

BIER WG
Internet-Draft
Intended status: Standards Track
Expires: January 4, 2020

Quan Xiong
Greg Mirsky
ZTE Corporation
Fangwei Hu
Individual
Chang Liu
China Unicom
July 3, 2019

BIER BFD
draft-hu-bier-bfd-04.txt

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 January 4, 2020.

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
2.1. Terminology	3
2.2. Requirements Language	3
3. BIER BFD Encapsulation	3
4. BIER BFD Session Bootstrapping	3
4.1. BIER OAM Bootstrapping	3
4.2. IGP protocol Bootstrapping	4
4.2.1. IS-IS extension for BIER BFD	4
4.2.2. OSPF extension for BIER BFD	5
5. Discriminators and Packet Demultiplexing	6
6. Active Tail in BIER BFD	6
7. Security Considerations	7
8. Acknowledgements	7
9. IANA Considerations	7
10. References	7
10.1. Normative References	7
10.2. Informative References	8
Authors' Addresses	8

1. Introduction

Bit Index Explicit Replication (BIER) [RFC8279] provides 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 unicast 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 Msg Type field MUST be set to BIER BFD (TBD1 by IANA). 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], 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 wishes to run BIER BFD. The BFIR MUST advertise the BFD Discriminator along with the corresponding multipoint path to the set of BFERs. Bootstrapping a BIER BFD session MAY use BIER OAM message section 4.1 or the control plane section 4.2.

The BIER BFD bootstrapping MUST be repeated when the value of this discriminator being changed.

4.1. BIER OAM Bootstrapping

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 [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 session establishment.

The BFD discriminator TLV is a new TLV for BIER OAM TLV with the type (TBD2 by IANA) and the length of 4. The value contains the 4-byte local discriminator generated by BFIR for this session. This discriminator MUST subsequently be used as the My Discriminator field in the BIER BFD session packets sent by BFIR. The format is as follows.




Figure 1: BFD discriminator TLV

4.2. IGP protocol Bootstrapping

An alternative option to bootstrap the BIER BFD is to advertise the BFD information in control plane. This document defines a new BIER BFD Sub-sub-TLV carried in IS-IS and OSPF protocol.

The BFIR generates the My Discriminator value for each multicast flow and advertises it to the expecting BFERs which is indicated by the Bitstring which is carried in BIER BFD sub-sub-TLV. The corresponding BFERs SHOULD store the My Discriminator value for packet Demultiplexing.

4.2.1. IS-IS extension for BIER BFD

The new BIER BFD Sub-sub-TLV is carried within the BIER Info sub-TLV defined in [RFC8401]. The format is as follows.




Figure 2: BIER BFD Sub-sub-TLV for IS-IS extension

Type: TBD3 by IANA.

Length: Length of the BIER BFD Sub-sub-TLV for IS-IS extension, in bytes.

My Discriminator: A unique, nonzero discriminator value generated by BFIR for each multipoint path.

The BitString field carries the set of BFR-IDs of BFER(s) that the BFIR expects to establish BIER BFD session.

The BIFT-id represents a particular Bit Index Forwarding Table (BIFT) as per [RFC8279].

4.2.2. OSPF extension for BIER BFD

The new BIER BFD Sub-TLV is a sub-TLV of the BIER Sub-TLV defined in [RFC8444]. The format is as follows.




Figure 3: BIER BFD Sub-TLV for OSPF extension

Type: TBD4 by IANA.

Length: Length of the BIER BFD Sub-TLV for OSPF extension, in bytes.

Other fields in BIER BFD Sub-TLV is the same with section 4.2.1.

5. Discriminators and Packet Demultiplexing

As defined in [RFC8562], the BFIR sends BFD Control packets over the multipoint path via the BIER BFD session with My Discriminator set to the value assigned by the BFIR and the value of the Your Discriminator 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. The My Discriminator value is advertised in BIER BFD bootstrapping using one of 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 active tail mode, the BFIR may learn the state and connectivity of the BFERs. As per [RFC8563], the BFIR uses a combination of multicast Poll sequence messages and unicast Poll messages. The unicast messages must be sent over the path which is disjoint from the multicast distribution tree.

7. Security Considerations

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

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

No additional security issues are raised in this document beyond those that exist in the referenced BFD documents.

8. Acknowledgements

Authors would like to thank the comments and suggestions from Sandy Zhang, Jeffrey (Zhaohui) Zhang, Donald Eastlake 3rd.

9. IANA Considerations

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

Value	Description	Reference
TBD1	BIER BFD	[this document]
TBD2	BFD discriminator TLV	[this document]
TBD3	BIER BFD Sub-sub-TLV for IS-IS	[this document]
TBD4	BIER BFD Sub-TLV for OSPF	[this document]

Table 1

10. References

10.1. 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-05 (work in progress), April 2019.
- [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>>.

- [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>>.
- [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>>.
- [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>>.

10.2. Informative References

- [ISO9577] ISO/IEC TR 9577:1999,, "International Organization for Standardization "Information technology - Telecommunications and Information exchange between systems - Protocol identification in the network layer"", 1999.

Authors' Addresses

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

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

Greg Mirsky
ZTE Corporation
USA

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