

# MVPN Regional Segmentation

draft-zhang-bess-mvpn-regional-segmentation

Jeffrey Zhang, Juniper Networks

BESS, IETF 112

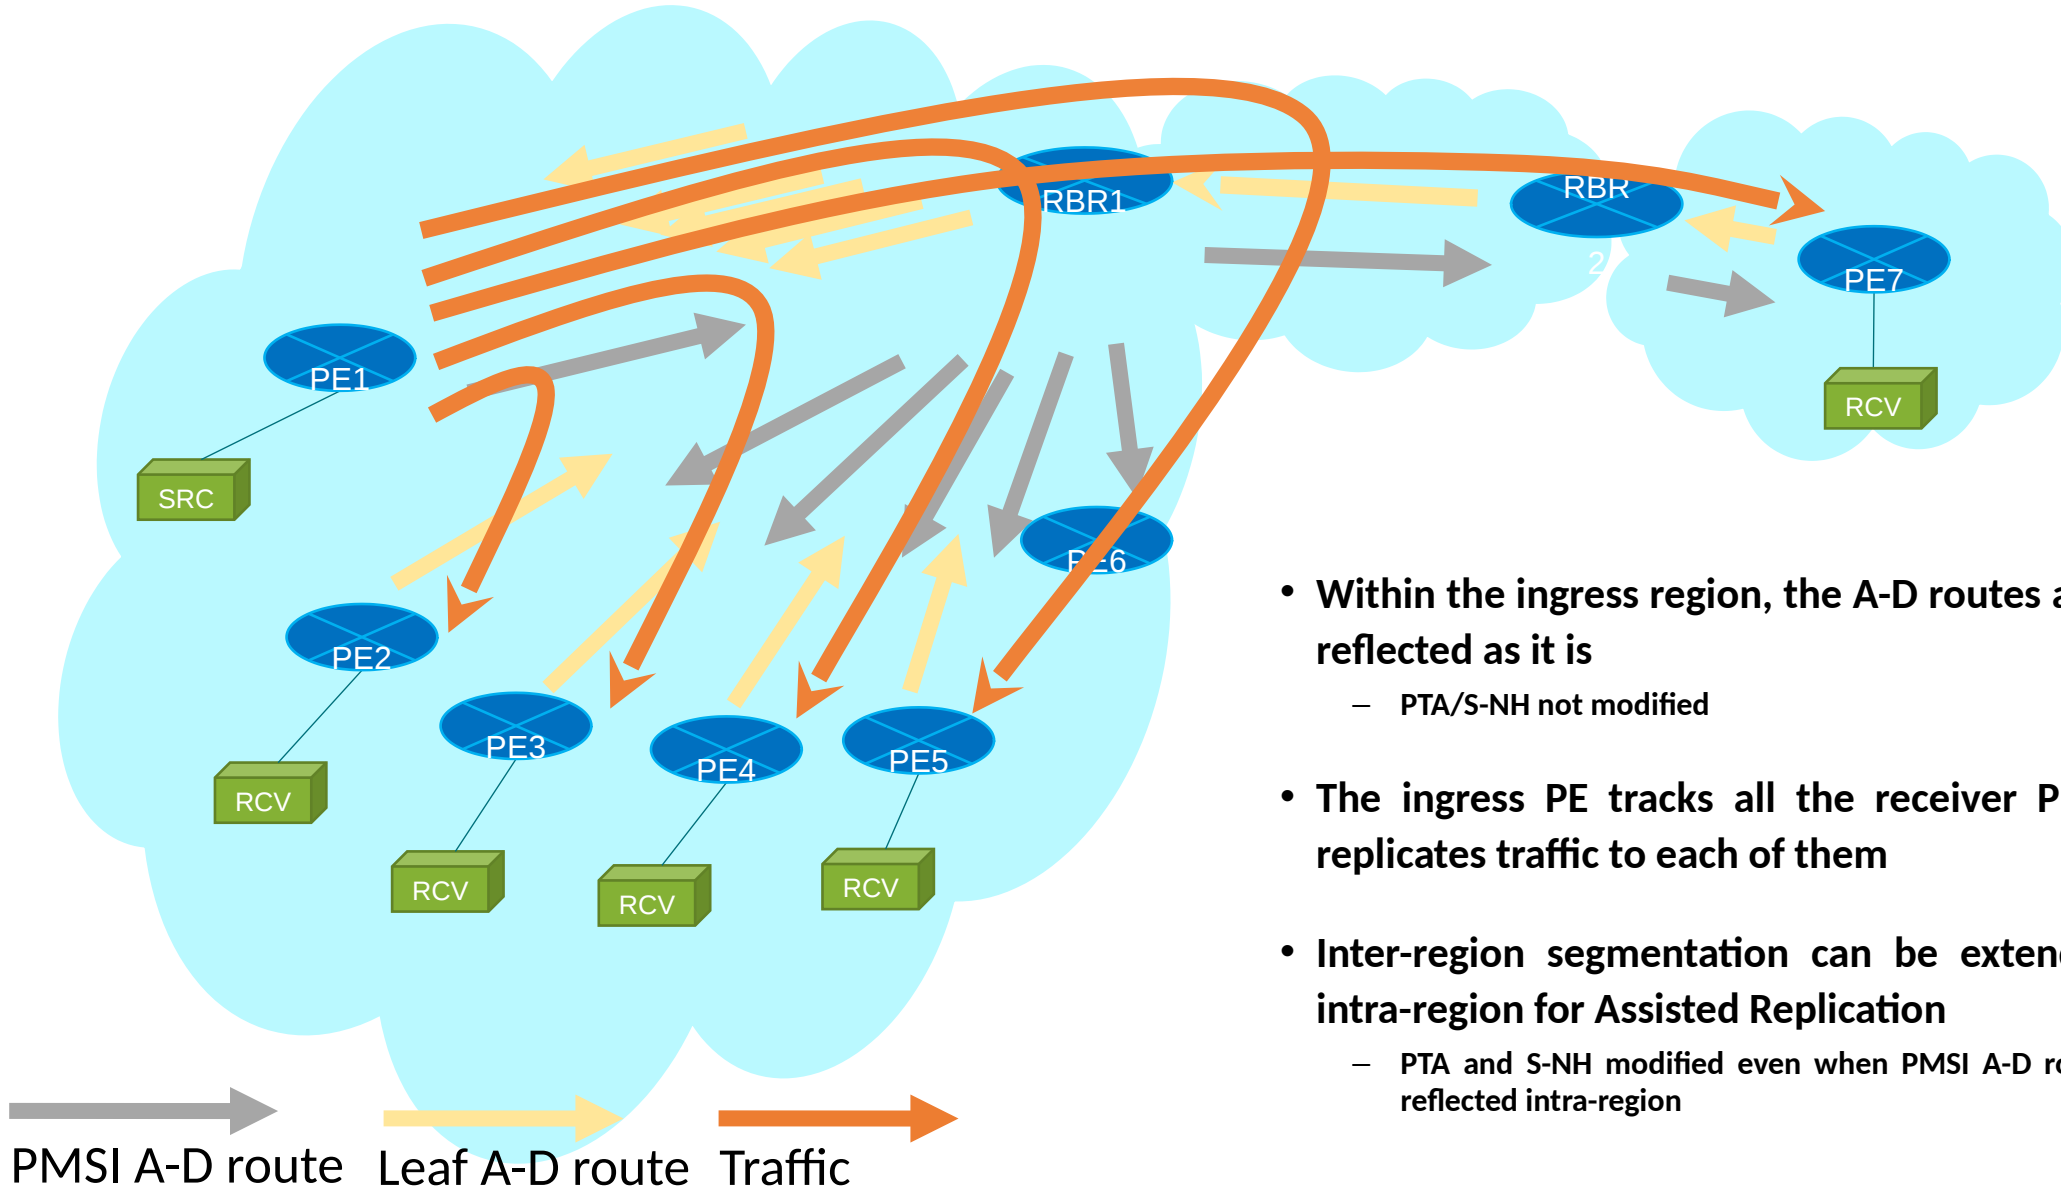
# Background: RFC7524

- MVPN/VPLS inter-area segmentation
  - *Inter-Area P2MP Segmented LSP*
  - A Provider tunnel consists of intra-area segments of different types and/or instances
    - ABR as RR modifies Tunnel Type/ID in PMSI Tunnel Attribute (PTA) when reflecting PMSI A-D routes to the next area
    - Also updates S-NH EC
      - Segmented Next-Hop Extended Community
      - To direct triggered Leaf A-D routes to the ABR
- RFC7524 is limited to “inter-area” and “LSP”
  - Can be applied to inter/intra-region
  - Can be applied to non-MPLS

