

Segmented MVPN Using IP Lookup for BIER

draft-xie-bier-mvpn-segmented-01

Jingrong Xie @Huawei

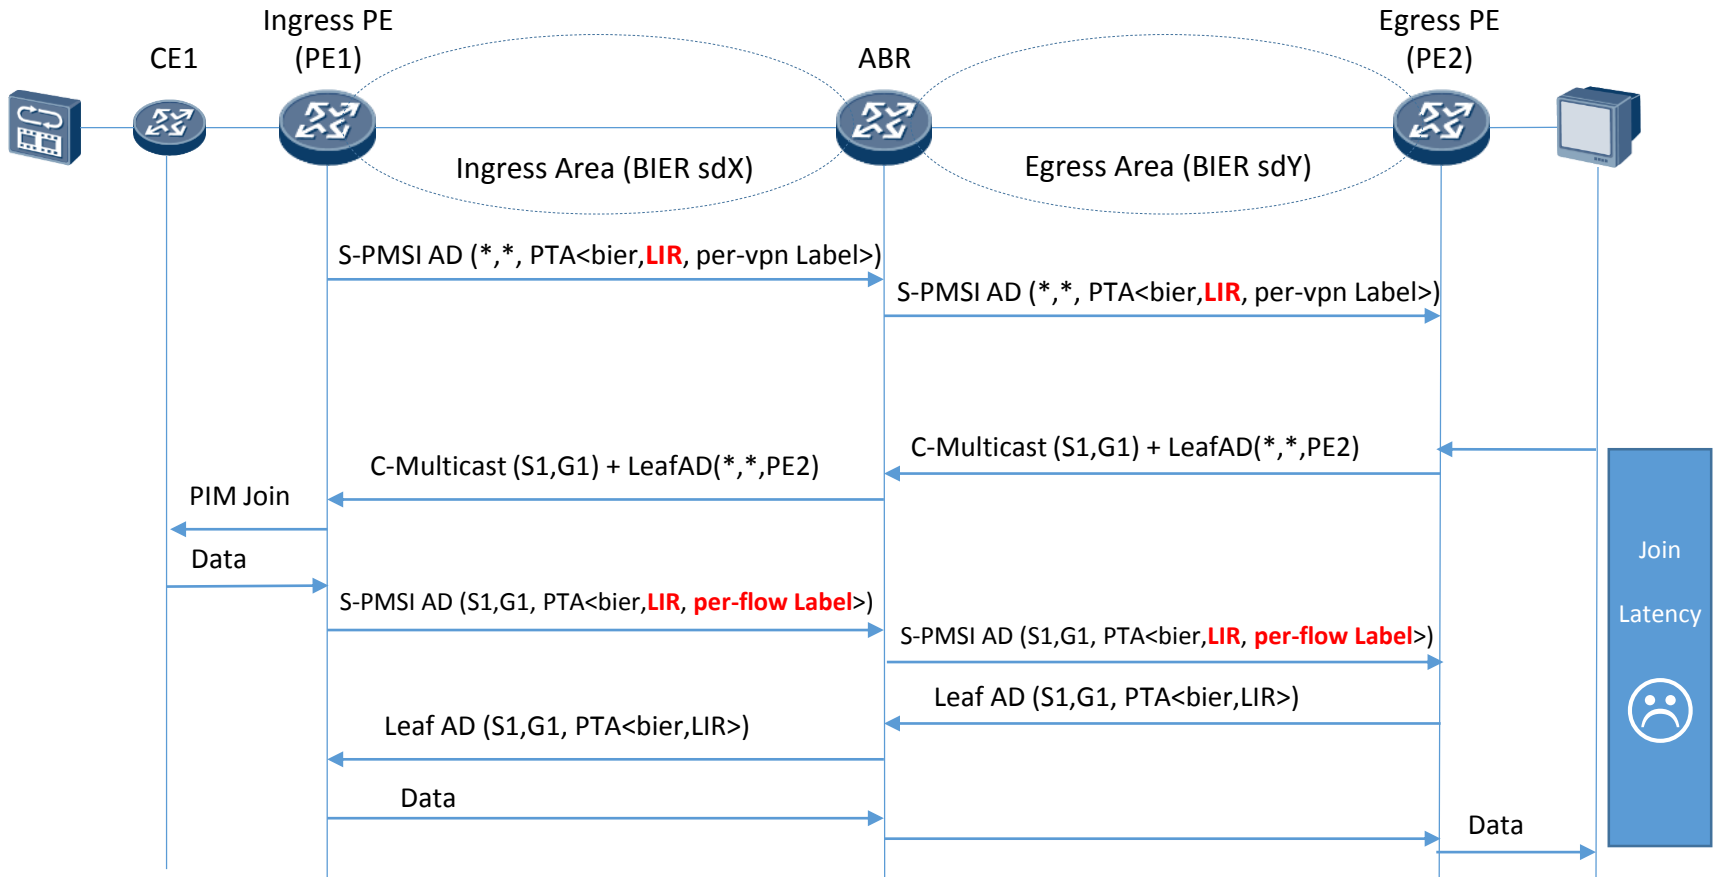
