

Network Work group
Internet-Draft
Intended status: Standards Track
Expires: July 26, 2018

N. Kumar
C. Pignataro
Cisco Systems, Inc.
N. Akiya
Big Switch Networks
L. Zheng
M. Chen
Huawei Technologies
G. Mirsky
Ericsson
January 22, 2018

BIER Ping and Trace
draft-ietf-bier-ping-03

Abstract

Bit Index Explicit Replication (BIER) 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 router adds a BIER header to the packet. The BIER header contains a bit-string in which each bit represents exactly one BFER to forward the packet to. The set of BFERs to which the multicast packet needs to be forwarded is expressed by setting the bits that correspond to those routers in the BIER header.

This document describes the mechanism and basic BIER OAM packet format that can be used to perform failure detection and isolation on BIER data plane.

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 July 26, 2018.

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	3
2.	Conventions used in this document	3
2.1.	Terminology	3
2.2.	Requirements notation	3
3.	BIER OAM	4
3.1.	BIER OAM message format	4
3.2.	Return Code	6
3.3.	BIER OAM TLV	8
3.3.1.	Original SI-BitString TLV	8
3.3.2.	Target SI-BitString TLV	9
3.3.3.	Incoming SI-BitString TLV	10
3.3.4.	Downstream Mapping TLV	10
3.3.5.	Responder BFER TLV	13
3.3.6.	Responder BFR TLV	14
3.3.7.	Upstream Interface TLV	14
3.3.8.	Reply-To TLV	15
4.	Procedures	16
4.1.	BIER OAM Processing	16
4.2.	Per BFER ECMP Discovery	16
4.3.	Sending BIER Echo Request	17
4.4.	Receiving BIER Echo Request	18
4.5.	Sending Echo Reply	20
4.6.	Receiving Echo Reply	21
5.	IANA Considerations	21
5.1.	Message Types, Reply Modes, Return Codes	21
5.2.	TLVs	21

6.	Security Considerations	22
7.	Acknowledgement	22
8.	Contributing Authors	22
9.	References	22
9.1.	Normative References	22
9.2.	Informative References	23
	Authors' Addresses	24

[1.](#) Introduction

[RFC8279] introduces and explains BIER 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 router adds a BIER header to the packet. The BIER header contains a bit-string in which each bit represents exactly one BFER to forward the packet to. The set of BFERs to which the multicast packet needs to be forwarded is expressed by setting the bits that correspond to those routers in the BIER header.

This document describes the mechanism and basic BIER OAM packet format that can be used to perform failure detection and isolation on BIER data plane without any dependency on other layers like IP layer.

[2.](#) Conventions used in this document

[2.1.](#) Terminology

BFER - Bit Forwarding Egress Router

BFIR - Bit Forwarding Ingress Router

BIER - Bit Index Explicit Replication

ECMP - Equal Cost Multi-Path

OAM - Operation, Administration and Maintenance

SI - Set Identifier

[2.2.](#) Requirements notation

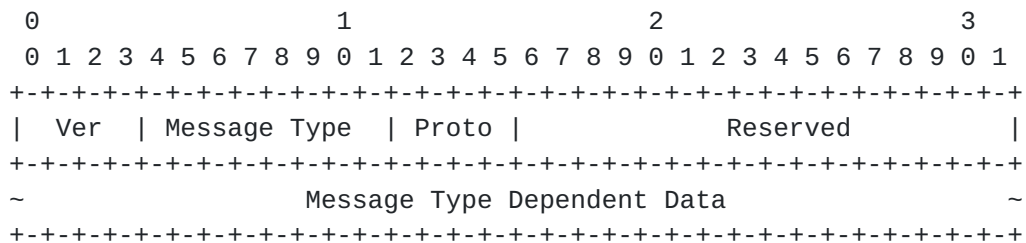
The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

3. BIER OAM

BIER OAM is defined in a way that it stays within BIER layer by following directly the BIER header without mandating the need for IP header. [RFC8296] defines a 4-bit field as "Proto" to identify the payload following BIER header. When the payload is BIER OAM, the "Proto" field will be set to 5 as defined in [RFC8296]

3.1. BIER OAM message format

The BIER OAM packet header format that follows BIER header is as follows:



Ver

Set to 1.

