

Support C-Bidir with Ingress Replication

draft-ietf-l3vpn-mvpn-bidir-ingress-replication

**Jeffrey Zhang
Yakov Rekhter
Andrew Dolganow**

89th IETF, London

Summary

Presented in 87th IETF, Berlin

Accepted as WG Item recently

Inter-area P-tunnel Segmentation was out of scope

- Now we've figured out how to support it

Present the inter-area segmentation support

- Assuming that the audience is familiar with the base technology

Seek WG feedback/consensus before updating the draft

Simulating “Partial Mesh of MP2MP P-tunnels” with IR

Each partition has its own MP2MP P-tunnel simulated with IR

- A partition is the set of PEs selecting the same UMH wrt a C-RPA

One S-PMSI A-D route from the UMH wrt a C-RPA

- Identifying the MP2MP tunnel
- PTA specifies IR, with a label for other PEs to send traffic to the UMH

One Leaf A-D route responded from each PE in the same partition/tunnel

- Imported by all PEs: Not just by the S-PMSI originator
 - Even those not in the same partition; but those will just ignore the route
- PTA includes a label corresponding to the UMH
 - For other PEs to send traffic to the Leaf A-D route’s originator

Inter-area segmentation: Regular IR for C-unidir

Ingress PEs replicate to local PEs and ABRs

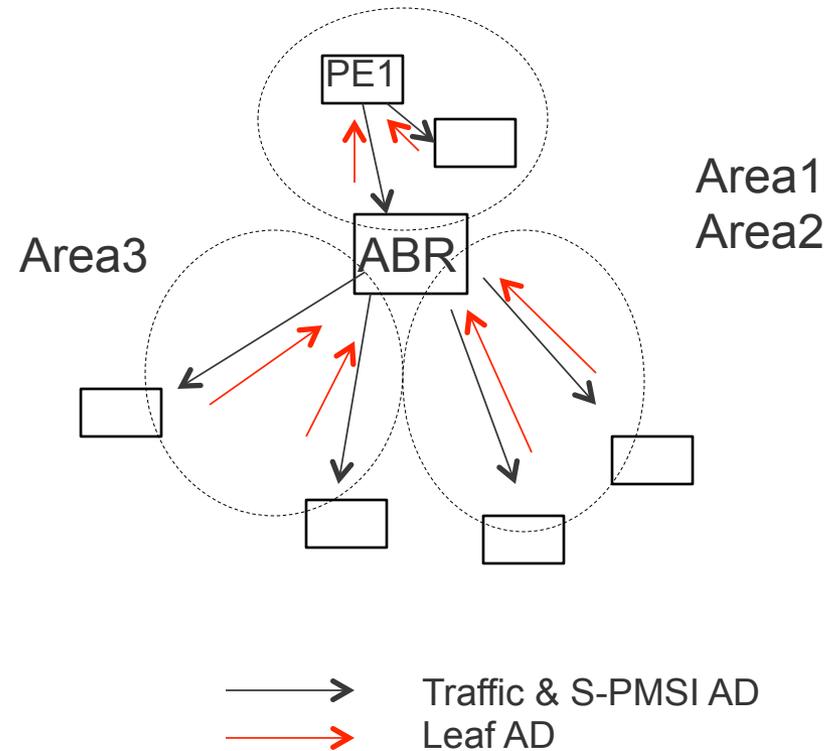
ABRs replicate to downstream PEs/ABRs via label switching

ABRs reflect S-PMSI AD routes with S-NH-EC

Leaf AD routes sent to upstream ABRs

- Follow the S-NH-EC in S-PMSI AD routes

ABRs reflect Leaf-AD routes further upstream, with updated RT pointing to the upstream