Liang Geng @China Mobile

Lei Wang @China Mobile

Mike McBride @Huawei

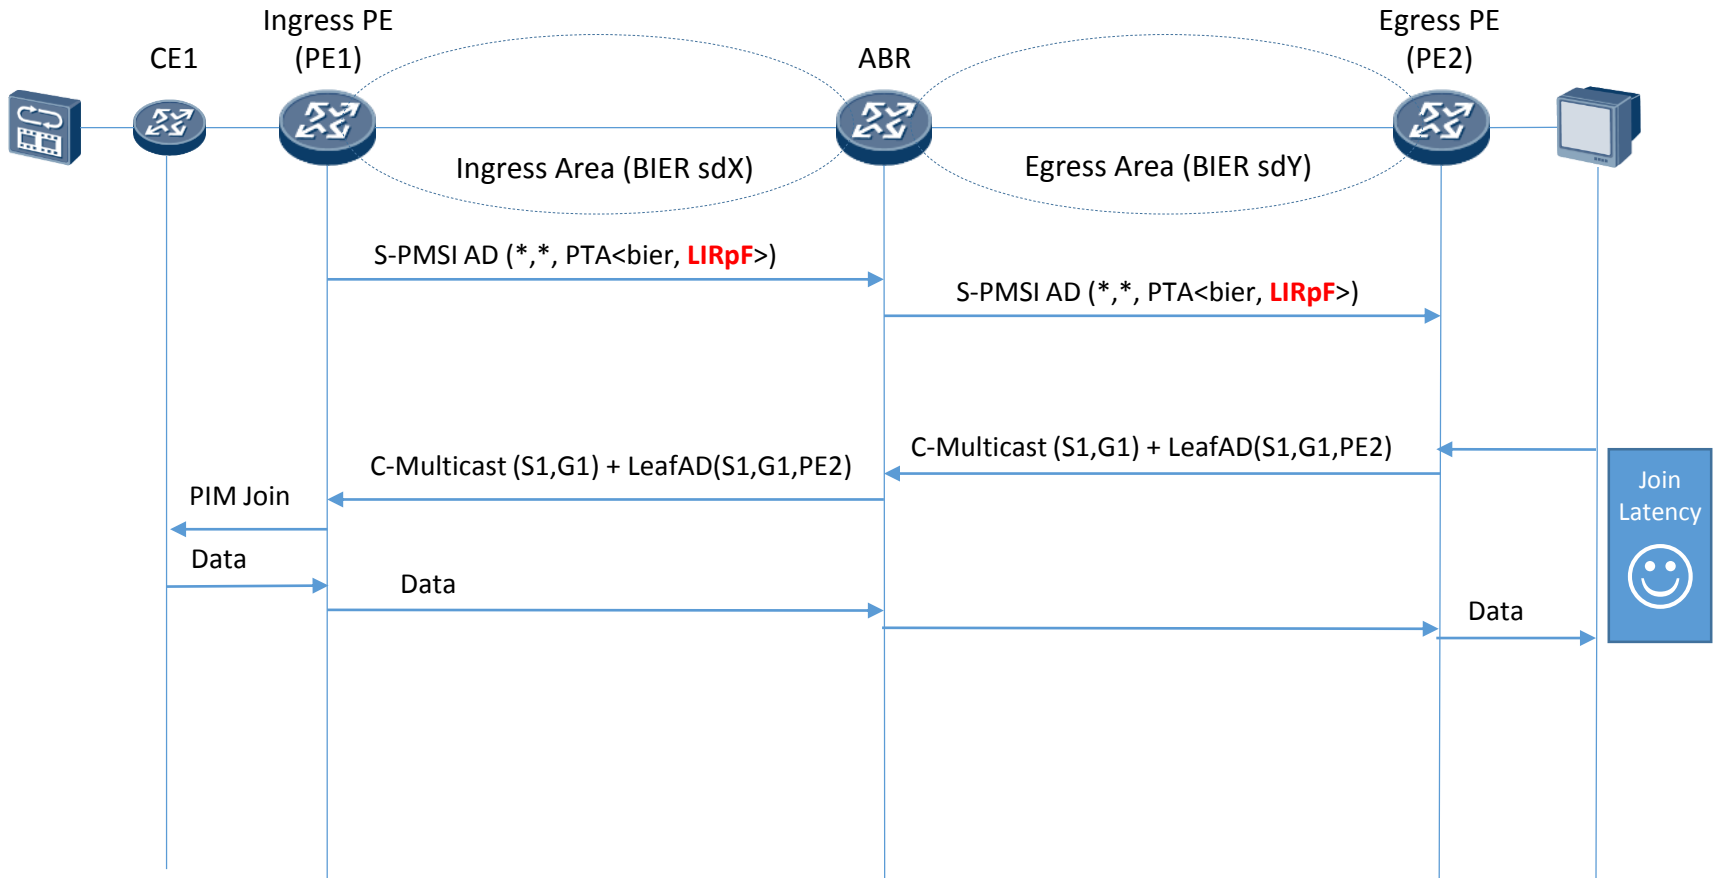
Gang Yan @Huawei

