BIER Workgroup
Internet Draft
Intended status: Standard Track

H. Bidgoli
J. Kotalwar
Nokia
Z.Zhang
Juniper Networks
Eddie Leyton
Vrizon
Mankamana Mishra
I. Wijanands
Cisco System, Inc.

Expires: May 6, 2020 November 3, 2019

M-LDP Signaling Through BIER Core
draft-hb-bier-mldp-signaling-over-bier-01

Abstract

Bit Index Explicit Replication (BIER) is an architecture that provides multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain multicast related per-flow state. Neither does BIER require an 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. Such 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 the according set of bits switched on in BIER packet header.

This document describes the procedure needed for mLDP tunnels to be signaled over and stitched through a BIER core, allowing LDP routers to run traditional Multipoint LDP services through a BIER core.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Table of Contents

1. Introduction 3
2. Conventions used in this document 3
2.1. Definitions 3
3. mLDP Signaling Through BIER domain 4
3.1. Ingress BBR procedure 5
3.1.1. Automatic tLDP session creation 5
3.1.1. ECMP Method on IBBR 6
3.2. Egress BBR procedure method 6
3.2.1. IBBR procedure upon arriving upstream assigned label . 6
4. Datapath Forwarding 7
4.1. Datapath traffic flow 7
5. Recursive FEC 7
6. IANA Considerations 7
7. Security Considerations 7
8. References 7
8.1. Normative References 7
8.2. Informative References 8
7. Acknowledgments 8

Authors' Addresses 8

1. Introduction

Some operators that are using mLDP P2MP LSPs for their multicast transport would like to deploy BIER technology in some segment of their network. This draft explains a method to signal mLDP services and stitch the mLDP datapath labels through a BIER domain, with minimal disruption and operational impact to the mLDP domain.

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 RFC 2119 [RFC2119].

2.1. Definitions

Some of the terminology specified in [I-D.draft-ietf-bier-architecture-05] is replicated here and extended by necessary definitions:

BIER:

Bit Index Explicit Replication (The overall architecture of forwarding multicast using a Bit Position).

BFR:

Bit Forwarding Router (A router that participates in Bit Index Multipoint Forwarding). A BFR is identified by a unique BFR-prefix in a BIER domain.

BFIR:

Bit Forwarding Ingress Router (The ingress border router that inserts the Bit Map into the packet). Each BFIR must have a valid BFR-id assigned. BFIR is term used for dataplain packet forwarding.

BFER:

Bit Forwarding Egress Router. A router that participates in Bit Index Forwarding as leaf. Each BFER must be a BFR. Each BFER must have a valid BFR-id assigned. BFIR is term used for dataplain packet forwarding.

BBR:

BIER Boundary router. The router between the LDP domain and BIER domain.

IBBR:

Ingress BIER Boundary Router. The ingress router from signaling point of view. It maintains mLDP adjacency toward the LDP domain and determines if the mLDP FEC needs to be signaled across the BIER domain via targeted ldp.

EBBR:

Egress BIER Boundary Router. The egress router in BIER domain from signaling point of view. It terminates the targeted ldp signaling through BIER domain. It also keeps track of all IBBRs that are part of this p2mp tree

BFT:

Bit Forwarding Tree used to reach all BFERs in a domain.

BIFT:

Bit Index Forwarding Table.

BIER sub-domain:

A further distinction within a BIER domain identified by its unique sub-domain identifier. A BIER sub-domain can support multiple BitString Lengths.

BFR-id:

An optional, unique identifier for a BFR within a BIER sub-domain.

3. mLDP Signaling Through BIER domain

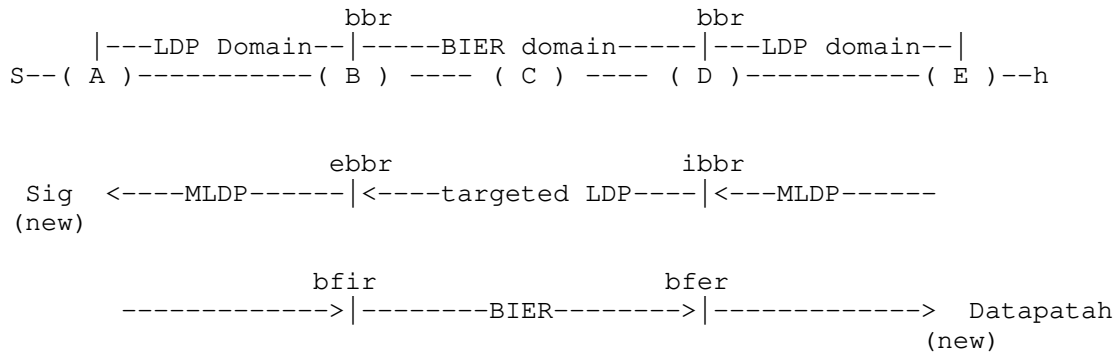

Figure 1: bier boundary router

As per figure 1, point-to-multipoint and multipoint-to-multipoint LSPs established via mLDP [RFC6388] can be signaled through a bier domain via targeted LDP sessions. This procedure is explained in [RFC7060] (Using LDP Multipoint Extension on Targeted LDP Sessions).

This documents provides some details and defines some needed procedures.

3.1. Ingress BBR procedure

The Ingress BBR (IBBR) is connected to the mLDP on one side and a bier domain on the other side. To connect the LDP domains via BIER domain IBBR needs to establish a targeted LDP session with EBBR closest to the root of the P2mp or mp2mp LSP. To do so IBBR will follow procedures in [RFC7060] in particular the section "6. targeted mLDP with Multicast Tunneling".

