

# BIER for EVPN BUM Traffic

draft-zhang-bier-evpn-00

Jeffrey Zhang

Tony Przygienda

Ali Sajassi

Jorge Rabadan

99<sup>th</sup> IETF, Prague

# Overview

- 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 draft-ietf-bier-mvpn

# PMSI Tunnel Attribute

- Like with MVPN x-PMSI routes, EVPN IMET routes (RFC 7432) and x-PMSI (draft-ietf-bess-evpn-bum-procedure-updates) carry PMSI Tunnel Attribute (PTA)
- PTA identifies the tunnel advertised in the routes
  - Tunnel Type
  - Tunnel Identifier
  - MPLS Label
  - Flags

# PTA with BIER for EVPN

- Tunnel Type: BIER
- Tunnel ID: Subdomain-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

# Leaf Tracking

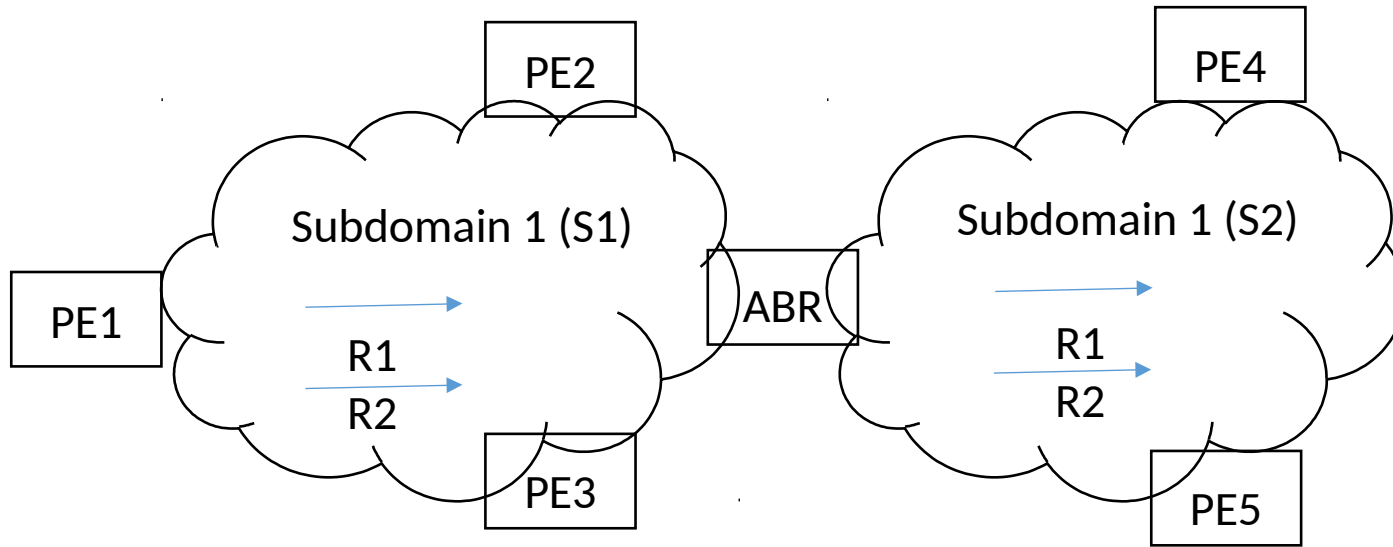
- IMET/SMET routes provide explicit tracking functionality
  - IMET routes carry PTA but SMET routes do not
- If it is desired not to use SMET route for every flow, S-PMSI route can solicit Leaf routes for certain flows
  - Only those flows will use selective forwarding while others will use flooding
    - 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
- Other use cases of x-PMSI/Leaf routes are documented in
  - draft-ietf-bess-evpn-bum-procedures-update
  - draft-zzhang-bess-mvpn-evpn-cmcast-enhancements

# Data Plane & Multi-homing

- VXLAN/NVGRE: BIER Header + VXLAN/NVGRE 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
  - Notice that two upstream-assigned labels are used if traffic is from a multi-homed ES

# Segmentation

- Segmentation may be used in the following situations for MPLS
  - When it is required or desired to use different tunnel types or instances in different AS/areas (regions)
  - For BIER in particular, segmentation may be used to divide a larger BIER domain to multiple smaller sub-domains
    - So that smaller BitString can be used w/o using multiple sets
  - It is recommended that PEs on the same ES be in the same segmentation region
- Segmentation points update the PTA when re-advertising x-PMSI 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



- 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



# Next steps

- Request WG comments
- Polish draft and request adoption afterwards