

BIER WG
Internet-Draft
Intended status: Standards Track
Expires: 30 August 2025

Q. Xiong
ZTE Corporation
G. Mirsky
Ericsson
F. Hu
Individual
C. Liu
China Unicom
G. Mishra
Verizon Inc.
26 February 2025

BIER BFD
draft-ietf-bier-bfd-08

Abstract

Point to multipoint (P2MP) BFD is designed to verify multipoint connectivity. This document specifies the application of P2MP BFD in BIER 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 30 August 2025.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Conventions used in this document	3
2.1. Terminology	3
2.2. Requirements Language	3
3. BIER BFD Encapsulation	3
4. BIER BFD Session Bootstrapping	3
4.1. Bootstrapping a BIER BFD Session Using BIER Ping	4
4.2. BGP Bootstrapping	4
5. Discriminators and Packet Demultiplexing	5
6. Active Tail in BIER BFD	5
6.1. Unsolicited Head Notification Mode	5
7. Security Considerations	6
8. Acknowledgements	6
9. IANA Considerations	6
9.1. BIER OAM Message Type	6
9.2. BFD Discriminator TLV	7
9.3. BGP's BFD Discriminator Attribute Mode of P2MP Session with Active Tails	7
10. References	7
10.1. Normative References	7
Authors' Addresses	9

1. Introduction

Bit Index Explicit Replication (BIER) [RFC8279] provides the forwarding of multicast data packets through a multicast domain. It does so without requiring any explicit tree-building protocol and without requiring intermediate nodes to maintain any per-flow state.

[RFC8562] defines a method of using Bidirectional Forwarding Detection (BFD) to monitor and detect failures between the sender (head) and one or more receivers (tails) in multipoint or multicast networks. [RFC8563] describes active tail extensions to the BFD protocol for multipoint networks.

This document describes the procedures for using such mode of BFD protocol to monitor connectivity between a multipoint sender, Bit-Forwarding Ingress Router (BFIR), and a set of one or more multipoint receivers, Bit-Forwarding Egress Routers (BFERs). The BIER BFD only supports the unidirectional multicast. This document defines the use of P2MP BFD as per [RFC8562], and active tail as per [RFC8563] for BIER-specific domain.

2. Conventions used in this document

2.1. Terminology

This document uses the acronyms defined in [RFC8279] along with the following:

BFD: Bidirectional Forwarding Detection.

OAM: Operations, Administration, and Maintenance.

P2MP: Point to Multi-Point.

2.2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. BIER BFD Encapsulation

BIER BFD encapsulation uses the BIER OAM packet format defined in [I-D.ietf-bier-ping]. The value of the Message Type field MUST be set to BIER BFD (TBD1 will be assigned by IANA Section 9.1). BFD Control packet, defined in Section 4 [RFC5880], immediately follows the BIER OAM header. The operation of Multipoint BFD with the BFD Control Packet is described in [RFC8562].

4. BIER BFD Session Bootstrapping

As defined in [RFC8562], a BIER BFD session MAY be established to monitor the state of the multipoint path. The BIER BFD session could be created for each multipoint path and the set of BFERs over which the BFIR is requested to run BIER BFD. The BFIR, according to Section 5.7 of [RFC8562], MAY bootstrap the BFD session using a BIER OAM message (Section 4.1) or the control plane (Section 4.2). Either method MUST refer to the root of the multipoint path and the value of My Discriminator associated with the path to the set of BFERs. The

BIER BFD bootstrapping MUST be repeated when the value of My Discriminator is changed.

4.1. Bootstrapping a BIER BFD Session Using BIER Ping

The BIER OAM could be used for bootstrapping the BIER BFD session. The BFIR sends the BIER OAM Echo request message carrying a BFD discriminator TLV which immediately follows the Target SI-Bitstring TLV (Section 3.3.2 of [I-D.ietf-bier-ping]).

The Target SI-Bitstring TLV MUST be used to carry the set of BFER information (including Sub-domain-id, Set ID, BS Len, Bitstring) for the purpose of the session establishment.

