

Optimized Ingress Replication solution for EVPN

draft-rabadan-l2vpn-evpn-optimized-ir-01

Jorge Rabadan (ALU)
Senthil Sathappan (ALU)
Wim Henderickx (ALU)
Mukul Katiyar (Juniper)
Ravi Shekhar (Juniper)
Nischal Sheth (Juniper)
Wen Lin (Juniper)
Ali Sajassi (Cisco)
Aldrin Isaac (Bloomberg)
Mudassir Tufail (Citibank)

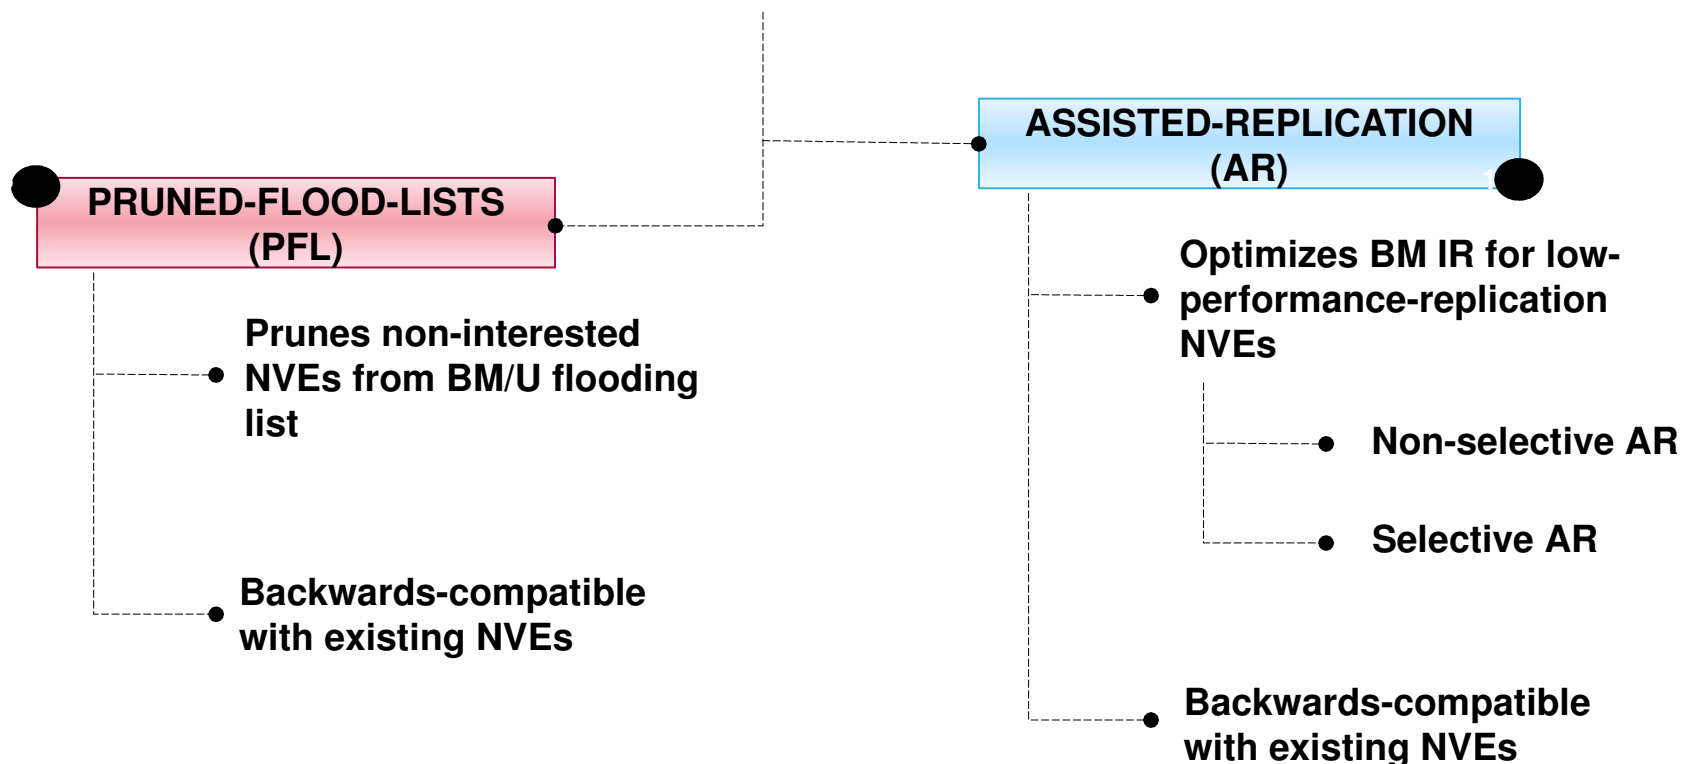
IETF 93, July 2015
Prague, Czech Republic

