BIER Workgroup   H. Bidgoli, Ed.
Internet Draft   A. Dolganow
Intended status: Standard Track   J. Kotalwar
                                                             Nokia
                                                        Fengman Xu
                                                           Verizon
                                                       IJ. Wijnand
                                                  Mankamana Mishra
                                               Cisco Systems, Inc.
                                                          Z. Zhang
                                                  Juniper Networks

Expires: October 19, 2018   April 17, 2018

                   PIM Signaling Through BIER Core
                   draft-ietf-bier-pim-signaling-02

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 PIM Joins and Prunes
   to be signaled through a BIER core. Allowing PIM routers to run
   traditional PIM multicast 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.

Copyright Notice

Table of Contents

1. Introduction

   Greenfield deployment of BIER might not be possible for some large
   network. These networks deploy traditional PIM multicast services in
   GRT or in mvpns such as multicast vpns rfc 6037. Typically, each
   portion of these large networks have their own mandates and
   requirements.

   It might be desirable to deploy BIER technology in some part of these
   networks to replace PIM signaling. In such cases downstream PIM
   states need to be signaled over BIER core to the source.

   This draft explains the procedure to signal PIM joins and prunes
   through a BIER core, as such enable provisioning of traditional pim
   services through a BIER core.

   It should be noted that these type of BIER deployments are usually
   with in a single IGP area. As such the procedures in this draft is
   concentrating on a single BIER IGP area.

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. rfc8279] 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 BM into the packet). Each BFIR must have a valid
>   BFR-id assigned. In this draft BIER will be used for
>   forwarding and tunneling of control plain packet (i.e. PIM)
>   and forwarding dataplane packets. BFIR is term used for
>   dataplane 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. In this draft BIER
>   will be used for forwarding and tunneling of control plain
>   packet (i.e. PIM) and forwarding dataplain packets. BFIR is
>   term used for dataplain packet forwarding.

BBR:

>   BIER Boundary router. The router between the PIM domain and
>   BIER domain. Maintains PIM adjacency for all routers attached
>   to it on the PIM domain and terminates the PIM adjacency
>   toward the BIER domain.

IBBR:

>   Ingress BIER Boundary Router. The ingress router from
>   signaling point of view. It maintains PIM adjacency toward the
>   PIM domain and determines if PIM joins and prunes arriving
>   from PIM domain need to be signaled across the BIER domain. If
>   so it terminates the PIM adjacency toward the BIER domain and
>   signals the PIM joins/prunes through the BIER core.

EBBR:

>   Egress BIER Boundary Router. The egress router in BIER domain
>   from signaling point of view. It terminates the BIER packet
>   and forwards the signaled joins and prunes into PIM Domain.

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]. PIM Signaling Through BIER domain

```
                    bbr                     bbr
      |--pim Domain--|-----bier domain-----|--pim domain--|
  S--( A )----------( B ) ---- ( C ) ---- ( D )----------( E )--h


                    ebbr                    ibbr
  Sig    <-----PIM-----|<--Bier Tunneling----|<----PIM------
  (new)


                    bfir                    bfer
      ------------->|--------BIER-------->|-------------> Datapatah
                                                         (no change)
```
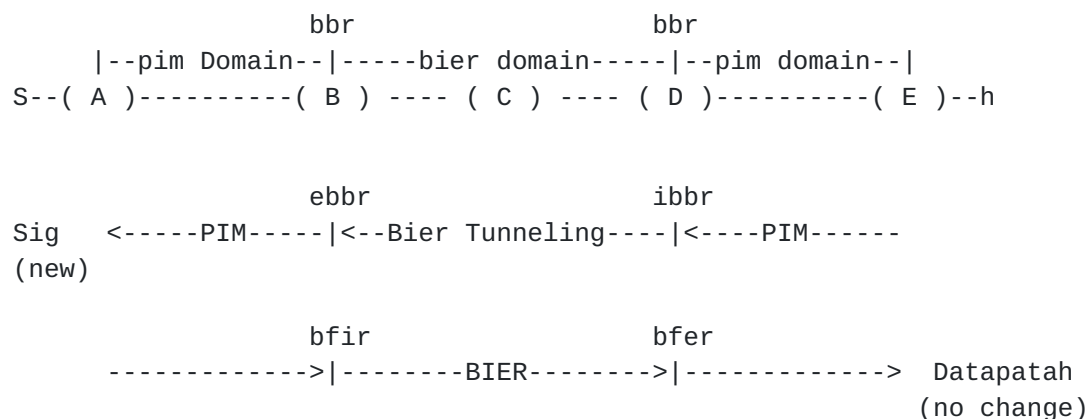
Figure 1: bier boundary router