Figure 1 displays the format of BFD Discriminator TLV.




Figure 1: BFD Discriminator TLV

where:

- * Type indicates BFD Discriminator TLV. The value (TBD2) is to be allocated by IANA (Section 9.2).
- * Length MUST be set to 4.
- * My Discriminator - four-octet long field. The value is the local discriminator generated by BFIR for this session. This discriminator MUST be used as the My Discriminator field in the BIER BFD Control packets sent by the BFIR.

4.2. BGP Bootstrapping

[RFC9026] describes the applicability of [RFC8562] in the detection of a Multicast Virtual Private Network. The new BGP Attribute, BFD Discriminator, can be used to bootstrap a p2mp BFD session according to [RFC8562]. To bootstrap a p2mp BFD session with active tails according to [RFC8563] IANA is requested to allocate a new value (TBD5) for the BFD Mode (Section 9.3).

5. Discriminators and Packet Demultiplexing

As defined in [RFC8562], the BFIR sends BFD Control packets to the set of BFERs over the multipoint path via the BIER BFD session with My Discriminator is set to the value assigned by the BFIR, and the value of the Your Discriminator is set to zero. The set of BFERs MUST demultiplex BFD packets based on a combination of the source address, My Discriminator value. The source address is BFIR-id and BIER MPLS Label (MPLS network) or BFIR-id and BIFT-id (Non-MPLS network) for BIER BFD. My Discriminator value is advertised in BIER BFD bootstrapping using one of the options described in Section 4.

6. Active Tail in BIER BFD

[RFC8563] defined an extension for Multipoint BFD, which allows the head to discover the state of a multicast distribution tree for any sub-set of tails. For BIER BFD in the active tail mode, the BFIR may learn the state and connectivity of the BFERs by allowing the BFERs to notify the BFIR. As per [RFC8563] provides detailed information on how the BFIR can use multipoint Poll sequence message or a combination of multicast and unicast Poll sequence messages to determine the state of the multicast tree. Also, [RFC8563] describes that a BFER can transmit an unsolicited unicast Poll sequence message to the BFIR (note that a unicast message must be sent over a path which is disjoint from the multicast distribution tree).

6.1. Unsolicited Head Notification Mode

[I-D.ietf-mpls-p2mp-bfd] provides detailed information on using the unsolicited notification method for P2MP MPLS LSP, which is also applicable to BIER.

In Section 5.2.1 [RFC8563] is noted that "the tail sends unsolicited BFD packets in response to the detection of a multipoint path failure" but without the specifics on the information in the packet and frequency of transmissions. This document defines the procedure of the active tail with unsolicited notifications for BIER as specified below.

Upon detecting the failure, a BFER sends a BFD control packet with the following settings:

- * the Poll (P) bit is set;
- * the Status (Sta) field is set to Down value;
- * the Diagnostic (Diag) field is set to Control Detection Time Expired value;

- * the value of the Your Discriminator field is set to the value the BFER has been used to demultiplex that BFD multipoint session;
- * BFD Control packet is encapsulated in IP/UDP with the destination IP address of the BFIR and the UDP destination port number set to 4784 per [RFC5883]
- * the BFD Control packets are transmitted at the rate of one per second until either the BFER receives valid for this BFD session control packet with the Final (F) bit is set from the BFIR or the defect condition clears.

To improve the likelihood of notifying the BFIR of the failure, the BFER SHOULD transmit three BFD Control packets defined above in with pseudo-random intervals between packets within a one-second interval.

A BFIR that has received the BFD Control packet, as described above, sends the unicast IP/UDP encapsulated BFD control packet with the Final (F) bit set to the BFER.

7. Security Considerations

For BIER OAM packet processing security considerations, see [I-D.ietf-bier-ping].

For general multipoint BFD security considerations, see [RFC8562].

