

BIER Workgroup
Internet Draft
Intended status: Standard Track

H. Bidgoli, Ed.
A. Dolganow
J. Kotalwar
Nokia

Expires: January 3, 2019

July 2, 2018

M-LDP Signaling Through BIER Core
draft-hj-bier-mldp-signaling-00

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 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.

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." The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on October 8, 2017.

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 (<http://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. Definitions	3
3. M-LDP Signaling Through BIER domain	5
3.1. Ingress BBR procedure	5
3.2. Assigning the BIER sub-domain Tree Label	6
3.2.1. Method 1: IBBR as extraction point to SDN controller	6
3.2.2. Method 2, EBBR assigning the BTL	7
3.2.2.1. BIER packet construction at IBBR for Method 2	7
3.2.2.2. Signaling mLDP through the BIER domain procedure	8
3.3. SDN Controller procedure for updating BBRs with BTL	8
3.4. BGP procedure for signaling BTL for method 2	9
3.5. EBBR procedure method	9
3.6. Label release	10
3.6.1 Label release in method 1	10
3.6.2 Label release in method 2	10
4. Datapath Forwarding	10
4.1. BFIR tracking of FEC	10
4.2. Datapath traffic flow	10
5. Recursive FEC	11
6. IANA Considerations	11
7. Security Considerations	11
8. References	11
8.1. Normative References	11
8.2. Informative References	11

7. Acknowledgments	12
Authors' Addresses	12

1. Introduction

This draft extends draft-ietf-bier-pim-signaling to mLDP.

Some operators would like to deploy BIER technology in some segment of their network. This draft explains a method to signal mLDP services and stitch it to a BIER domain, with minimal disruption and operational impact to the mLDP domain.

This draft explains the procedures needed to signal and uniquely identify a mLDP P2MP LSP in a BIER 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. BFER is term used for dataplain packet forwarding.

BBR:

BIER Boundary router. The router between the LDP domain and BIER domain. Maintains mLDP adjacency for all routers attached to it on the mLDP domain and terminates the mLDP adjacency toward the 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. If so it terminates the mLDP adjacency toward the BIER domain and signals the mLDP FEC through the BIER core. The router also signals the FEC withdraw or release.

EBBR:

Egress BIER Boundary Router. The egress router in BIER domain from signaling point of view. It terminates the BIER packet sends the mLDP FEC to LDP module to be signaled through the LDP domain.T:

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:

3. M-LDP Signaling Through BIER domain



IBBR will create LDP adjacency to all LDP routers attach to it toward the LDP domain.

When a mLDP label mapping or withdraw arrives, the IBBR first determines whether the root of the FEC is reachable through the BIER domain. As an example, this root is located on a disjoint LDP domain that is reachable through the BIER domain. If so IBBR will forward the FEC to a BIER label assigning entity to allocated a BIER domain label (BTL) which would uniquely represent this P2MP FEC in the BIER domain. As it will be explained in upcoming sections this "BIER label assigning entity" could be an SDN controller or EBBR.

On forwarding plane the IBBR will track all the LDP interfaces on the attach LDP domain which are signaling the same FEC. It creates a stitching point (ILM entry) between the assigned BTL and labels received via the label mapping message on LDP interfaces in LDP domain. This stitching could be a swap or pop and push. With the same token if a label withdraw arrives the IBBR will remove the stitch mapping and forward the FEC and the action to the "BIER label assigning entity" for appropriate action.

3.2. Assigning the BIER sub-domain Tree Label

There could be 2 methods for assigning a BTL. First via a SDN controller and second via the EBBR. In either method for scalability the BTL needs to be assigned from a label pool represented by EBBR prefixID and its BIER sub-domain <EBBR, SD>.

3.2.1. Method 1: IBBR as extraction point to SDN controller

In the first method IBBR will terminate the MDLP signaling from LDP domain. For label mapping/withdraw signaling packets, IBBR will extract the FEC and the action and forward it to the SDN Controller with its own BIER Prefix. The message format is <<IBBR,SD>, FEC, action>.

The SDN Controller has the entire view of the end to end network or at minimum the view of the BIER domain network. BGP-LS could be used to build this network view. The controller will examine the FEC and find the EBBR closes to the root. The procedure to find the EBBR closest to the root is described in ietf-draft-bier-pim-signaling. After the SDN Controller determines the EBBR it will determine whether this is the first occurrence of this specific mLDP FEC, if so it will assign a BTL from the <EBBR, SD> pool to the FEC to uniquely represent the P2MP tree in the BIER domain.