Message Type

This document defines the following Message Types:

Type	Value Field
-----	-----
1	BIER Echo Request
2	BIER Echo Reply

Proto

This field is used to define if there is any data packet immediately following the OAM payload which is used for passive OAM functionality. This field is set to 0 if there is no data packet following OAM payload.

The Echo Request/Reply 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                             ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

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 Protocol format. Any other value SHOULD be considered as sanity check failure.

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 packet

When Reply mode is set to 1, the receiver will not send any reply. This is used for unidirectional path validation. The Reply mode by default would be set to 2 and the Responder BFR will encapsulate the Echo reply payload with IP header. When the Initiator intend to validate the return BIER path, the Reply mode is set to 3 so that the Responder BFR will encapsulate the Echo Reply with BIER header.

Return Code

Set to zero if Type is "BIER Echo Request". Set to one of the value defined in [section 3.2](#), if Type is "BIER Echo Reply".

Return subcode

To Be updated.

Sender's Handle, Sequence number and Timestamp

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.

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

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

Carries the TLVs as defined in [Section 3.3](#).

[3.2](#). Return Code

Responder uses Return Code field to reply with validity check or other error message to Initiator. It does not carry any meaning in Echo Request and MUST be set to zero.

The Return Code can be one of the following:

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	DDMAP Mismatch

"No return code" will be used by Initiator in the Echo Request. This Value MUST NOT be used in Echo Reply.

"Malformed Echo Request received" will be used by any BFR if the received Echo Request packet is not properly formatted.

When any TLV included in the Echo Request is not understood by receiver, the Return code will be set to "One or more of the TLVs was not understood" carrying the respective TLVs.

When the received header BitString in Echo Request packet contains only its own Bit-ID, "Replying BFR is the only BFER in header BitString" is set in the reply. This implies that the receiver is BFER and the packet is not forwarded to any more neighbors.

When the received header BitString in Echo Request packet contains its own Bit-ID in addition to other Bit-IDs, "Replying BFR is one of the BFER in header BitString" is set in the reply. This implies that the responder is a BFER and the packet is further forwarded to one or more neighbors.

Any transit BFR will send the Echo Reply with "Packet-Forward-Success", if the TLV in received Echo Request are understood and forwarding table have forwarding entries for the BitString. This is used by transit BFR during traceroute mode.

When Echo Request is received with multipath info for more than one BFER, the return-code is set to "Invalid Multipath Info Request".

If the BitString cannot be matched in local forwarding table, the BFR will use "No matching entry in forwarding table" in the reply.

If the BIER-MPLS label in received Echo Request is not the one assigned for SI in Original SI-BitString TLV, "Set-Identifier Mismatch" is set inorder to report the mismatch.

BS Len is set based on the length of BitString as defined in [\[RFC8296\]](#)

The BitString field carries the set of BFR-IDs that Initiator will include in the BIER header. This TLV MUST be included by Initiator in Echo Request packet

[3.3.2.](#) Target SI-BitString TLV

The Target SI-BitString TLV carries the set of BFER from which the Initiator expects the reply from. This TLV has the following format:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Type = 2           |           Length = variable           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Set ID   | Sub-domain ID |BS Len|   Reserved   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           BitString  (first 32 bits)           ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           BitString  (last 32 bits)           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Set ID field is set to the Set Identifier to which the BitString belongs to. This value is derived as defined in [\[RFC8279\]](#)

Sub-domain ID is set to the Sub domain value to which BFER in BitString belongs to.

BS Len is set based on the length of BitString as defined in [\[RFC8296\]](#)

The BitString field carries the set of BFR-IDs of BFER(s) that Initiator expects the response from. The BitString in this TLV may be different from the BitString in BIER header and allows to control the BFER responding to the Echo Request. This TLV MUST be included by Initiator in BIER OAM packet if the Downstream Mapping TLV ([section 3.3.4](#)) is included.

3.3.3. Incoming 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. This TLV has the following format:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type = 3      |      Length = variable      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Set ID   | Sub-domain ID |BS Len|  Reserved   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      BitString (first 32 bits)      ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      BitString (last 32 bits)      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Set ID field is set to the Set Identifier to which the BitString belongs to. This value is derived as defined in [[RFC8279](#)]

Sub-domain ID is set to the Sub domain value to which BFER in BitString belongs to.

BS Len is set based on the length of BitString as defined in [[RFC8296](#)]

The BitString field copies the BitString from BIER header of the incoming Echo Request. A Responder BFR SHOULD include this TLV in Echo Reply if the Echo Request is received with I flag set in Downstream Mapping TLV.