As per figure 1, the procedures of PIM signaling is done at the BIER
boundary router. The BIER boundary router (BBR) are connected to PIM
capable routers toward the pim domain and BIER routers toward the
bier domain. PIM routers in pim domain continue to send PIM state
messages to the BBR. The BBR will create pim adjacency between all
the PIM routers attach to it on the pim domain. That said the BBR
does not propagate all PIM packets natively into the BIER domain.
Instead when it determines that the PIM join or prune messages needs
to be signaled through the BIER domain it will tunnel the PIM packet
through the BIER network. This tunneling is only done for signaling
purposes and not for creating a PIM adjacency between the two

disjoint pim domains through the bier domain.

The terminology ingress BBR (ibbr) and egress BBR (ebbr) are relative from signaling point of view.

The ingress BBR will determine if an arriving pim join or prune needs to be signaled across the bier domain. While the egress BBR will determine if the bier packet is a signaling packet and propagate the packet to its attach pim domain.

The BFER and BFIR are BBR from datapath point of view. It should be noted the new procedures in this draft are only applicable to signaling and there are no changes from datapath point of view.

3.1. Ingress BBR procedure

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

When a PIM join or prune for certain (S,G) arrives, the IBBR first determines weather the join or prune is meant for a source that is reachable through the bier domain. As an example, this source is located on a disjoint PIM domain that is reachable through the BIER domain. If so the ibbr will try to resolve the source via an ebbr closest to the source.

The procedure to find the ebbr (BFIR from datapath point of view) can be via many mechanisms explained in more detail in upcoming sections.

After discovering the EBBR and its BFR-ID (flooded via IGP BIER extension), the IBBR will construct the BIER header via the BIFT. The signaling packet, in this case the PIM join/prune packet, is encapsulated in the BIER header and transported through BIER domain to EBBR.

On forwarding plane the IBBR will track all the PIM interfaces on the attach pim domain which are interested in a certain (S,G). It creates multicast states for arriving (S,G)s from pim domain, with incoming interface (RPF) as BIER "tunnel" interface and outgoing interface as the pim domain interface(s) on which PIM Join(s) were received on. If there is another PIM Join for the same multicast (S,G) entry on another interface arriving from pim domain, that interface gets added in the outgoing interface list as well.

3.2. **Procedure to determine EBBR on IBBR**

As it was explained in previous section, IBBR needs to determine the
EBBR closest to the source. This is needed to encode the BIER header
BitString field for forwarding of the signaling packet. There can be
many mechanism to determine the EBBR. This section explain some
routing methods that can be used to achieve this. It should be noted
that in most cases the BIER domain is a single IGP area. The PIM
domains are part of the same IGP area as BIER domain(single area) or
are stitched to the BIER domain via an ABR or ASBR. in either case
the BBRs are all located in the same area as bier domain.

3.2.1. **SPF**

On IBBR SPF procedures can be used to find the EBBR closest to the
source.

Assuming the BIER domain is consist of all BIER forwarding routers,
SPF calculation can be used to find the source. A post process can
find the EBBR (BFR closes to the source) by walking the SPT starting
from the source and finding the first Vertex that is a BFR.

3.2.1.2. **Route summarization**

An IGP area could be a BIER sub-domain. In this case the EBBRs and
IBBRs could be ABRs (or non-ABRs). If the route to the source is a
summarized/leaked route by an ABR, the advertising router field for
the summarized route identifies the EBBR.

3.2.1.3. **Constrain shortest path first**

In some cases the BIER domain contains non bier forwarding routers.
It might be desirable to find a SPT with only BFR routers. In this
case CSPF can be used to find the path trough the BIER domain via BFR
only routers. In this case the constrain is to find the path to the
source using BFRs.

