BIER for EVPN BUM Traffic

draft-zzhang-bier-evpn-02

Jeffrey Zhang
Tony Przygienda
Ali Sajassi
Jorge Rabadan

IETF106, Nov 2019
Singapore

BIER for EVPN services
A short refresh

BIER as provider/underlay tunnel for EVPN BUM traffic

- EVPN-MPLS or EVPN-VXLAN/NVGRE

Very similar to BIER for MVPN

- Some text borrowed verbatim from RFC8556

The document covers EVPN specific aspects including:

- Leaf Tracking
- Allowed data planes and EVPN multi-homing techniques for each encapsulation
- Segmentation

BIER for EVPN services

PMSI Tunnel Attribute (PTA)

Tunnel Type: BIER

Tunnel ID: Subdomain-ID + BFIR-ID + BFIR Prefix

- A “BIER tunnel” is considered as an aggregation tunnel
- One tunnel is used for many BDs

Label:

- MPLS label upstream-assigned by BFIR
  - Per Bridge Table (BT) when there is no segmentation
  - Per PMSI in case of segmentation
- Globally-unique VNI/VSI in case of VXLAN/NVGRE

Flags

- Leaf Information Required (LIR) bit
  - Requesting Leaf routes in response to x-PMSI A-D routes
- Leaf Information Required per Flow (LIR-pF) bit
  - Requesting Leaf A-D routes for more specific flows

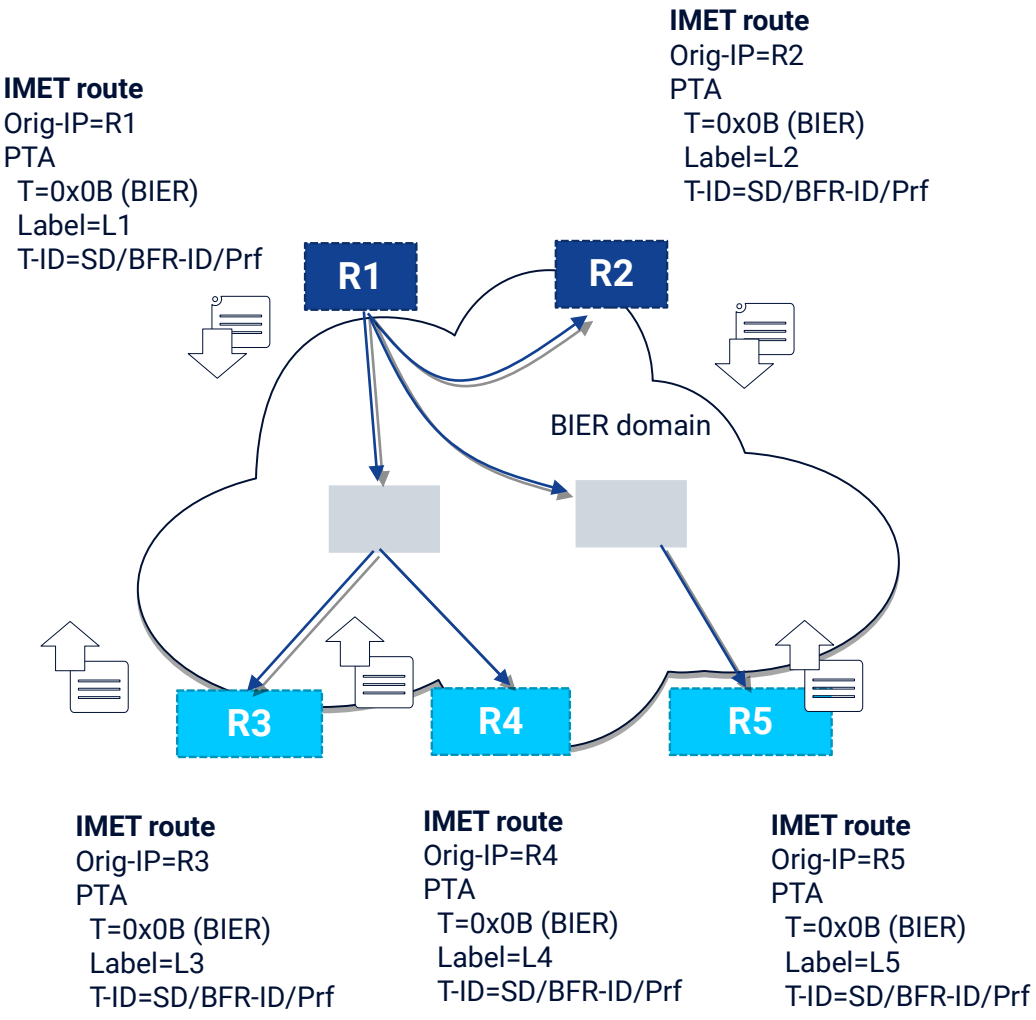

BIER for EVPN services
Leaf Tracking

IMET/SMET routes provide explicit tracking functionality

- IMET routes carry PTA but SMET routes do not

If it is desired not to use SMET route for every flow, S-PMSI route can solicit Leaf routes for certain flows

- Only those flows will use selective forwarding while others will use flooding
  - A compromise between optimized forwarding and reduced control plane state
- Leaf routes can be solicited by one-to-one mapped S-PMSI routes with LIR bit
- Or, a single S-PMSI route with LIR-pF bit could trigger many Leaf A-D routes
  - E.g, a (C-*,C-G) S-PMSI route with LIR-pF bit triggers individual (C-S,C-G) Leaf A-D routes