# Region vs. Area

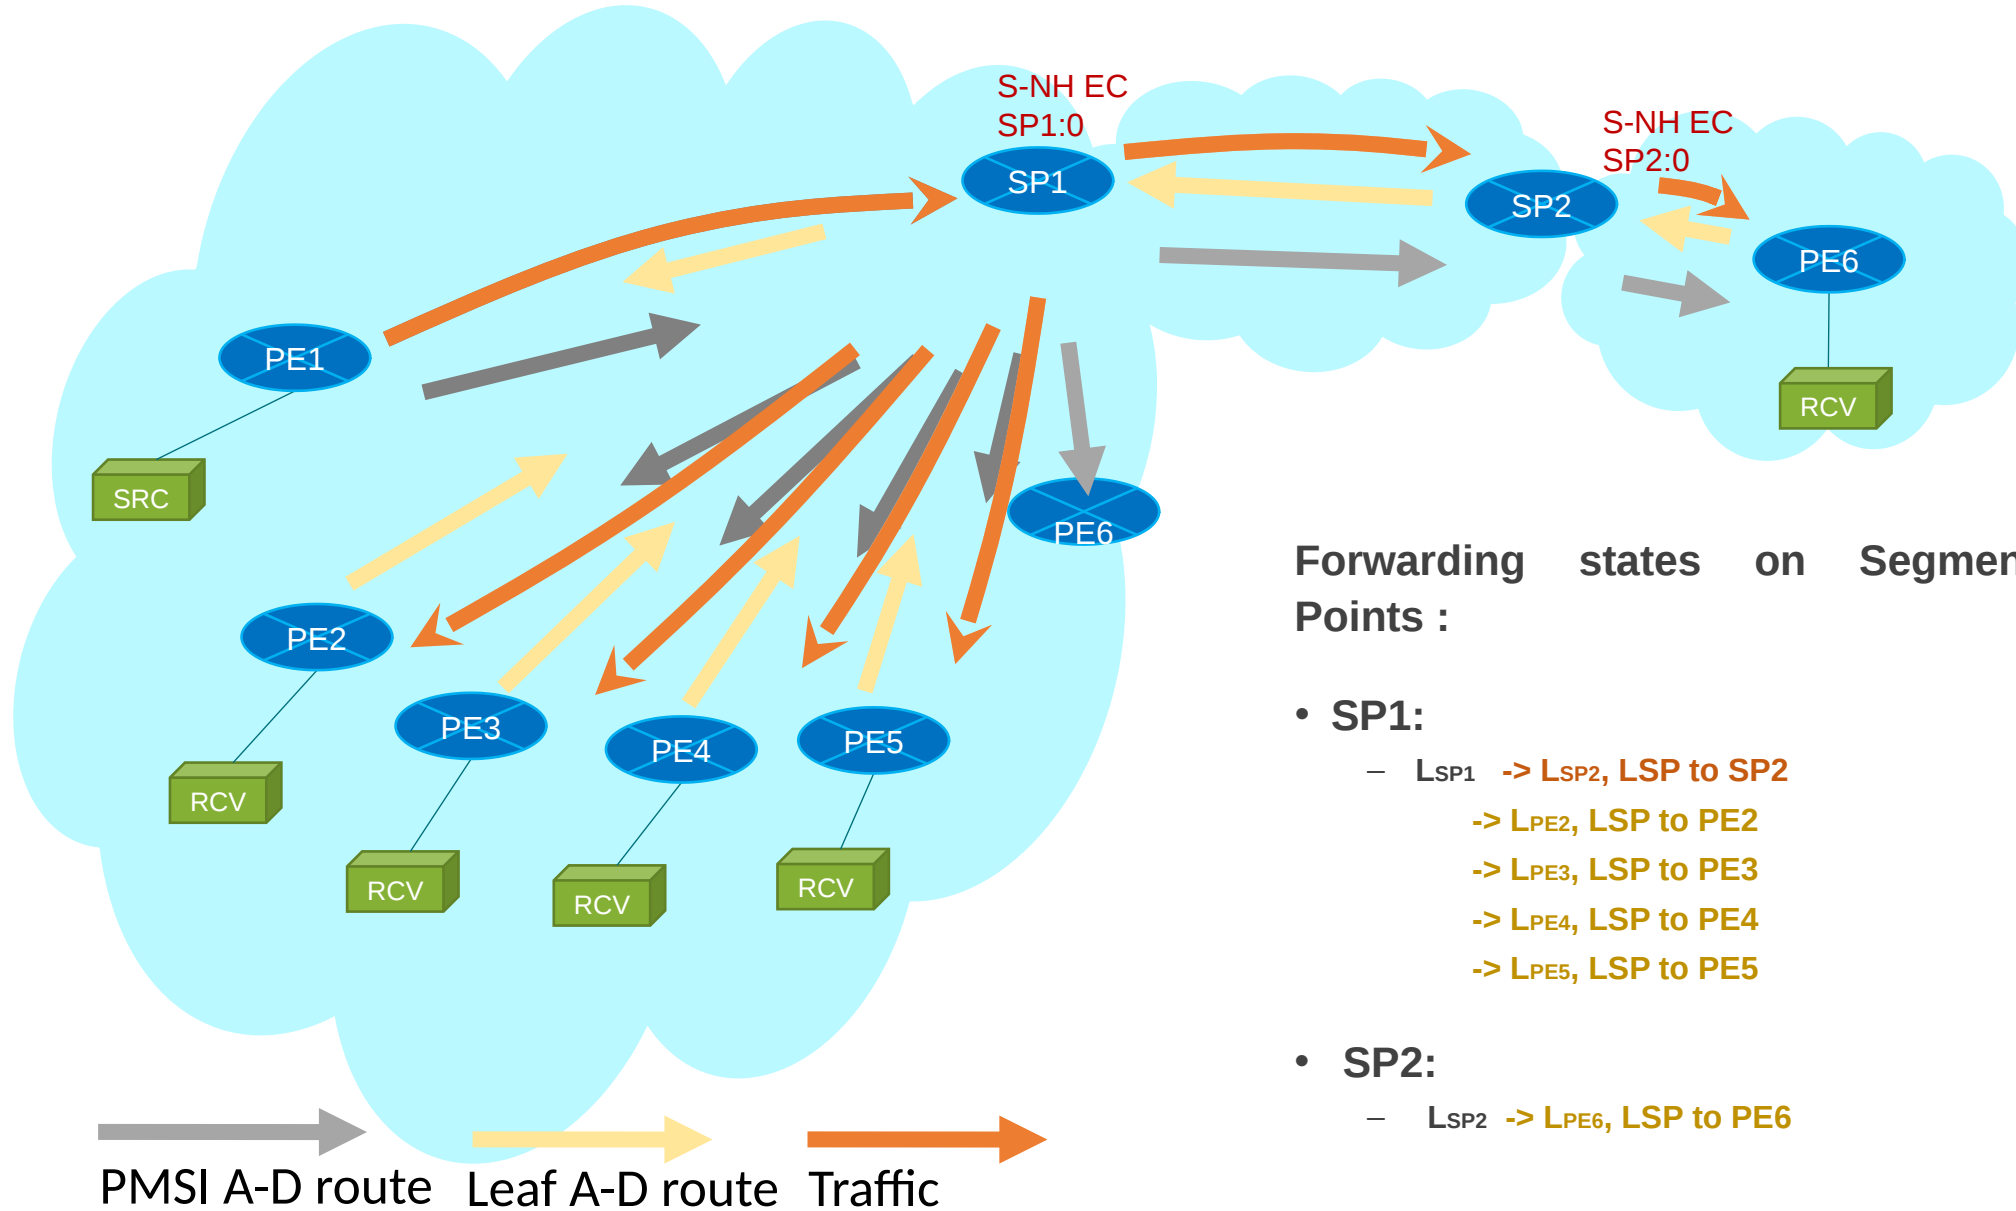
- To modify the PTA and S-NH, typically a BGP neighbor group is used for all peers that should get the same PTA/S-NH
  - This allows segmentation at Regional Border Routers (RBRs)
- A region can be arbitrarily defined via BGP neighbor group
  - A IGP sub-area
  - An AS plus its external link
  - Several ASes
- No new signaling needed
- This concept is already adopted in draft-ietf-bess-evpn-bum-procedure-updates