LIR explicit-tracking for Segmented BIER



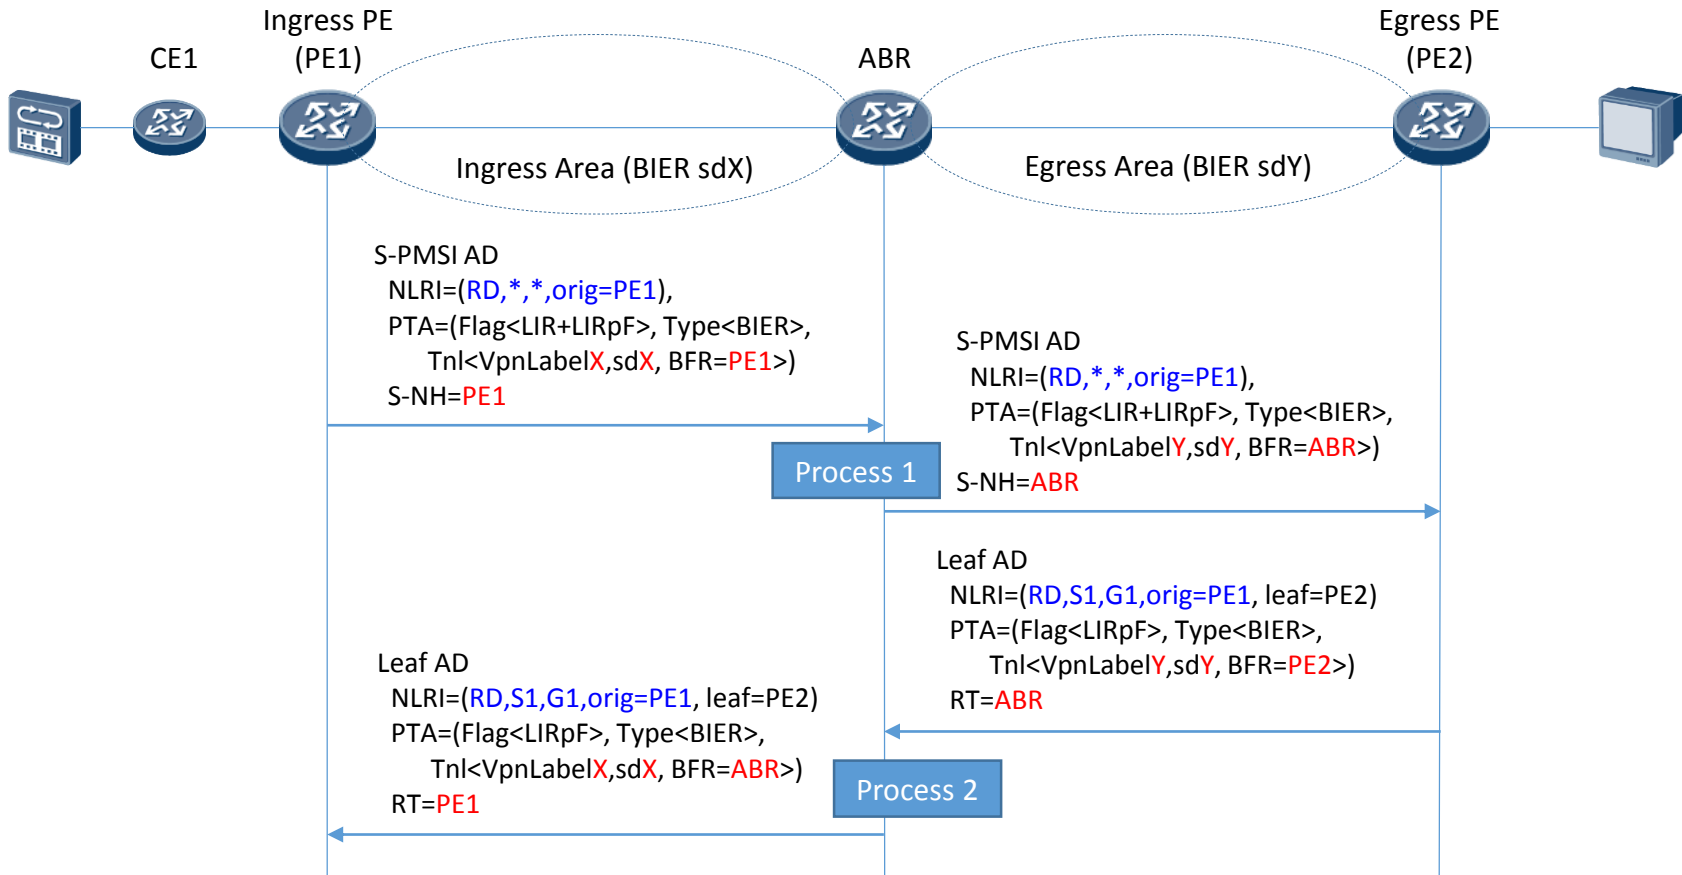
- Besides, the SPMSI(S,G) routes are 'flooded' to routers that even don't want.