BIER for EVPN services
Data Plane and Multi-homing

VXLAN/NVGRE/GENEVE:

BIER Header + VXLAN/NVGRE/GENEVE header + Ether frame

- No IP/UDP header used; new BIER proto/payload type in BIER header indicates that VXLAN/NVGRE header is following the BIER header
- Split-horizon with local-bias works fine as BFIR-ID in the BIER header identifies the sending PE

MPLS:

BIER Header + PMSI/BT Label [+ ESI label] + Ether frame

- Two upstream-assigned labels are used if traffic is from a multi-homed ES

BIER for EVPN services
Segmentation

Segmentation may be used in the following situations for MPLS

- If different tunnel types in different AS/areas is required
- If a larger BIER domain needs to be divided into multiple smaller sub-domains
  - So that smaller BitString can be used w/o using multiple sets
- PEs attached to the same ES should be in the same segmentation region

Segmentation points update the PTA when re-advertising IMET/S-PMSI AD routes to specify new tunnel type/ID

- Label switching is done at the segmentation points
- If tunnel aggregation is used in the next region (e.g., BIER), a new label is assigned and advertised by the segmentation point for the PMSI

VXLAN:

What is new in revision 02?

- If no PHP (I-D.ietf-bier-php), VXLAN header goes directly below the BIER header (BIER header indicates the next protocol – saves overhead and ip multicast allocation)
- If PHP, an IP/UDP header MUST be added between the BIER and VXLAN headers (so that the IP/UDP header indicates the next-protocol)

Optional use of the “Auxiliary Information” in the PTA in case of PHP

- Tunnel encapsulation TLV (as defined in I-D.ietf-idr-tunnel-encaps)
- Includes the tunnel type (VXLAN/NVGRE/GENEVE) and tunnel endpoint (must be a multicast IP address)
- The BFIR can optionally signal the IP multicast address to be used when sending BIER VXLAN/NVGRE/GENEVE packets to the BFERs

Next steps

Request feedback

The Authors believe the document is ready for WG LC

Thank you

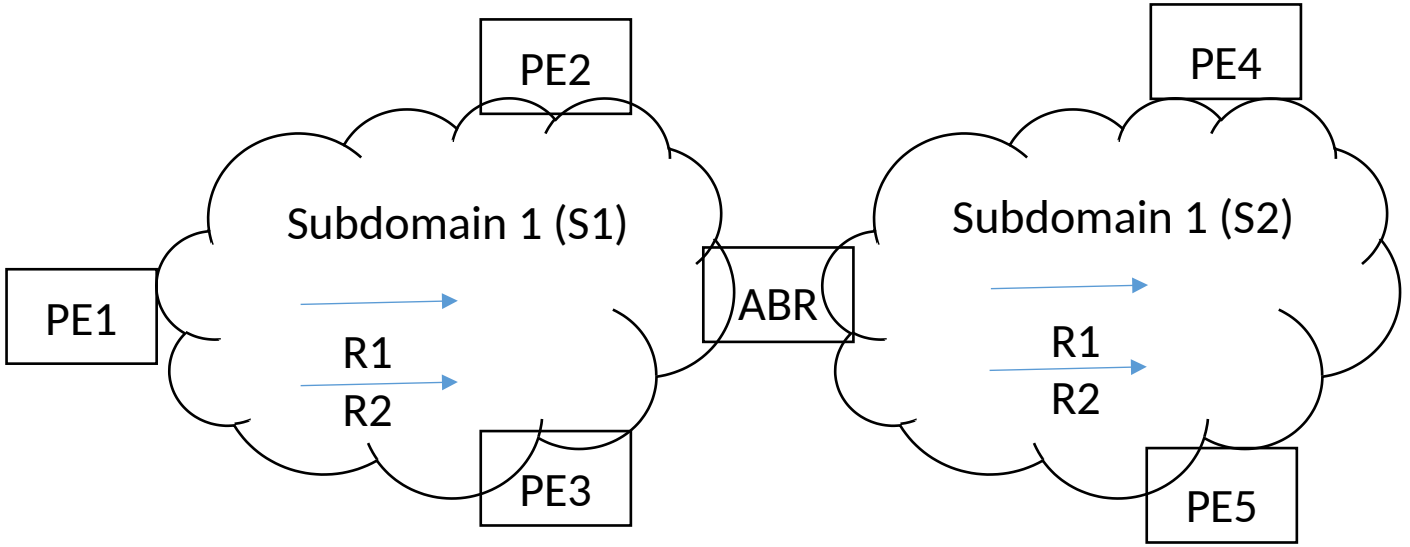

- PTA in S-PMSI R1 advertised by PE1: BIER, S1, PE1, L1, LIR
- PTA in S-PMSI R2 advertised by PE1: BIER, S1, PE1, L2, LIR
- PTA in S-PMSI R1 re-advertised by ABR: BIER, S2, ABR, L3, LIR
- PTA in S-PMSI R2 re-advertised by ABR: BIER, S2, ABR, L4, LIR

- PE1 sends BIER packets with BitString including PEs/ABRs in S1
- After ABR decapsulates BIER packets received in S1, L1 or L2 is exposed
- ABR label switches L1 to L3 (or L2 to L4), and then encapsulates with BIER header and send into S2
    - Because of label switching (no ether lookup), different labels must be advertised in the PTA for different routes