

# MVPN/EVPN-BUM Segmented Forwarding

Jeffrey Zhang, Jingrong Xie

IETF 104<sup>th</sup>, Prague

draft-zhang-bess-mvpn-evpn-segmented-forwarding

draft-xie-bess-mvpn-segmented-updates

draft-xie-bier-mvpn-segmented

# Background

- MVPN/EVPN PMSI – Provider Multicast Service Interface
  - A virtual interface for sending/receiving customer multicast traffic through provider network
  - Instantiated with provider tunnels
  - Advertised with Inclusive/Selective-PMSI A-D routes
    - A PMSI Tunnel Attribute (PTA) in the routes specify tunnel type/ID
    - EVPN IMET route is the equivalent of I-PMSI A-D routes
- MVPN/EVPN PMSI Tunnel Segmentation
  - Different tunnel type/instance used in different AS/area/regions
    - For administrative or technical reasons
  - When segmentation points re-advertise an x-PMSI A-D route from an upstream region to downstream regions, it updates the PTA to specify the tunnel used in the downstream region

# Forwarding at Segmentation Points

- Assumed to be label switching
  - Traffic arriving on an upstream segment label switched to downstream segments
  - Like unicast inter-as Option-B
  - Requires per-PMSI label:
    - If Ingress Replication (IR) is used, label is advertised in the PTA of Leaf A-D route
      - Leaf A-D routes are (normally) in response to S-PMSI A-D routes
    - If mLDP/RSVP-TE P2MP tunnel w/o aggregation is used, tunnel label is the per-PMSI label
    - If mLDP/RSVP-TE P2MP tunnel aggregation, or BIER is used, the per-PMSI label is advertised in the PTA of x-PMSI A-D routes
- ***X-PMSI A-D routes are needed in case of P2MP/BIER***
  - There could be a lot of S-PMSI A-D routes if one wants to use selective tunnels for individual flows

# Reducing S-PMSI A-D Routes

- For RSVP-TE P2MP, IR and BIER, ingress PEs track egress PEs for selective forwarding
  - Explicit Leaf Tracking; via Leaf A-D routes
    - EVPN SMET routes are equivalent of Leaf A-D routes
  - S-PMSI A-D routes have Leaf Information Required (LIR) flag set in PTA, triggering Leaf A-D routes from egress PEs that needs to receive traffic
- A LIR-pF flag in a (\*,G)/(\*,\*) wildcard S-PMSI can be used to trigger more specific Leaf A-D routes without more specific S-PMSI A-D routes
  - This allows leaf tracking for individual flows w/o individual S-PMSI A-D routes
    - And use of individual selective tunnels for those flows
    - EVPN SMET routes can be viewed as if triggered by implicit (\*,\*) S-PMSI A-D routes with the LIR-pF flag
  - But that does not work with segmentation
    - Because there are no corresponding S-PMSI A-D routes to advertise per-PMSI label

# draft-xie-bier-mvpn-segmented

- Draft-ietf-bier-mvpn explicitly disallows LIR-pF with segmentation
- Draft-xie-bier-mvpn-segmented documents the use of IP forwarding at segmentation points
  - To allow LIR-pF with segmentation for the BIER & MVPN specific use case
- Lots of discussions
  - IP forwarding is much more heavy weighted
  - Instead of IP forwarding, problem can be solved by triggering individual S-PMSI A-D routes from Leaf A-D routes that are triggered by LIF-pF flag
    - This gives per-PMSI labels to allow label switching
    - Counter argument is that it leads to more control plane state
      - Counter argument for that: forwarding state is less compared to IP forwarding option
    - Another counter argument is that this can cause traffic loss because of added signaling
      - Counter argument for that: the delay can be mitigated by switching to individual selective tunnels after a delay (traffic continue to follow on less specific tunnel for a while before switching)
  - This is also not specific to BIER or MVPN

# draft-zzhang-bess-mvpn-evpn-segmented-forwarding

- The problem/solution/discussions are:
  - applicable to both MVPN and EVPN BUM, not specific to BIER
- draft-zzhang-bess-mvpn-evpn-segmented-forwarding
  - Explaining that switching has been the de-facto forwarding option at segmentation points
    - And why LIR-pF was not allowed with segmentation
  - Documenting that IP forwarding could be used if really desired
    - Pros and cons of label switching vs. ip forwarding
  - Laying out basic rules for label allocation to allow:
    - Label switching whenever desired and possible (even with the IP forwarding option)
    - IP forwarding w/o requiring VRFs on segmentation points

# draft-xie-bess-mvpn-segmented-updates

- Draft-xie-bier-mvpnd-segmente replaced with draft-xie-bess-mvpn-segmented-updates
  - Since it is not BIER specific
- The authors of the three drafts agree to consolidate to draft-zzhang-bess-mvpn-evpn-segmented-forwarding

# Next Steps

- Polish draft-zzhang
  - Likely to take BIER specific text from draft-xie into draft-zzhang as examples
- Seek comments
- Seek adoption when it is ready (not yet)