3.2.2. **Indirect next-hop**

On IBBR the source can be resolved via EBBR. Once EBBR is identified,
its BIER prefix-id and bfr-id is determined. This information can be
used to build the BIER header with EBBR bit-index. The next-hop of
the source on IBBR can be set to EBBR via multiple methods, including
Static Route and BGP.

3.2.2.1. **Static Route**

On IBBR there can be a static route configured for the source, with

source next-hop set as EBBR BIER prefix id.

3.2.2.2. Interior Border Gateway Protocol (iBGP)

Consider the following topology:

```
                 bbr                      bbr
                 ebbr                     ibbr
        |--pim Domain--|-----bier domain-----|--pim domain--|
     S--( A )----------( B ) ---- ( C ) ---- ( D )----------( E )--h
```
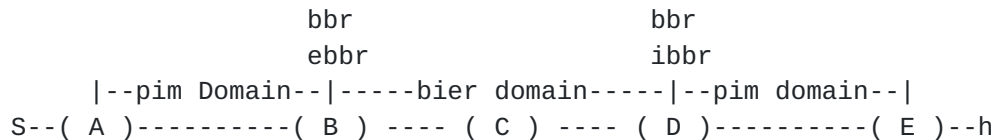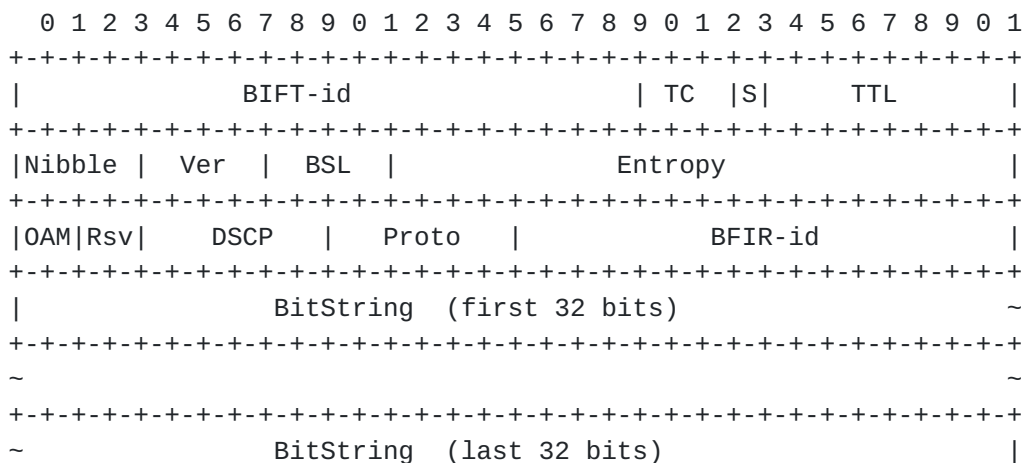
                            Figure 2

Suppose BGP is enable between EBBR (B) and IBBR (D) and the PIM
Domain routes are redistributed to the BIER domain via BGP. This
would include the Multicast Source IP address (S), which resides in
the PIM Domain. In such case BGP should use the same loopback
interface as its next-hop as the BBR prefix-id. This will ensure that
all PIM domain routes, including the Multicast Source IP address (S)
are resolve via BBR's bier prefix id as thier next-hop. When the host
(h) triggers a PIM join message to IBBR (D), IBBR tries to resolve
(S). It resolves (S) via BGP installed route and realizes its next-
hop is EBBR (B). IBBR will use this next-hop (B) to do a lookup in
the BIFT and find its corresponding BIER bit index in the BIFT. Next
IBBR will build the BIER header with corresponding EBBR bit index and
tunnel the PIM signaling message toward EBBR. This procedure is
inline with RFC6826 mLDP in-band signaling section 2.

3.3. BIER packet construction at IBBR

The BIER header will be encoded with the BFR-id of the IBBR(with
appropriate bit set in the bitstring) and the PIM 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
PIM 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.

3.4. Signaling PIM through the BIER domain procedure

Throughout the BIER domain the BIER forwarding procedure is in par
with RFC 8279. No BIER router will examine the BIER packet
encapsulating the PIM signaling packet. As such there is no multicast
state build in the BIER domain.

The packet will be forwarded through the BIER domain until it reaches
the BER with matching BFR-ID as in the BIERHeader.Bitstring. This BER
(EBBR) will remove the BIER header and examine the PIM IPv4 or IPv6
signaling packet farther.

