

# EVPN BUM Procedure Updates

draft-zzhang-bess-evpn-bum-procedure-updates-01

Jeffrey Zhang

Wen Lin

Jorge Rabadan

Keyur Patel

95<sup>th</sup> IETF

# Summary

- Selective Multicast based on S-PMSI/Leaf Auto Discovery (AD) routes
  - Based on RFC 7117 (VPLS Multicast)
    - Similar to MVPN procedures
- Inter-region Tunnel Segmentation
  - A region could be an AS, an area, or even part of an area
    - Region Border Routers (RBRs) are the segmentation points
  - Similar to MVPN inter-AS/area segmentation
  - Based on PMSI/Leaf AD routes

# PMSI/Leaf AD Procedures

- PMSI - Provider Multicast Service Interface
  - Inclusive/Selective PMSI
  - MVPN/VPLS-MCAST concepts applicable to EVPN as well
    - Inclusive/Selective Multicast Ethernet Tag (IMET/SMET) routes correspond to I/S-PMSI AD routes
- An I/S-PMSI AD route from an ingress PE announces the tunnel that instantiates the PMSI
  - Provider Tunnel Attribute specifies tunnel information
- An S-PMSI AD route includes (S,G) information
  - where both S and G could be wildcard
- An egress PE joins the tunnel that instantiates an S-PMSI by:
  - Sending PIM join or mLDP label mapping, or
  - Sending corresponding Leaf AD routes for other tunnel types
    - Ingress Replication, RSVP-TE P2MP, BIER

# Tunnel Segmentation 1/2

- An end-end PE-PE tunnel can consist of segments of intra-AS/area tunnels
- Different tunnel types/instances may be required/desired in different AS/area
  - E.g. mLDP in AS1, RSVP-TE P2MP in AS2, BIER in AS3
  - For technical or administrative reasons
- Segmented IR can reduce replication fan-out
  - Each segment only includes PEs and border routers in the local AS/area
- Segmented BIER can reduce the BitStringLength in case of a large BIER domain
  - A 1024-node BIER domain can be divided into four 256-node sub-domains
    - BitStringLength of 256 instead of 1024
- MVPN: RFC 6514 (inter-as segmenation) and 7524 (inter-area segmentation)
- EVPN: this draft-zzhang-bess-evpn-bum-procedure-updates
  - Inter-region – where a region can be an AS, an area, or part of an area

# Tunnel Segmentation 2/2

- When an RBR re-advertises an I/S-PMSI route to the next region, it modifies the tunnel type/ID advertised by the route to what's used in that region.
- PEs/RBRs in that region joins the tunnel segment for that region by
  - Sending PIM join or mLDP label mapping, or
  - Sending corresponding Leaf AD routes for other tunnel types
    - Targeted at the upstream RBR

# I-PMSI AD route

- An I-PMSI AD route can be per-PE/region
- Per-PE I-PMSI AD route – identifies a PE
  - MVPN type-1 route (Intra-AS I-PMSI AD route)
  - EVPN type-3 route (IMET: Inclusive Multicast Ethernet Tag route)
- Per-region I-PMSI AD route – identifies a region
  - Allows aggregation from PEs to regions
    - Per-PE I-PMSI routes are not propagated across region
      - They're summarized into per-region I-PMSI routes
  - MVPN type-2 route (Inter-AS I-PMSI AD route)
    - NLRI encodes an AS number
  - With EVPN a region could be an AS, area, or sub-area
    - NLRI encodes an 8-octet field encoded just like an EC
      - Allowing maximum flexibility of identifying a region
      - This is the main difference between revision -00 and -01

# Relationship with draft-sajassi-bess-evpn-igmp-ml-d-proxy

- Draft-sajassi covers two aspects
  1. Turn IGMP/MLD soft-state messages to EVPN SMET routes
  2. SMET routes allows Ingress PEs to send traffic on selective tunnels
    - w/o the need for S-PMSI AD routes
      - SMET routes are like analogous to Leaf AD routes
- Draft-zzhang and draft-sajassi are complementary wrt selective multicast
  - Not mutually exclusive
  - SMET routes in draft-sajassi work well if non-segmented IR/BIER is used for all flows
  - S-PMSI/Leaf AD routes and procedures in draft-zzhang are needed for other situations
    - Segmented tunnels
    - PIM/mLDP/RSVP-TE P2MP tunnel
    - IR/BIER selective tunnel is used only for some flows
      - This could be for the purpose of reducing the number of SMET routes

# Next Steps

- Seeking more feedback from WG
- Request WG adoption