An Initiator MUST NOT include this TLV in Echo Request.

3.3.4. Downstream Mapping TLV

This TLV has the following format:


```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Type = 4           |           Length = variable           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           MTU           | Address Type |           Flags           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Downstream Address (4 or 16 octets)           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Downstream Interface Address (4 or 16 octets)           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Sub-tlv Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
.
.           List of Sub-TLVs           .
.
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

MTU

Set to MTU value of outgoing interface.

Address Type

The Address Type indicates the address type and length of IP address for downstream interface. The Address type is set to one of the 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

Flags

The Flags field has the following format:

```

                                0 1 2 3 4 5 6 7
+---+---+---+---+---+---+---+---+
|           Rsvd           | I |
+---+---+---+---+---+---+---+---+

```


When I flag is set, the Responding BFR SHOULD include the Incoming SI-BitString TLV in Echo Reply message.

Downstream Address and Downstream Interface Address

If the Address Type is 1, the Downstream Address MUST be set to IPV4 BFR-Prefix of downstream BFR and Downstream Interface Address is set to downstream interface address.

If the Address Type is 2, the Downstream Address MUST be set to IPV4 BFR-Prefix of downstream BFR and Downstream Interface Address is set to the index assigned by upstream BFR to the interface.

If the Address Type is 3, the Downstream Address MUST be set to IPV6 BFR-Prefix of downstream BFR and Downstream Interface Address is set to downstream interface address.

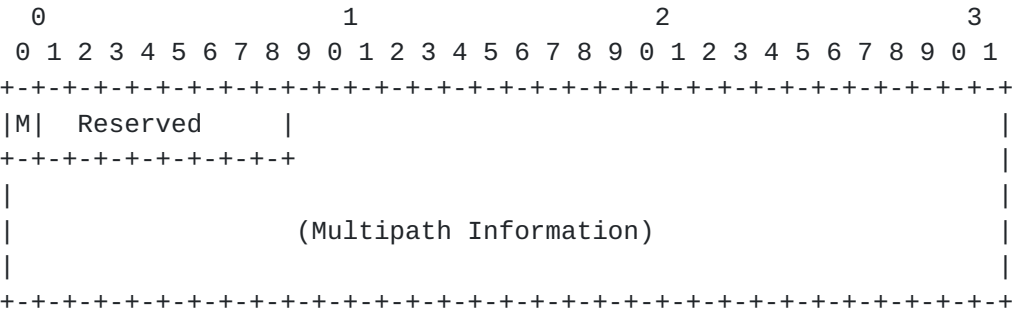
If the Address Type is 4, the Downstream Address MUST be set to IPV6 BFR-Prefix of downstream BFR and Downstream Interface Address is set to index assigned by upstream BFR to the interface.

3.3.4.1. Downstream Detailed Mapping Sub-TLVs

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

Sub-TLV Type	Value
1	Multipath Entropy Data
2	Egress BitString

3.3.4.1.1. Multipath Entropy Data



M Flag

This flag is set to 0 if all packets will be forwarded out through interface defined in Downstream Mapping TLV. When set to 1, Multipath Information will be defined with Bit masked Entropy data.

Multipath Information

Entropy Data encoded as defined in section x3.

3.3.4.1.2. Egress BitString

This Sub-TLV MAY be included by Responder BFR with the rewritten BitString in outgoing interface as defined in [section 6.1 of \[RFC8279\]](#)