The target LDP session can be established manually via configuration or via automated mechanism.

3.1.1. Automatic tLDP session creation

A tLDP session can be generated automatically from every IBBR to EBBR. As an example when a mLDP FEC arrives on the IBBR, it can automatically start a tLDP Session with the EBBR. In this case both IBBR and EBBR should be in auto-discovery mode and react to the arriving FEC or tLDP Signaling packets (i.e. targeted hellos, keep-alives etc...).

The Root node address in the mLDP FEC can be used to find the EBBR. To identify the EBBR same procedures as [RFC7060} section 2.1 can be

used or the procedures as explained in the [draft-ietf-bier-pim-signaling] appendix A. After fining the IBBR the tLDP session can be initiated from the IBBR to EBBR.

3.1.1. ECMP Method on IBBR

If IBBR finds multiple equal cost EBBRs on the path to the Root, it can use a vendor specific algorithm to choose between the EBBRs. These algorithms are beyond the scope of this draft. As an example the IBBR can use the smallest EBBR IP address to establish its mLDP signaling to.

3.2. Egress BBR procedure method

The Egress BBR (EBBR) is connected to the mLDP domain which the root of the P2MP or MP2MP LSP resides on. The EBBR should accept the tLDP session generated form IBBR. It should assign a unique "upstream assigned label" for each arriving FEC generated by IBBRs.

The EBBR should follow the [RFC7060] procedures with following modifications:

- The label assigned by EBBR cannot be Implicit Null. This is to ensure that identity of each p2mp and/or mp2mp tunnel in BIER domain is uniquely distinguished.

- The label can be assigned from a domain-wide Common Block (DCB) [I-D.zzhang-bess-mvpn-evpn-aggregation-label], as well as upstream assigned.

- The Interface ID TLV [RFC6389]includes a new BIER sub-domain sub-tlv (type TBD)

The EBBR will also generate a new label and FEC toward the ROOT on the mLDP domain. The EBBR Should stitch this generate label with the "upstream assigned label" to complete the p2MP or MP2MP LSP. This stitch point should be stored on the datapath (ILM) table for packet forwarding.

With same token the EBBR should track all the arriving FECs and the IBBRs that are generating these FECs. EBBR will use this information to build the bier header for each set of common FEC arriving from the IBBRs.

3.2.1. IBBR procedure upon arriving upstream assigned label

Upon receiving the "upstream assigned label", IBBR should create its own stitching instruction between the "upstream assigned label" and

the down stream label that was signaled to it. IBBR should download these instructions to the datapath.

4. Datapath Forwarding

4.1. Datapath traffic flow

On BFIR when the MPLS label for P2MP/MP2MP LSP arrives a lookup in ILM table is done and the label is swapped with tLDP upstream assigned label. The BFIR will note all the BFERs that are interested in specific p2mp/mp2mp LSP (as per section 3.2). BFIR will put the corresponding BIER header with bit index set for all IBBRs interested in this P2MP LSP. BFIR will set the BIERHeader.Proto = MPLS and will forward the BIER packet into BIER domain.

In the BIER domain normal BIER forwarding procedure will be done, as per [RFC 8279]

The IBBRs will receive the BIER packet, will look at the protocol of BIER header (MPLS). BFER will remove the BIER header and will do a lookup in the ILM table for the upstream assigned label and perform its corresponding action.

It should be noted that these procedures are valid if BFIR is the ILER and/or BFER is the ELER as per [RFC 7060]

5. Recursive FEC

The above procedures also will work with a mLDP recursive FEC. The root used to determine the EBBR is the outer root of the FEC. The entire recursive FEC needs to be preserve when it is forwarded via tLDP and the label request.

6. IANA Considerations

This document contains no actions for IANA.

7. Security Considerations

TBD

8. References

8.1. Normative References

[BIER_ARCH] Wijnands, IJ., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication",

internet-draft draft-ietf-bier-architecture-08, October 2016.

8.2. Informative References

[BIER_MVPN] Rosen, E., Ed., Sivakumar, M., Wijnands, IJ., Aldrin, S., Dolganow, A., and T. Przygienda, "Multicast VPN Using Bier", internet-draft draft-ietf-bier-mvpn-08, January 2017.

[ISIS_BIER_EXTENSIONS] Ginsberg, L., Przygienda, T., Aldrin, S., and Z. Zhang, "BIER Support via ISIS", internet-draft draft-ietf-bier-isis-extensions-06.txt, March 2017.

[OSPF_BIER_EXTENSIONS] Psenak, P., Kumar, N., Wijnands, IJ., Dolganow, A., Przygienda, T., Zhang, Z., and S. Aldrin, "OSPF Extensions for Bit Index Explicit Replication", internet-draft draft-ietf-ospf-bier-extensions-09.txt, March 2017.

7. Acknowledgments Authors would like to acknowledge Jingrong Xie for his comments and help on this draft.

Authors' Addresses

Hooman Bidgoli (editor)
Nokia
600 March Rd.
Ottawa, Ontario K2K 2E6
Canada

Email: hooman.bidgoli@nokia.com

Jayant Kotalwar
Nokia
380 N Bernardo Ave,
Mountain View, CA 94043
US

Email: jayant.kotalwar@nokia.com

Zhaohui Zhang
Juniper Networks

EMail: zzhang@juniper.net

IJsbrand Wijnands
Cisco Systems

EMail: ice@cisco.com

Eddie Leyton
Vrizon

Email: Edward.leyton@verizonwireless.com

Mankamana Mishra
Cisco System
821 alder drive
Milpitas California
USA

Email: mankamis@cisco.com