Inter-area segmentation: IR for C-bidir 1/2

Per-area partitions

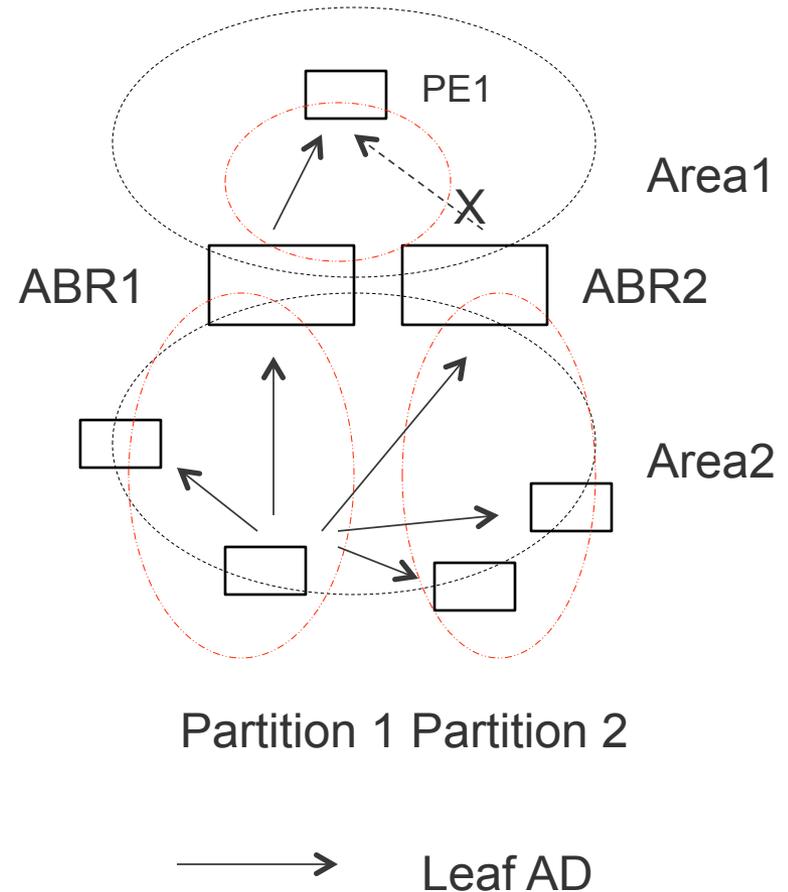
- Ingress area (wrt C-RPA): PEs/ABRs picking the same UMH
- Downstream area: PEs/ABRs that pick the same active S-PMSI AD route

Leaf AD routes are imported by all PEs/ABRs in an area

- Even those not in the same partition

Only the ABR in the same partition reflects the Leaf AD routes upstream

- It re-advertised the corresponding S-PMSI AD routes downstream



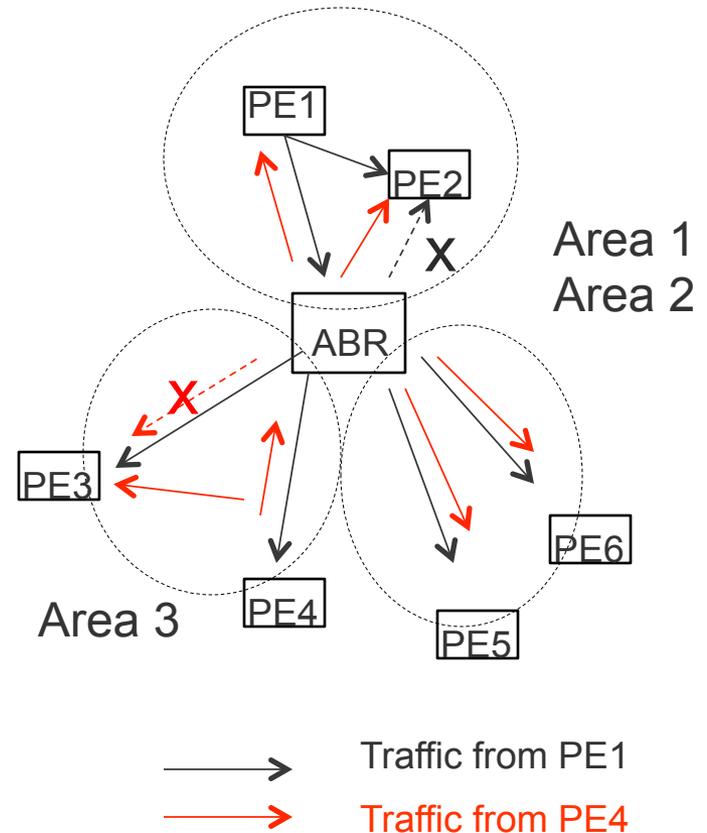
Inter-area segmentation: IR for C-bidir 2/2

ABRs label switch bidirectional traffic, requiring split-horizon

- ABR label switches traffic from PE4 to PE1/2/5/6
 - But not to PE3
- ABR label switches traffic from PE1 to PE3/4/5/6
 - But not to PE2

When ABRs re-advertise S-PMSI AD routes to different downstream areas, different labels must be used

- So that it can install different label switching routes to switch traffic accordingly



Plan

Seek comments

Update draft with inter-area segmentation support

- Per-area partitions
- Only the ABR in the same partition reflects Leaf AD routes upstream
- ABRs advertise different labels when re-advertising S-PMSI AD routes to different downstream areas

BTW –

Please review & comment draft-rosen-l3vpn-mvpn-ir as well

- Ingress Replication Tunnels in Multicast VPN