Authors: H. Bidgoli, Ed.   J. Kotalwar   I. Wijnands
         Nokia   Nokia   Cisco System
         M. Mishra   Z. Zhang   E. Leyton
         Cisco System   Juniper Networks   Verizon

M-LDP Signaling Through BIER Core

Abstract

   Consider an end to end Multipoint LDP (mLDP) network, where it is
   desirable to deploy BIER in portion of this network. It might be
   desirable to deploy BIER with minimum disruption to the mLDP network
   or redesign of the network.

   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 mLDP 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). 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 16 May 2022.

Copyright Notice

Table of Contents

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 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 have a valid BFR-id assigned. BFER 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

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,
   all BFERs and BFIRs need to be assigned a BFR-id.

3. mLDP Signaling Through BIER domain

```
                    bbr                       bbr
     |---LDP Domain--|-----BIER domain-----|---LDP domain--|
 S--( A )-----------( B ) ---- ( C ) ---- ( D )-----------( E )--h


                    ebbr                      ibbr
 Sig  <----MLDP------|<----targeted LDP----|<---MLDP------
 (new)

                    bfir                      bfer
     ------------->|--------BIER-------->|-------------> Datapatah
                                                             (new)
```

              Figure 1: BIER boundry 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 details and defines some needed procedures.

    .

3.1. Ingress BBR procedure

   The Ingress BBR (IBBR) is connected to the mLDP domain on downstream
   and a bier domain on the upstream. 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

   tLDP session can be signaled automatically from every IBBR to the
   appropriate EBBR. When mLDP FEC arrives to IBBR from LDP domain,
   IBBR can automatically start a tLDP Session to the EBBR closest to
   the Root node. Both IBBR and EBBR should be in auto-discovery mode

and react to the arriving tLDP Signaling packets (i.e. targeted hellos, keep- alives etc...) to establish the session automatically.

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.

3.1.2. 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

The Egress BBR (EBBR) is connected to the upstream mLDP domain. 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)
   [draft-zzhang-bess-mvpn-evpn-aggregation-label]

  *The Interface ID TLV, as per [RFC6389] should 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 LDP domain. The EBBR Should stitch this generate label with the "upstream assigned label" to complete the P2MP or MP2MP LSP.

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 for 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 signaled label.

4. Datapath Forwarding

4.1. Datapath traffic flow

On BFIR when the MPLS label for P2MP/MP2MP LSP arrives from
upstream, 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 stream. 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 [RFC8279]

The BFERs 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 also valid if BFIR is
the ILER and/or BFER is the ELER as per [RFC7060]

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 Consideration

adf

   1. A new BIER sub-domain sub- tlv for the interface ID TLV to be
      assigned by IANA

7. Security Considerations

TBD

8. Acknowledgments

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

9. References

9.1. Normative References

[draft-ietf-bier-pim-signaling]
                            "H, Bidgoli, F. Xu, J. Kotalwar, I.
          Wijnands, M. Mishra, Z. Zhang", 29 July 2020.

[draft-zzhang-bess-mvpn-evpn-aggregation-label] "Z. Zhang, E. Rosen,
          W. Lin, Z. Li, I. Wijnands, "MVPN/EVPN Tunnel Aggregation
          with Common Labels"", 27 April 2018.

[RFC6389]   "R Aggarwal, JL. Le Roux, "MPLS Upstream Label Assignment
          for LDP"", November 2011.

[RFC7060]   "M. Napierala, E. Rosen, I. Wijnands", November 2013.

[RFC8279]   "I. Wijnands, E. Rosen, A. ADolganow, T. Prygienda, S.
          Aldrin", November 2017.

9.2. Informative References

[RFC8401]   "Ginsberg, L., Przygienda, T., Aldrin, S., and Z. Zhang,
          "BIER Support via ISIS"", June 2018.

[RFC8444]   "Psenak, P., Kumar, N., Wijnands, IJ., Dolganow, A.,
          Przygienda, T., Zhang, Z., and S. Aldrin, "OSPF
          Extensions for Bit Index Explicit Replication"", June
          2018.

[RFC8556]   "Rosen, E., Ed., Sivakumar, M., Wijnands, IJ., Aldrin,
          S.,Dolganow, A., and T. Przygienda, "Multicast VPN Using
          Index Explicit Replication (BIER)", April 2018.

Authors' Addresses

Hooman Bidgoli (editor)
Nokia
Ottawa
Canada

Email: hooman.bidgoli@nokia.com

Jayant Kotalwar
Nokia
Montain View,
United States of America

Email: jayant.kotalwar@nokia.com

IJsbrand Wijnands
Cisco System
Diegem
Belgium

      Email: ice@cisco.com

   Mankamana Mishra
   Cisco System
   Milpitas,
   United States of America

      Email: mankamis@cisco.com

   Zhaohui Zhang
   Juniper Networks
   Boston,
   United States of America

      Email: zzhang@juniper.com

   Eddie Leyton
   Verizon

      Email: Edward.leyton@verizonwireless.com