

**BIER-TE Ping and Trace**  
**draft-chen-bier-te-ping-00**

Abstract

Bit Index Explicit Replication (BIER)-TE shares architecture and packet formats with BIER as described in [\[I-D.ietf-bier-architecture\]](#). 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 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 <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 17, 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

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">2</a>
<a href="#">2.</a>	<a href="#">Conventions used in this document</a>	<a href="#">3</a>
<a href="#">3.</a>	<a href="#">BIER-TE OAM Packet format</a>	<a href="#">3</a>
<a href="#">3.1.</a>	<a href="#">Target FEC Stack</a>	<a href="#">5</a>
<a href="#">3.1.1.</a>	<a href="#">BIER-TE forward_connected TLV</a>	<a href="#">6</a>
<a href="#">3.1.2.</a>	<a href="#">BIER-TE local_decap TLV</a>	<a href="#">7</a>
<a href="#">3.1.3.</a>	<a href="#">BIER-TE forward_routed TLV</a>	<a href="#">7</a>
<a href="#">3.2.</a>	<a href="#">Downstream Mapping TLV</a>	<a href="#">8</a>
<a href="#">3.2.1.</a>	<a href="#">Downstream Mapping Sub-TLVs</a>	<a href="#">8</a>
<a href="#">3.2.1.1.</a>	<a href="#">Multipath Entropy Data</a>	<a href="#">9</a>
<a href="#">3.2.1.2.</a>	<a href="#">Egress BitString sub-TLV</a>	<a href="#">9</a>
<a href="#">3.2.1.3.</a>	<a href="#">FEC Stack Change Sub-TLV</a>	<a href="#">9</a>
<a href="#">3.3.</a>	<a href="#">Original SI-BitString TLV</a>	<a href="#">10</a>
<a href="#">3.4.</a>	<a href="#">Target SI-BitString TLV</a>	<a href="#">11</a>
<a href="#">3.5.</a>	<a href="#">Responder BFER TLV</a>	<a href="#">11</a>
<a href="#">3.6.</a>	<a href="#">Responder BFR TLV</a>	<a href="#">11</a>
<a href="#">3.7.</a>	<a href="#">Reply-To TLV</a>	<a href="#">11</a>
<a href="#">4.</a>	<a href="#">BIER-TE OAM Processing</a>	<a href="#">11</a>
<a href="#">4.1.</a>	<a href="#">Sending BIER Echo Request</a>	<a href="#">11</a>
<a href="#">4.2.</a>	<a href="#">Receiving BIER Echo Request</a>	<a href="#">12</a>
<a href="#">4.3.</a>	<a href="#">Sending Echo Reply</a>	<a href="#">14</a>
<a href="#">4.4.</a>	<a href="#">Receiving Echo Reply</a>	<a href="#">15</a>
<a href="#">5.</a>	<a href="#">Security Consideration</a>	<a href="#">15</a>
<a href="#">6.</a>	<a href="#">Acknowledgements</a>	<a href="#">15</a>
<a href="#">7.</a>	<a href="#">IANA Considerations</a>	<a href="#">15</a>
<a href="#">8.</a>	<a href="#">References</a>	<a href="#">15</a>
<a href="#">8.1.</a>	<a href="#">Normative references</a>	<a href="#">15</a>
<a href="#">8.2.</a>	<a href="#">Informative references</a>	<a href="#">16</a>
	<a href="#">Authors' Addresses</a>	<a href="#">17</a>

## [1. Introduction](#)

[I-D.ietf-bier-architecture] introduces and explains BIER-TE architecture that provides optimal 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 basic BIER-TE OAM packet format that can be used to perform Ping and Traceroute on BIER-TE network.

This document is a supplement to [[I-D.kumarzheng-bier-ping](#)]. BIER-MPLS [[I-D.ietf-bier-mpls-encapsulation](#)] defines a 4-bit field as "Proto" to identify the payload following BIER header. When the payload is BIER-TE OAM, the "Proto" field will be set to 6 as defined in this document.

## 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 is similar with the BIER OAM header as described in [[I-D.kumarzheng-bier-ping](#)].

The BIER-TE OAM packet header format is 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Ver | Echo Req/Rep | Proto |                               Reserved |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| QTF | RTF | Reply mode | Return Code |           Return Subcode |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Sender's Handle |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Sequence Number |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               TimeStamp Sent |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               TimeStamp Sent |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               TimeStamp Received |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               TimeStamp Received |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                               TLVs                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Ver:Set to 1.