A single failure could affect a significant number of BFERs, thus causing a spike in the number of BFD Control packets with notifications, as defined in Section 6.1. To mitigate the overloading of the control plane, an implementation MUST control the number of BFD Control packets passed to the control plane for processing.

8. Acknowledgements

The authors would like to thank the comments and suggestions from Sandy Zhang, Jeffrey (Zhaohui) Zhang, Donald Eastlake 3rd, Reshad Rahman, and Les Ginsberg.

9. IANA Considerations

9.1. BIER OAM Message Type

IANA is requested to assign a new type from the BIER OAM Message Type registry as follows:

Value	Description	Reference
TBD1	BIER BFD	[this document]

Table 1

9.2. BFD Discriminator TLV

IANA is requested to assign a new type from the BIER OAM TLV registry as follows:

Value	Description	Reference
TBD2	BFD discriminator TLV	[this document]

Table 2

9.3. BGP's BFD Discriminator Attribute Mode of P2MP Session with Active Tails

IANA is requested to assign a new value from "BFD Mode" subregistry (defined in [RFC9026]) in the "Border Gateway Protocol (BGP) Parameters" registry according to Table 3:

Value	Description	Reference
TBD5	P2MP BFD Session with Active Tails	[this document]

Table 3

10. References

10.1. Normative References

[I-D.ietf-bier-ping]
Nainar, N. K., Pignataro, C., Chen, M., and G. Mirsky,
"BIER Ping and Trace", Work in Progress, Internet-Draft,
draft-ietf-bier-ping-15, 8 November 2024,
<<https://datatracker.ietf.org/doc/html/draft-ietf-bier-ping-15>>.

`[I-D.ietf-mpls-p2mp-bfd]`

Mirsky, G., Mishra, G. S., and D. E. Eastlake,
"Bidirectional Forwarding Detection (BFD) for Multipoint
Networks over Point-to-Multi-Point MPLS Label Switched
Path (LSP)", Work in Progress, Internet-Draft, draft-ietf-
mpls-p2mp-bfd-11, 19 February 2025,
<[https://datatracker.ietf.org/doc/html/draft-ietf-mpls-
p2mp-bfd-11](https://datatracker.ietf.org/doc/html/draft-ietf-mpls-p2mp-bfd-11)>.

- [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>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection
(BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010,
<<https://www.rfc-editor.org/info/rfc5880>>.
- [RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection
(BFD) for Multihop Paths", RFC 5883, DOI 10.17487/RFC5883,
June 2010, <<https://www.rfc-editor.org/info/rfc5883>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC
2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174,
May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [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>>.
- [RFC8562] Katz, D., Ward, D., Pallagatti, S., Ed., and G. Mirsky,
Ed., "Bidirectional Forwarding Detection (BFD) for
Multipoint Networks", RFC 8562, DOI 10.17487/RFC8562,
April 2019, <<https://www.rfc-editor.org/info/rfc8562>>.
- [RFC8563] Katz, D., Ward, D., Pallagatti, S., Ed., and G. Mirsky,
Ed., "Bidirectional Forwarding Detection (BFD) Multipoint
Active Tails", RFC 8563, DOI 10.17487/RFC8563, April 2019,
<<https://www.rfc-editor.org/info/rfc8563>>.
- [RFC9026] Morin, T., Ed., Kebler, R., Ed., and G. Mirsky, Ed.,
"Multicast VPN Fast Upstream Failover", RFC 9026,
DOI 10.17487/RFC9026, April 2021,
<<https://www.rfc-editor.org/info/rfc9026>>.

Authors' Addresses

Quan Xiong
ZTE Corporation
China
Email: xiong.quan@zte.com.cn

Greg Mirsky
Ericsson
Email: gregimirsky@gmail.com

Fangwei Hu
Individual
Email: hufwei@gmail.com

Chang Liu
China Unicom
No.9 Shouti Nanlu
Beijing
100048
China
Phone: +86-010-68799999-7294
Email: liuc131@chinaunicom.cn

Gyan Mishra
Verizon Inc.
Email: gyan.s.mishra@verizon.com