```

0                               1                               2                               3
+-+-+-+-+
|   Set ID   | Sub-domain ID |BS Len|   Reserved   |
+-+-+-+-+
|               BitString  (first 32 bits)           ~
+-+-+-+-+
~
~
+-+-+-+-+
|               BitString  (last 32 bits)            |
+-+-+-+-+

```

3.3.5. 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. This TLV have the following format:

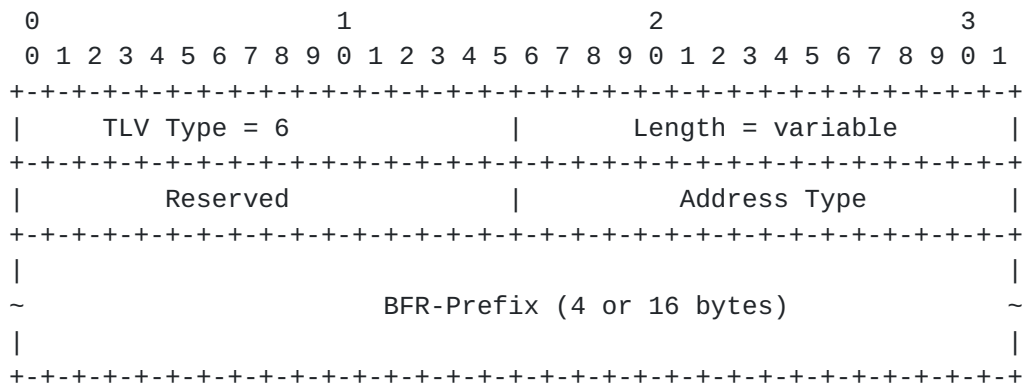
0										1										2										3																			
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9										
+-+-+...+-+-+ (40 bits)																																																	
										Type = 5																				Length																			
+-+-+...+-+-+ (40 bits)																																																	
										Reserved																				BFR-ID																			
+-+-+...+-+-+ (40 bits)																																																	

BFR-ID

The BFR-ID field carries the BFR-ID of replying BFER. This TLV MAY be included by Responding BFER in BIER Echo Reply packet.

3.3.6. 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 BFR-ID. This TLV have the following format:



Length

The Length field varies depending on the Address Type.

Address Type

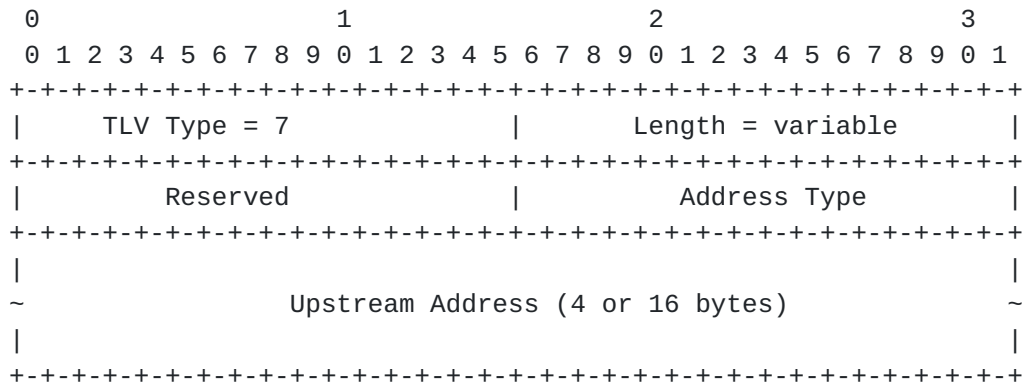
Set to 1 for IPv4 or 2 for IPv6.

BFR-Prefix

Carries the local BFR-Prefix of the replying BFR. This TLV MAY be included by Responding BFR in BIER Echo Reply packet.

3.3.7. Upstream Interface TLV

The Upstream Interface TLV will be included by the replying BFR in Echo Reply. This is used to identify the incoming interface and the BIER-MPLS label in the incoming Echo Request. This TLV have the following format:



Length

The Length field varies depending on the Address Type.

Address Type

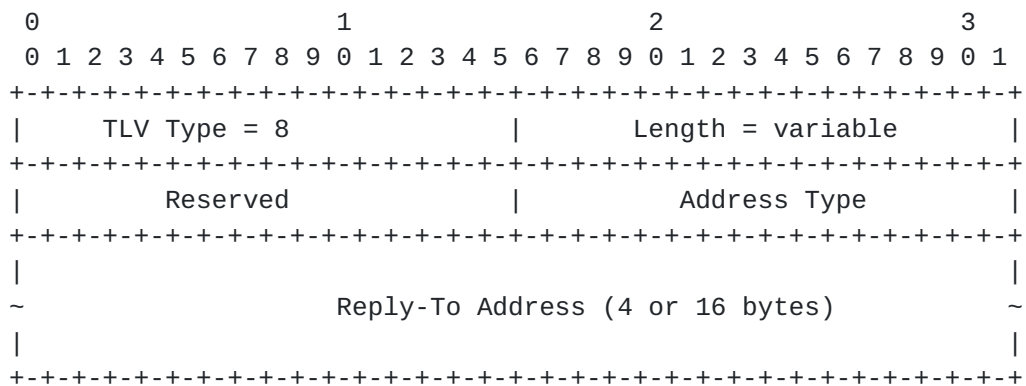
Set to 1 for IPv4 numbered, 2 for IPv4 Unnumbered 3 for IPv6 numbered or 4 for IPv6 Unnumbered.

Upstream Address

As defined in [Section 3.3.4](#)

3.3.8. 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 2. The IP address will be used to generate Echo Reply. This TLV have the following format:



Length

The Length field varies depending on the Address Type.

Address Type

Set to 1 for IPv4 or 2 for IPv6.

Reply-To Address

Set to any locally configured address to which the Echo reply should be sent to.

4. Procedures

This section describes aspects of Ping and traceroute operations. While this document explains the behavior when Reply mode is "Reply via BIER packet", the future version will be updated with details about the format when the reply mode is "Reply via IP/UDP packet".

4.1. BIER OAM Processing

A BIER 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.

4.2. Per BFER ECMP Discovery

As defined in [[RFC8279](#)], BIER follows unicast forwarding path and allows load balancing over ECMP paths between BFIR and BFER. BIER OAM MUST support ECMP path discovery between a BFIR and a given BFER and MUST support path validation and failure detection of any particular ECMP path between BFIR and BFER.

[RFC8296] proposes the BIER header with Entropy field that can be leveraged to exercise all ECMP paths. Initiator/BFIR will use traceroute message to query each hop about the Entropy information for each downstream paths. To avoid complexity, it is suggested that the ECMP query is performed per BFER by carrying required information in BIER OAM message.

Initiator MUST include Multipath Entropy Data Sub-TLV in Downstream Mapping TLV. It MUST also include the BFER in BitString TLV to which the Multipath query is performed.

Any transit BFR will reply back with Bit-masked Entropy for each downstream path as defined in [[RFC8029](#)]

4.3. Sending BIER Echo Request