LIR-pF explicit-tracking for Segmented BIER



- Accordingly, the unwanted SPMSI(S,G) routes are eliminated.
- The same benefit as Un-Segmented BIER MVPN.

Control Plane Process on ABR



Process 1: per-vpn info: FEC=(RD,PE1), upstream (X, PE1), downstream(Y, ABR)

Process 2: per-flow info: FEC=(RD,PE1,S1,G1), upstream(X,PE1), downstream(Y,ABR,PE2)

Per-flow state building

Process 2 : Build the control-plane state for per-flow on ABR


- ABR receive Leaf-AD, and form the **downstream state**: <RD, S1, G1, PE1>
<Leaf=PE2> (Attrs=BFR-id)
- ABR send Leaf-AD to PE1, and form the **upstream state**: <RD, S1, G1, PE1>
(umh=PE1)
- It is very similar to the PIM Join, or mLDP Mapping, which build a state driven by downstream join.
- Control-plane always need keeping a Per-flow state <RD, S1, G1, PE1>, including **upstream** and **downstream(s)** parts.
- The <RD, S1, G1, PE1> is an **Per-flow FEC** (I'd like to call it a BIER-FEC like RFC7524).
- The <RD, PE1> is an implicit VRF identifier for an ABR (call **Per-vpn FEC**).

Per-flow state for forwarding

Process 2: Build the forwarding state for per-flow on ABR

- **upstream state:** <RD, S1, G1, PE1> (umh=PE1)
- **downstream state:** <RD, S1, G1, PE1> <Leaf=PE2> (Attrs=BFR-id)
- Do a mapping of <RD,PE1> to <virtual VRF identifier> **locally on ABR**, then
 - Disposition Process : (BIER Label<of sd X >, BFIR-id<PE1>, VpnLabel X , virtual-VRF-identifier)
 - Re-Imposition Process : (virtual-VRF-identifier, S1, G1, sd Y , VpnLabel Y , BitString=PE2)
 - The Re-imposition process need an IP lookup ---- actually an MFIB lookup.

Think a little more



- Case 1 (the above diagram):
 - Ingress Area using P2MP ---- 'less specific replication', for example, SPMSI(*,*) for rep.
 - Egress Area using BIER ---- 'most specific replication', or say, 'per-flow specific replication'.
 - The LIR-pF explicit-tracking is still available ---- so do the LIR explicit-tracking.
- Case 2 (the opposite to the above diagram):
 - Ingress Area using BIER ---- 'per-flow specific replication' generally.
 - Egress Area using P2MP ---- 'less specific replication', for example, SPMSI(*,*) for rep.
 - Trade-off difficulty ----not only for LIR-pF, but also for LIR.
 - If Ingress Area(BIER) uses 'less specific replication' BIER ----not optimized replication.
 - If Egress Area(P2MP) uses 'most specific replication' Per-flow SPMSI ----possibly overloaded.

Thanks !