# Ingress Replication w/o Assisted Replication



- **Within the ingress region, the A-D routes are just reflected as it is**
  - PTA/S-NH not modified
- **The ingress PE tracks all the receiver PEs and replicates traffic to each of them**
- **Inter-region segmentation can be extended to intra-region for Assisted Replication**
  - PTA and S-NH modified even when PMSI A-D routes are reflected intra-region

# Assisted Replication via Intra-Region Segmentation



Forwarding states on Segmentation Points :

- **SP1:**
  - LSP1 -> LSP2, LSP to SP2
  - > LPE2, LSP to PE2
  - > LPE3, LSP to PE3
  - > LPE4, LSP to PE4
  - > LPE5, LSP to PE5
  
- **SP2:**
  - LSP2 -> LPE6, LSP to PE6

# Bud Node Support

- Bud Node
  - A RBR/SP (Segmentation Point) stitches upstream segment to downstream segments – label swap
  - A RBR may also have local receivers – IP forwarding
- An implementation typically can replicate an incoming packet for both label swapping to downstream segments and IP forwarding in local VRF
- Sometimes an implementation may not be able to do that
  - It needs one copy for label swapping and another copy for IP forwarding

# Requesting Extra Copy for Local IP Forwarding

- This is only to accommodate bud nodes that need the extra copy
- Upstream tunnel being RSVP P2MP
  - RSVP P2MP signaling allows for an upstream PHOP node to send extra copy
- Upstream tunnel being IR/BIER/mLDP
  - Extra copy needs to be tunneled directly from upstream PE/RBR
  - This is requested via a Tunnel Encapsulation Attribute (TEA) in a Leaf A-D route sent from the bud node to its upstream PE/RBR
    - Sent even if Leaf Information Required flag is not set in the PMSI route
    - With a Tunnel TLV plus a Tree Label sub-TLV
      - Tree Label is used to identify the local VRF

# Summary

- Inter-area segmentation in RFC 7524 extended to:
  - Inter-region, where a region is defined as a BGP Neighbor Group
  - Intra-region, for Assisted Replication
- Allow a bud node to request extra copy of traffic for local forwarding
  - In case of IR/BIER/mLDP
  - Via TEA added to Leaf A-D route
- Comments appreciated