Initiator MUST set the Message Type as 1 and Return Code as 0. The Proto field in OAM packet MUST be set to 0. The choice of Sender's Handle and Sequence number is a local matter to Initiator and SHOULD increment the Sequence number by 1 for every subsequent Echo Request. The QTF field is set to Initiator's local timestamp format and TimeStamp Sent field is set to the time that the Echo Request is sent.

Initiator MUST include Original SI-BitString TLV. Initiator MUST NOT include more than one Original SI-BitString TLV. Initiator infers the Set Identifier value and Sub-domain ID value from the respective BitString that will be included in BIER header of the packet and includes the values in "SI" and Sub-Domain ID fields respectively.

In Ping mode, Initiator MAY include Target SI-BitString TLV to control the responding BFER(s) by listing all the BFERs from which the Initiator expects a response. In trace route mode, Initiator MAY include Target SI-Bitstring TLV to control the path trace towards any specific BFER or set of BFERs. 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 unset the respective BFR-id from Target SI-BitString for any subsequent Echo Request.

When the Reply mode is set to 2, Initiator MUST include Reply-To TLV ([section 3.3.8](#)) in the Echo Request. Initiator MUST also listen to the UDP port defined in this TLV and process any segment received with destination port as value defined in the TLV and sent to control plane for BIER OAM payload processing.

Initiator MAY include Downstream Mapping TLV ([section 3.3.4](#)) 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 id.

Initiator MUST encapsulate the OAM packet with BIER header and MUST set the Proto as 5 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.

4.4. Receiving BIER Echo Request

Sending a BIER OAM Echo Request to control plane for payload processing is triggered as mentioned in [section 4.1](#).

Any BFR on receiving Echo Request MUST send Echo Reply with Return Code set to "Malformed Echo Request received", if the packet fails sanity check. If the packet sanity check is fine, it SHOULD initiate the below set of variables:

Reply-Flag

This flag is initially set to 1.

Interface-I

The incoming interface on which the Echo Request was received. This may be used to validate the DDMAP info and to populate the Upstream Interface TLV.

BIER-Label-L

The BIER-MPLS Label received as the top label on received Echo Request. This may be used to validate if the packet is traversing the desired Set Identifier and sub-domain path.

Header-H

The BIER header from the received Echo Request. This may be used to validate the DDMAP info and to populate the Incoming SI-BitString TLV. In addition, it can be used to perform Entropy calculation considering different field in header and reply via Multipath Entropy Data Sub-TLV.

BFR MUST initialize Best-return-code variable to null value.

BFR will populate Interface-I with interface over which the Echo Request is received, top label in the MPLS stack of the received Echo Request to BIER-Label-L, and the BIER header to BIER-Header. 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, reset Reply-Flag=0 and go to [section 4.5](#). Else:

- o If the BIER-Label-L does not correspond to the local label assigned for {sub-domain, BitStringLength, SI} in Original SI-BitString TLV, Set the Best-return-code to "Set-Identifier Mismatch" and Go to [section 4.5](#).

- * /* This detects any BIER-Label to {sub-domain, BitStringLen, SI} synchronization problem in the upstream BFR causing any unintended packet leak between sub-domains */
- o Set the Best-return-code to "One or more of the TLVs was not understood", if any of the TLVs in Echo Request message is not understood. Go to [section 4.5](#).
- o If the BitString in Header-H does not match the BitString in Egress BitString Sub-TLV of DDMAP TLV, set the Best-return-code to "DDMAP Mismatch" and go to [section 4.5](#). When there are more than one DDMAP TLV in the received Request packet, the Downstream Address and Downstream Interface Address should be matched with Interface-I to identify the right DDMAP TLV and then perform the BitString match.