Proto:Set to 0 for Echo Request/Reply header.

QTF:Querier Timestamp Format.When set to 2, the Timestamp Sent field is (in seconds and microseconds, according to the Initiator's clock) in NTP format [[RFC5905](#)]. When set to 3, the timestamp format is in IEEE 1588-2008 (1588v2) Precision Time Protocol format. Any other value SHOULD be considered as sanity check failure.

RTF:Responder Timestamp Format. When set to 2, the Timestamp Received field is (in seconds and microseconds, according to the Initiator's clock) in NTP format [[RFC5905](#)]. When set to 3, the timestamp format is in IEEE 1588-2008 (1588v2) Precision Time.

Reply mode:The Reply mode is set to one of the below:

Value	Meaning
-----	-----
1	Do not Reply
2	Reply via IPv4/IPv6 UDP packet.
3	Reply via BIER-TE packet

Return Code:Set to zero if Type is " BIER Echo Request". Set to the following value, if Type is "BIER Echo Reply".

Value	Value Meaning
-----	-----
0	No return code
1	Malformed Echo Request received
2	One or more of the TLVs was not understood
3	Replying BFR is the only BFER in header Bitstring
4	Replying BFR is one of the BFER in header Bitstring
5	Packet-Forward-Success
6	Invalid Multipath Info Request
8	No matching entry in forwarding table.
9	Set-Identifier Mismatch
10	Replying BFR is not in the path to any target BFER
11	Mapping for this FEC is not the given bitposition in

bitstring

Return subcode:To Be updated.

Sender's Handle:The Sender's Handle is filled by the Initiator, and returned unchanged by responder BFR. This is used for matching the replies to the request.



Sequence number: The Sequence number is assigned by the Initiator and can be used to detect any missed replies.

Timestamp: The Timestamp Sent is the time when the Echo Request is sent. The TimeStamp Received in Echo Reply is the time (accordingly to responding BFR clock) that the corresponding Echo Request was received. The format depends on the QTF/RTF value.

TLVs have the following format:

[illegible]

Length is the length of the Value field in octets. The Value field depends on the TLV Type.

A description of the Types and Values for TLLVs are given below:

Type#	value field
1	Target FEC Stack
2	Downstream Mapping
3	Original SI-BitString TLV
4	Target SI-BitString TLV
5	Responder BFER TLV
6	Responder BFR TLV
7	Reply-To TLV

### 3.1. Target FEC Stack

A Target FEC Stack is a list of sub-TLVs.



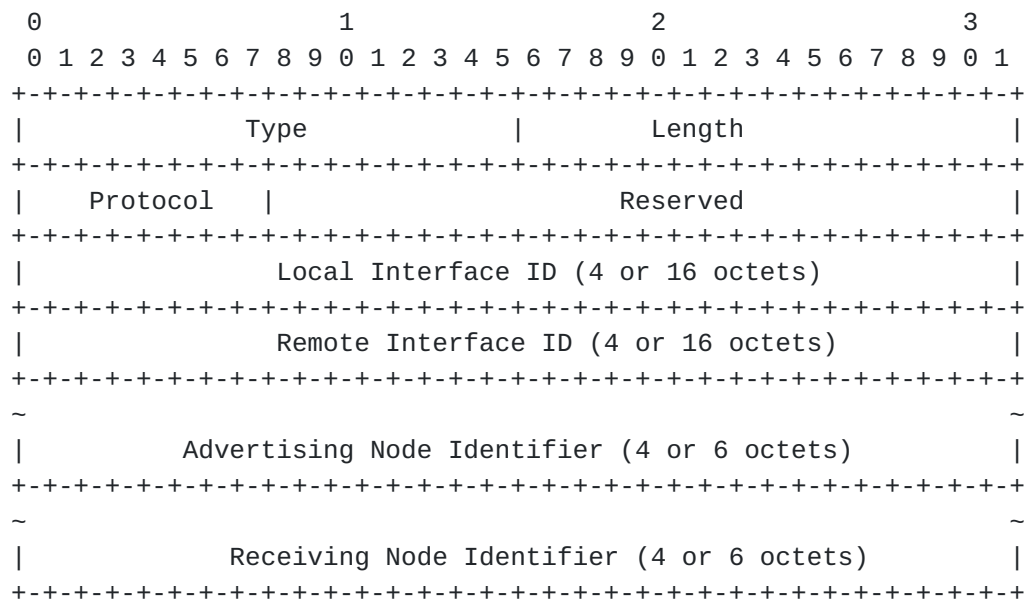


Sub-Length	Value Field
-----	-----
29	BIER-TE forward_connected TLV
30	BIER-TE local_decap TLV
31	BIER-TE forward_routed TLV

Other FEC Types will be defined as needed.

