-code is "Replying BFR is one of the BFER in header BitString", it MUST include Responder BFER TLV.
- o If the received Echo Request had DSMAP with Multipath Entropy Data Sub-TLV, Responder BFR MUST include DSMAP for each outgoing interface over which the packet will be replicated and include the respective Multipath Entropy Data Sub-TLV.
- o If the received Echo Request had DSMAP without Multipath Entropy Data Sub-TLV, Responder BFR MUST include DSMAP for each outgoing interface over which the packet will be replicated.
- o When the Best-return-code is "Replying BFR is the only BFER in header BitString", it MUST include Responder BFER TLV.
- o When the Reply mode in received Echo Request is set to " Reply via BIER packet", Responder appends BIER header listing the BitString with the BFIR's local-decap id and set the Proto to "OAM" and set the BFIR value to 0.
- o When the Reply mode in received Echo Request is set to "Reply via IPv4/IPv6 UDP packet", Responder encapsulates with IP/UDP header. The UDP destination port MUST be set to TBD1 and source port MAY be randomly selected from the dynamic range of port numbers. The source IP is any local address of the responder and destination IP is derived from Reply-To TLV.

4.4. Receiving Echo Reply

- o Initiator on receiving Echo Reply will use the Sender's Handle to match with Echo Request sent. If no match is found, Initiator MUST ignore the Echo Reply.
- o If receiving Echo Reply have Downstream Mapping, Initiator SHOULD copy the same to subsequent Echo Request(s).
- o If one of the Echo Reply is received with Return Code as "Replying BFR is one of the BFER in header BitString", it SHOULD remove the BFER's local-decap ID from Target SI-BitString for any subsequent Echo Request.

5. Security Considerations

TBD.

6. IANA Considerations

This document request UDP port TBD1 to be allocated by IANA for BIER-TE Echo.

This document request the IANA for creation and management of below registries and sub-registries:

Return codes defined in this document are the following:

Value	value Meaning
10	Replying BFR is not in the path to any target BFER
11	Mapping for this FEC is not the given BitPosition in BitString

6.1. TLVs

The TLVs and Sub-TLVs requested by this document for IANA consideration are the following:

6.2. Target FEC Stack

Sub-Type	Value Field
29	BIER-TE forward-connected
30	BIER-TE local-decap
31	BIER-TE forward-routed

6.3. Downstream Detailed Mapping Sub-TLVs

This section defines the optional Sub-TLVs that can be included in Downstream Mapping TLV.

Sub-TLV Type	Value
3	FEC stack change

7. Normative references

[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-04](#) (work in progress), October 2018.

[I-D.ietf-bier-te-arch]

Eckert, T., Cauchie, G., Braun, W., and M. Menth, "Traffic Engineering for Bit Index Explicit Replication (BIER-TE)", [draft-ietf-bier-te-arch-01](#) (work in progress), October 2018.

- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, [RFC 792](#), DOI 10.17487/RFC0792, September 1981, <<https://www.rfc-editor.org/info/rfc792>>.
- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", [RFC 4379](#), DOI 10.17487/RFC4379, February 2006, <<https://www.rfc-editor.org/info/rfc4379>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", [RFC 5905](#), DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.
- [RFC6424] Bahadur, N., Kompella, K., and G. Swallow, "Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels", [RFC 6424](#), DOI 10.17487/RFC6424, November 2011, <<https://www.rfc-editor.org/info/rfc6424>>.
- [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>>.
- [RFC8444] Psenak, P., Ed., Kumar, N., Wijnands, IJ., Dolganow, A., Przygienda, T., Zhang, J., and S. Aldrin, "OSPFv2 Extensions for Bit Index Explicit Replication (BIER)", [RFC 8444](#), DOI 10.17487/RFC8444, November 2018, <<https://www.rfc-editor.org/info/rfc8444>>.

Authors' Addresses

Ran Chen
ZTE Corporation

Email: chen.ran@zte.com.cn

Greg Mirsky
ZTE Corporation

Email: gregimirsky@gmail.com

Shaofu Peng
ZTE Corporation

Email: peng.shaofu@zte.com.cn