From this point on the SDN Controller keeps track of all new IBBRs which are interested in this P2MP FEC. The SDN Controller will create

a mapping <<<EBBR, SD> (BTL), FEC, action>>, IBBRs>. The SDN Controller will update the relevant BBRs as described in the SDN Controller section.

3.2.2. Method 2, EBBR assigning the BTL

Alternatively the IBBR can signal the <FEC, action> to EBBR via bier in-band signaling in BIER domain. This signaling could be a new BIER TLV or using the mLDP packet as a signaling packet, in par with draft-ietf-bier-pim-signaling.

IBBR will use the ROOT of the mLDP FEC to find EBBR. The procedure to find EBBR is identical to ietf-draft-bier-pim-signaling. After identifying the EBBR, IBBR will encapsulate the mLDP signaling in a BIER packet with the correct EBBR bit set in the bier header and forward the signaling packet into the BIER sub-domain.

The EBBR will examine the signaling packet and will allocate a BTL from its <EBBR, SD> pool. The EBBR then constructs <<<<EBBR, SD> (BTL), FEC, action>, IBBRs> mapping. EBBR needs to signal the <<<<EBBR, SD> (BTL), FEC> to the relevant IBBRs. This EBBR signaling can be done via SDN Controller or BGP, explained in upcoming sections.

3.2.2.1. BIER packet construction at IBBR for Method 2

Assuming the mLDP packet is forwarded from IBBR to EBBR for signaling, the BIER header will be encoded with the BFR-id of the IBBR(with appropriate bit set in the bitstring) and the mLDP signaling packet is then encapsulated in the packet.

```

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
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     BIFT-id                                     | TC | S |         TTL         |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Nibble | Ver | BSL |                                     Entropy                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
| OAM | Rsv | DSCP | Proto |                                     BFIR-id                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     BitString (first 32 bits)                                     ~
+-----+-----+-----+-----+-----+-----+-----+-----+
~
+-----+-----+-----+-----+-----+-----+-----+-----+
~                                     BitString (last 32 bits)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

BIERHeader.Proto = IPv4 or IPv6

BIERHeader.BitString= Bit corresponding to the BFR-ID of the EBBR

BIERHeader.BFIR-id = BFR-Id of the BER originating the encapsulated LDP packet, i.e. the IBBR.

Rest of the values in the BIER header are determined based on the network (MPLS/non-MPLS), capabilities (BSL), and network configuration.

3.2.2.2. Signaling mLDP through the BIER domain procedure

Throughout the BIER domain the BIER forwarding procedure is in par with RFC 8279. No BIER transit router will examine the BIER packet encapsulating the mLDP signaling packet.

The packet will be forwarded through the BIER domain until it reaches the BBR with matching BFR-ID as in the BIERHeader.Bitstring. This BBR (EBBR) will remove the BIER header and examine the mLDP signaling packet farther.

3.3. SDN Controller procedure for updating BBRs with BTL

The BBRs can use openflow message to send the mLDP FEC info to the SDN Controller. On the other direction the controller can use BGP SR-TE signaling (or pcep) to download the mapping information to the BBRs. Detail are outside the scope of this document but a new BGP SR-TE NRLI and TLVs would be needed.

In either method the SDN Controller will track the <<<<EBBR, SD> (BTL), FEC, action>, List of <IBBRs>> and download the <<EBBR, SD> (BTL), FEC, action> to the relevant BBRs.

In method 1 the relevant BBRs are EBBR and the IBBRs that are interested in the P2MP LSP. It should be noted since in this method EBBR does not know about all the IBBRs which are interested in a mLDP FEC, as such the SDN Controller should download to EBBR the <<EBBR, SD> (BTL), FEC, action> mapping with the list of all interested IBBRs. In addition to this, every time an IBBR receives the mLDP FEC it should be signaled to SDN Controller via described method and added to the list. SDN Controller should update the EBBR with that info.

In method 2 the EBBR has assigned this mapping and is aware of all IBBRs that are interested in the FEC. EBBR should signal the <<<<EBBR, SD> (BTL), FEC, action>, List of <IBBRs>> to SDN Controller. SDN Controller downloads the mapping to only IBBRs that are interested in the P2MP lsp. As an example the list of <IBBRs>. EBBR will update the SDN controller with each arriving new IBBR that

wants to join the P2MP LSP. The SDN controller will update these new IBBRs with the relevant mapping.

3.4. BGP procedure for signaling BTL for method 2

As described in the "Method 2, EBBR assigning the BTL" section, BGP can be used to signal the <<EBBR, SD> (BTL), FEC> to IBBR. This signaling can be via MP-BGP MVPN address family with a new NLRI route type Intra-AS BTL Route