### [3.1.1.](#) BIER-TE forward\_connected TLV

The format is as below:



#### Local Interface ID

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

#### Remote Interface ID

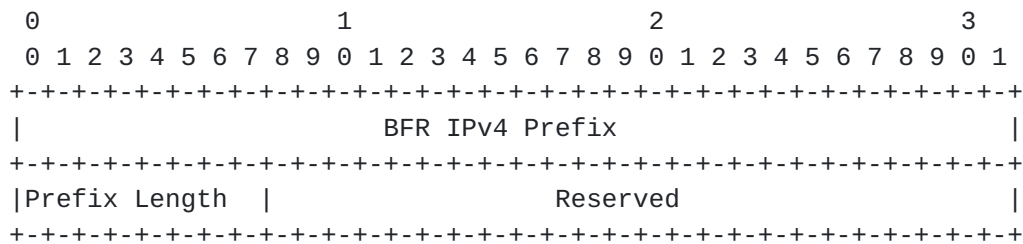
Remote Interface ID is assigned by remote BFR for a link on which Adjacency ID is bound. This field is set to remote link address (IPv4 or IPv6).

#### Advertising Node Identifier



The ipv4 format is as below:

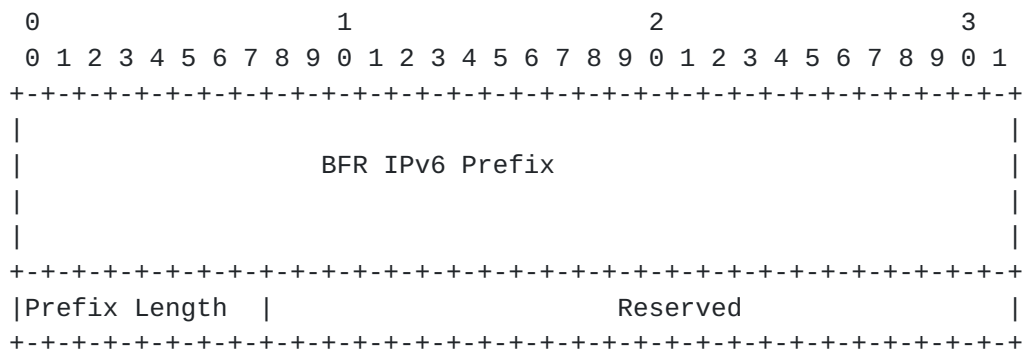




IPv4 Prefix: This field carries the IPv4 prefix.

Prefix Length is one octet, it gives the length of prefix in bits.

The ipv6 format is as below:



IPv6 Prefix: This field carries the IPv4 prefix.

Prefix Length is one octet, it gives the length of prefix in bits.

### 3.2. Downstream Mapping TLV

The TLV format is similar with Downstream Detailed Mapping TLV as described in [I-D.kumarzheng-bier-ping].

### 3.2.1. Downstream Mapping Sub-TLVs

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

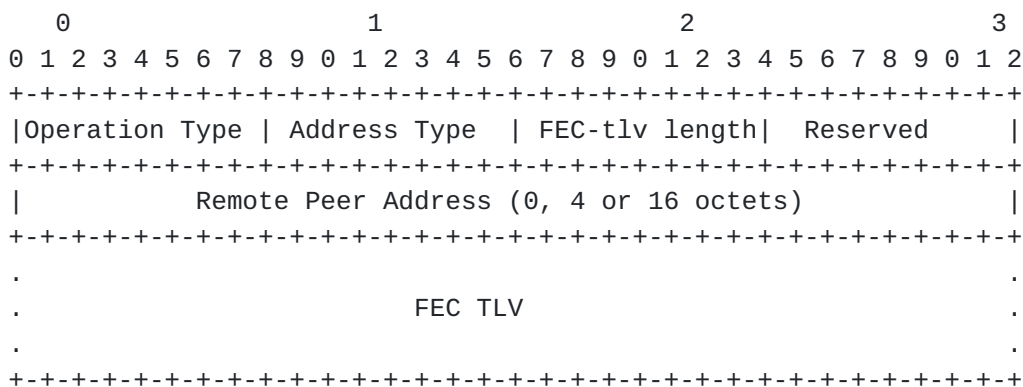
Sub-TLV Type	Value
1	Multipath Entropy Data
2	Egress BitString
3.	FEC stack change