 - * /* This detects any deviation between in BIER control plane and BIER forwarding plane in the previous hop that may result in loop or packet duplication. */
- o Set the Best-return-code to "Invalid Multipath Info Request", when the DDMAP TLV carries Multipath Entropy Data Sub-TLV and if the Target SI-BitString TLV in the received Echo Request carries more than 1 BFER id. Go to [section 4.5](#). Else, list the ECMP downstream neighbors to reach BFR-id in Target SI-BitString TLV, calculate the Entropy considering the BitString in Header-H and Multipath Entropy Data Sub-TLV from received Echo Request. Store the Data for each Downstream interface in temporary variable. Set the Best-return-code to 5 (Packet-Forward-Success) and goto [Section 4.5](#)
 - * /* This instructs to calculate the Entropy Data for each downstream interface to reach the BFER in Target SI-BitString TLV by considering the incoming BitString and Entropy Data.*/
- o Set the Best-return-code to "Replying router is the only BFER in BIER header Bitstring", and go to [section 4.5](#) if the responder is BFER and there is no more bits in BIER header Bitstring left for forwarding.
- o Set the Best-return-code to "Replying router is one of the BFER in BIER header Bitstring", and include Downstream Mapping TLV, if the responder is BFER and there are more bits in BitString left for forwarding. In addition, include the Multipath information as defined in [Section 4.2](#) if the received Echo Request carries Multipath Entropy Data Sub-TLV. Go to [section 4.5](#).

- o Set the Best-return-code to "No matching entry in forwarding table", if the forwarding lookup defined in [section 6.5 of \[RFC8279\]](#) does not match any entry for the received BitString in BIER header.
 - * /* This detects any missing BFR-id in the BIER forwarding table. It could be noted that it is difficult to detect such missing BFR-id while sending the Request with more than one BFR-id in BitString and so may need to just include the BFER id that is not responding to detect such failure.*/
- o Set the Best-return-code to "Packet-Forward-Success", and include Downstream Mapping TLV. Go to [section 4.5](#)

4.5. Sending Echo Reply

If Reply-Flag=0; BFR MUST release the variables and MUST not send any response to the Initiator. If Reply-Flag=1, proceed as below:

The Responder BFR SHOULD include the BitString from Header-H to Incoming SI-BitString TLV and include the Set ID, Sub-domain ID and BS Len corresponding to BIER-Label-L. Responder BFR SHOULD include the Upstream Interface TLV and populate the address from Interface-I.

When the Best-return-code is "Replying BFR is one of the BFER in header Bitstring", it MUST include Responder BFER TLV.

If the received Echo Request had DDMAP with Multipath Entropy Data Sub-TLV, Responder BFR MUST include DDMAP as defined in [Section 3.3.4](#) 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 DDMAP TLV.

If the received Echo Request had DDMAP without Multipath Entropy Data Sub-TLV, Responder BFR MUST include DDMAP as defined in [Section 3.3.4](#) for each outgoing interface over which the packet will be replicated. For each outgoing interface, respective Egress BitString MUST be included in DDMAP TLV.

When the Best-return-code is "Replying BFR is the only BFER in header Bitstring", it MUST include Responder BFER TLV.

Responder MUST set the Message Type as 2 and Return Code as Best-return-code. The Proto field MUST be set to 0.

The Echo Reply can be sent either as BIER-encapsulated or IP/UDP encapsulated depending on the Reply mode in received Echo Request.

When the Reply mode in received Echo Request is set to 3, Responder appends BIER header listing the BitString with BFIR ID (from Header-H), set the Proto to 5 and set the BFIR as 0. 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.