```

+-----+
|          SD      (8 octets)          |
+-----+
|          EBBR Router's IP Addr       |
+-----+

```

The PMSI tunnel attribute can be used to signal the (BTL, FEC) to IBBRs. A new "BTL" tunnel type needs to be assigned. The PMSI TUNNEL ATTRIBUTE MPLS label field can be used to encoded BTL. The Tunnel Identifier will contain the FEC.

In addition BGP SR-TE could be used, where the EBBR is generating the NLRI. As per SDN controller section this NLRI is a new NLRI, and the BTL will be part of the SID list (single SID).

3.5. EBBR procedure method

After identifying the BTL for the P2MP FEC (via method 1 or method 2) EBBR will assign a up stream label for the FEC and signal it toward root via mLDP signaling.

With same token the EBBR creates a multicast state with incoming interface as same interface that mLDP signaling packet was forwarded and outgoing interfaces of IBBRs BFIR-ids interested in the P2MP FEC.

EBBR will stitch the upstream label to the downstream BTL with out going interfaces the IBBRs interested in this particular P2MP tree and identified via IBBRs BFIR-id. The EBBR will also build a BIER reverse path forwarding table, using the IBBR BFIR-id. This is explained in section 4.1.

It should be noted EBBR will maintain LDP adjacency toward the LDP domain and all LDP routers which are connected to it.

At this point the end-to-end multicast traffic flow setup is complete.

3.6. Label release

Assuming label release is originated from a disjoint LDP domain, the EBBR needs to signal the release of BTL to all IBBRs interested in this particular mLDP FEC. Also the EBBR will remove the ILM entry for this FEC base on the label release.

3.6.1 Label release in method 1

In method 1 the EBBR can signal the label release to the SDN Controller <<<<EBBR, SD> (BTL), FEC, action>, the SDN Controller will find the list of IBBRs interested in this FEC and will update them accordingly. The IBBR will in addition send a label release to all mLDP neighbors on LDP domain that were signaling this FEC. The EBBR and IBBR will remove the ILM entry for this FEC and the SDN Controller will remove the entry and release the BTL for this FEC as well.

3.6.2 Label release in method 2

In method 2 the same procedure as method 1 can be followed when the SDN Controller is used for signaling. In case of BGP signaling of BTL, EBBR will automatically withdraw the BTL route.

4. Datapath Forwarding

4.1. BFIR tracking of FEC

As explained before the BFIR has a ILM entry which stitches arriving P2MP label to the BIER sub-domain Tree Label.

The BFIR (EBBR) also track all the interested BFERs via arriving binding <<<EBBR, SD> (BTL), FEC, action>>, list of (IBBRs)> from SDN Controller (method 1) or the FEC signaling in (method 2). BFIR should build its multicast tree with incoming interface (IIF) as LDP interface (in LDP domain) and out going interfaces OIFs set as the <SD, BFR-IDs> of the interested BFERs (in BIER Domain).

4.2. Datapath traffic flow

On BFIR when the MPLS label for P2MP LSP arrives a lookup in ILM table is done. Base on the arriving label BFIR will find the stitching forwarding entry. BFIR will swap the incoming MPLS label with the assigned BTL. The swap action can also be a pop of mpls domain label and push of the BTL. BFIR will go through all the BTL's out going interface, (i.e. the IBBRs BFIF-id interested in this P2MP lsp). 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) and find the EBBR label pool base on the arriving packet BFR-ID and its sub-domain. BFER will remove the BIER header and will do a lookup in the <EBBR, SD> ILM for BTL. The BTL entry could be swap to the MPLS domain P2MP label or a pop of BST and push of MPLS domain P2MP. The MPLS domain will forward the packet as per MPLS forwarding procedure to all the MPLS OIFs on the IBBR.

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 sent from IBBR to EBBR via the controller or the inband BIER signaling.

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 <Add any acknowledgements>

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

Andrew Dolganow
Nokia
750D Chai Chee Rd
06-06, Viva Business Park
Singapore 469004

Email: Andrew.dolganow@nokia.com