### 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. Original SI-BitString TLV

The Incoming SI-BitString TLV will be included by Responder BFR in Reply message and copies the BitString from BIER header of incoming Echo Request message. The format and usage is similar with Original SI-BitString TLV as defined in [[I-D.kumarzheng-bier-ping](#)].



### **3.4. Target SI-BitString TLV**

The Target SI-BitString TLV carries the set of BFER's local\_decap adjacency from which the Initiator expects the reply from. The format and usage is similar with Target SI-BitString TLV as defined in [[I-D.kumarzheng-bier-ping](#)].

### **3.5. Responder BFER TLV**

The Responder BFER TLV will be included by the BFER replying to the request. This is used to identify the originator of BIER Echo Reply. The format and usage is similar with Responder BFER TLV as defined in [[I-D.kumarzheng-bier-ping](#)].

### **3.6. Responder BFR TLV**

The Responder BFR TLV will be included by the transit BFR replying to the request. This is used to identify the replying BFR without BFRID. The format and usage is similar with Responder BFR TLV as defined in [[I-D.kumarzheng-bier-ping](#)].

### **3.7. Reply-To TLV**

The Reply-To TLV MAY be included by the Initiator BFR in Echo Request. This is used by transit BFR or BFER when the reply mode is the IP address will be used to generate Echo Reply. The format and usage is similar with Reply-To TLV as defined in [[I-D.kumarzheng-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, 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 Initiator MAY include Downstream Mapping TLV in the Echo Request to query additional information from transit BFRs and BFERs. In 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 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, which the push order is the same with adjacency BitPosition of 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="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 BitString in Header-H does not match the BitString in Egress BitString Sub-TLV of DSMAP TLV, set the Best-return-code to ERR-TBD and Go to [section 4.3](#).
- o If the forwarding lookup defined in section 6.5 of [\[I-D.ietf-bier-architecture\]](#) 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 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 adjacency and local\_decap adjacency 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 adjacency which match the BIFT entry, if the BIFT entry indicate that not local decapsulation but continue forwarding the OAM packet, FEC change sub-TLV should not carried in DSMAP TLV. If the BIFT entry indicate that 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 match the local\_decap BIFT, and there is 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 in 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. For each outgoing interface, respective Egress BitString MUST be included in DSMAP 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. For each outgoing interface, respective Egress BitString MUST be included in DSMAP TLV.



- 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 3, Responder appends BIER header listing the BitString with the BFIR's local\_decap id and set the Proto to 6 and set the BFIR as 0.
- o When the Reply mode in received Echo Request is set to 2, Responder encapsulates with IP/UDP header. The UDP destination port MUST be set to TBD1 and source port MAY be set to TBD1 or other random local value. 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 Consideration**

The section will be added in next version.

#### **6. Acknowledgements**

TBD.

#### **7. IANA Considerations**

TBD.

#### **8. References**

##### **8.1. Normative 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-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-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.

[RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, [RFC 792](#), DOI 10.17487/RFC0792, September 1981, <<http://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, <<http://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, <<http://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, <<http://www.rfc-editor.org/info/rfc6424>>.

## 8.2. Informative references

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

Eckert, T. and G. Cauchie, "Traffic Engineering for Bit Index Explicit Replication BIER-TE", [draft-eckert-bier-te-arch-02](#) (work in progress), October 2015.

[I-D.kumarzheng-bier-ping]

Kumar, N., Pignataro, C., Akiya, N., Zheng, L., Chen, M., and G. Mirsky, "BIER Ping and Trace", [draft-kumarzheng-bier-ping-02](#) (work in progress), December 2015.



Authors' Addresses

Ran Chen  
ZTE Corporation  
No.50 Software Avenue,Yuhuatai District  
Nanjing, Jiangsu Province 210012  
China

Phone: +86 025 88014636  
Email: chen.ran@zte.com.cn

Shaofu Peng  
ZTE Corporation

Email: peng.shaofu@zte.com.cn