3.5. EBBR procedure

After receiving the BIER packet and determining this packet is a
signaling packet. As such the EBBR will remove the BIER header from
PIM packet and does a route lookup for the source of the pim packet,
if the source is on a local attach pim domain, it forwards the PIM
packet toward the source.

With same token the EBBR creates a multicast state with incoming
interface as same interface that PIM join packet was forwarded and
outgoing interfaces of BIER tunnel with BIER-Header.BFIR-id as one of
the BFER of the tunnel.

The EBBR will also build a BIER reverse path forwarding table, using
the BIERHeader.BFIR-id and the arriving PIM packet (S,G). This is
explained in section 4.1.

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

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

4. **Datapath Forwarding**

4.1. **BFIR tracking of (S,G)**

  For a specific Source and Group, BFIR (EBBR)should track all the
  interested BFERs via arriving PIM signaling from BIER Domain. BFIR
  should build its multicast tree with incoming interface (IIF) as PIM
  interface (in PIM domain) and out going interfaces OIFs set as the
  <SD, BFR-ID> of the interested BFERs (in BIER Domain).

4.2. **Datapath traffic flow**

  When the multicast data traffic arrives on the BFIR (EBBR) the router
  will find all the interested BFERs for that specific (S,G). The
  router then constructs the BIERHeader.BitString with all the BFER
  interested in the group and will forward the packet to the BIER
  domain. The BFER(s) will accept the packets and remove the BIER
  header and forward the multicast packet as per pre-build multicast
  state for (G) and its outgoing interfaces.

5. **PIM-ASM behavior**

  In case of PIM ASM the procedure for LEAFs joining RP is same as
  above. The unicast (source registration) traffic from source to RP
  will be flooded throughout the BIER domain as regular unicast traffic
  without BIER involvement.

6. **Applicability to MVPN**

  With just minor changes, the above procedures apply to MVPN as well,
  with BFIR/BFER/EBBR/IBBR being VPN PEs.

  All the PIM related procedures, and the determination of EBBR happens
  in the context of a VRF, following procedures for PIM-MVPN.

  When a PIM packet arrives from PIM domain attached to the vrf (IBBR),
  and it is determine that the source is reachable via the vrf through
  the BIER domain, a PIM signaling message is sent via BIER to the
  EBBR. In this case usually the PE terminating the PIM-MVPN is the
  EBBR. A label is imposed before the BIER header is imposed, and the
  "proto" field in the BIER header is set to 1 (for "MPLS packet with
  downstream-assigned label at top of stack"). The label is advertised
  by the EBBR/BFIR to associate incoming packets to its correct VRF. In
  many scenarios a label is already bound to the VRF loopback address
  on the EBBR/BFIR and it can be used.

  When a multicast data packet is sent via BIER by an EBBR/BFIR, a
  label is imposed before the BIER packet is imposed, and the "proto"

field in the BIER header is set to 1 (for "MPLS packet with downstream-assigned label at top of stack"). The label is assigned to the VPN consistently on all VRFs [draft-zzhang-bess-mvpn-evpn-aggregation-label].

If the more complicated label allocation scheme is needed for the data packets as specified in [draft-zzhang-bess-mvpn-evpn-aggregation-label], then additional PMSI signaling is needed as specified in [RFC6513].

To support per-area subdomain in this case, the ABRs would need to become VPN PEs and maintain per-VPN state so it is unlikely practical.

7. IANA Considerations

This document contains no actions for IANA.

8. Security Considerations

TBD

9. References

9.1. Normative References

[BIER_ARCH] Wijnands, IJ., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", rfc 8279, October 2016.

9.2. 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.

10. Acknowledgments <Add any acknowledgements>

Authors' Addresses

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

Email: hooman.bidgoli@nokia.com

Fengman Xu
Verizon
400 International PKWY
Richardson, Tx 75081
US

Email: fengman.xu@verizon.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

IJsbrand Wijnands
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Email: ice@cisco.com

Mankamana Mishra
Cisco System
821 alder drive,
Milpitas California
USA

Email: mankamis@cisco.com

Zhaohui Zhang
Juniper Networks

EMail: zzhang@juniper.net