4.6. Receiving Echo Reply

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.

If receiving Echo Reply have Downstream Mapping, Initiator SHOULD copy the same to subsequent Echo Request(s).

If one of the Echo Reply is received with Return Code as "Replying BFR is one of the BFER in header Bitstring", it SHOULD reset the BFR-id of the responder from Target SI-BisString TLV in subsequent Echo Request.

/* This helps avoiding any BFR that is both BFER and also transit BFR to continuously responding with Echo Reply.*/

5. IANA Considerations

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

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

5.1. Message Types, Reply Modes, Return Codes

This document request to assign the Message Types and Reply mode mentioned in [section 3.1](#), , Return code mentioned in [Section 3.2](#).

5.2. TLVs

The TLVs and Sub-TLVs defined in this document is not limited to Echo Request or Reply message types and is applicable for other message types. The TLVs and Sub-TLVs requested by this document for IANA consideration are the following:

Type	Sub-Type	Value Field
-----	-----	-----
1		Original SI-BitString
2		Target SI-BitString
3		Incoming SI-BitString
4		Downstream Mapping
4	1	Multipath Entropy Data
4	2	Egress BitString
5		Responder BFER
6		Responder BFR
7		Upstream Interface

6. Security Considerations

The security consideration for BIER Ping is similar to ICMP or LSP Ping. AS like ICMP or LSP ping, BFR may be exposed to Denial-of-service attacks and it is RECOMMENDED to regulate the BIER Ping packet flow to control plane. A rate limiter SHOULD be applied to avoid any attack.

As like ICMP or LSP Ping, a traceroute can be used to obtain network information. It is RECOMMENDED that the implementation check the integrity of BFIR of the Echo messages against any local secured list before processing the message further

7. Acknowledgement

The authors would like to thank Antoni Przygienda, Eric Rosen, Faisal Iqbal Jeffrey (Zhaohui) Zhang and Shell Nakash for thier review and comments.

8. Contributing Authors

TBD

9. References

9.1. Normative References

[I-D.ietf-bier-ospf-bier-extensions]
 Psenak, P., Kumar, N., Wijnands, I., Dolganow, A.,
 Przygienda, T., Zhang, Z., and S. Aldrin, "OSPF Extensions
 for BIER", [draft-ietf-bier-ospf-bier-extensions-10](#) (work
 in progress), December 2017.

- [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>>.
- [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>>.
- [RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", [RFC 8029](#), DOI 10.17487/RFC8029, March 2017, <<https://www.rfc-editor.org/info/rfc8029>>.
- [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

- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, [RFC 792](#), DOI 10.17487/RFC0792, September 1981, <<https://www.rfc-editor.org/info/rfc792>>.
- [RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", [BCP 161](#), [RFC 6291](#), DOI 10.17487/RFC6291, June 2011, <<https://www.rfc-editor.org/info/rfc6291>>.
- [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>>.

[RFC6425] Saxena, S., Ed., Swallow, G., Ali, Z., Farrel, A., Yasukawa, S., and T. Nadeau, "Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping", [RFC 6425](https://www.rfc-editor.org/info/rfc6425), DOI 10.17487/RFC6425, November 2011, <<https://www.rfc-editor.org/info/rfc6425>>.

Authors' Addresses

Nagendra Kumar
Cisco Systems, Inc.
7200 Kit Creek Road
Research Triangle Park, NC 27709
US

Email: naikumar@cisco.com

Carlos Pignataro
Cisco Systems, Inc.
7200 Kit Creek Road
Research Triangle Park, NC 27709-4987
US

Email: cpignata@cisco.com

Nobo Akiya
Big Switch Networks
Japan

Email: nobo.akiya.dev@gmail.com

Lianshu Zheng
Huawei Technologies
China

Email: vero.zheng@huawei.com

Mach Chen
Huawei Technologies

Email: mach.chen@huawei.com

Greg Mirsky
Ericsson

Email: gregory.mirsky@ericsson.com