Why do we need optimized-IR in EVPN-overlay networks?

- Because PIM-based or other trees are not always possible and many operators want to use ingress replication (IR) since:
 - IR is simple
 - IR provides independence between underlay and overlay
- Because IR is inefficient and too expensive for software-based NVEs
 - Tx issue ->> too many copies of the same BUM packet can seriously affect the performance of certain NVEs
 - Rx issue ->> NVEs may not be interested in BM and/or Unknown unicast

The Optimized-IR solution for EVPN-overlay networks

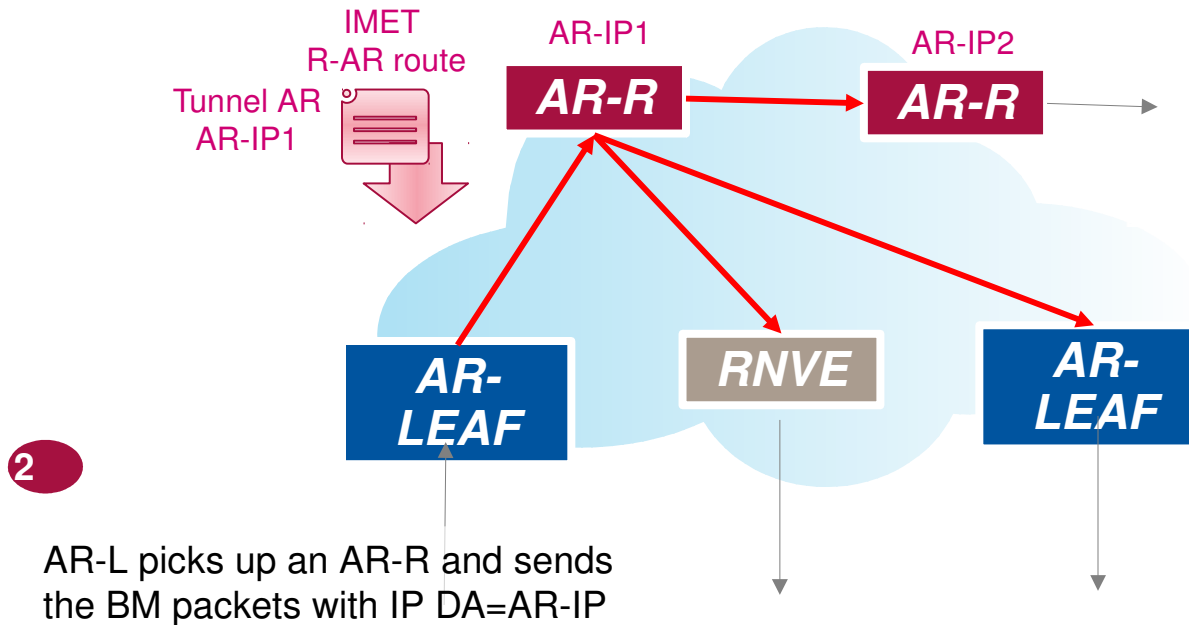
OPTIMIZED-IR



Non-selective AR (Assisted Replication)

1 AR-R advertise R-AR routes with tunnel type = AR, Tunnel id = AR-IP

3 AR-R does IR or AR based on the IP DA of the packets



2 AR-L picks up an AR-R and sends the BM packets with IP DA=AR-IP

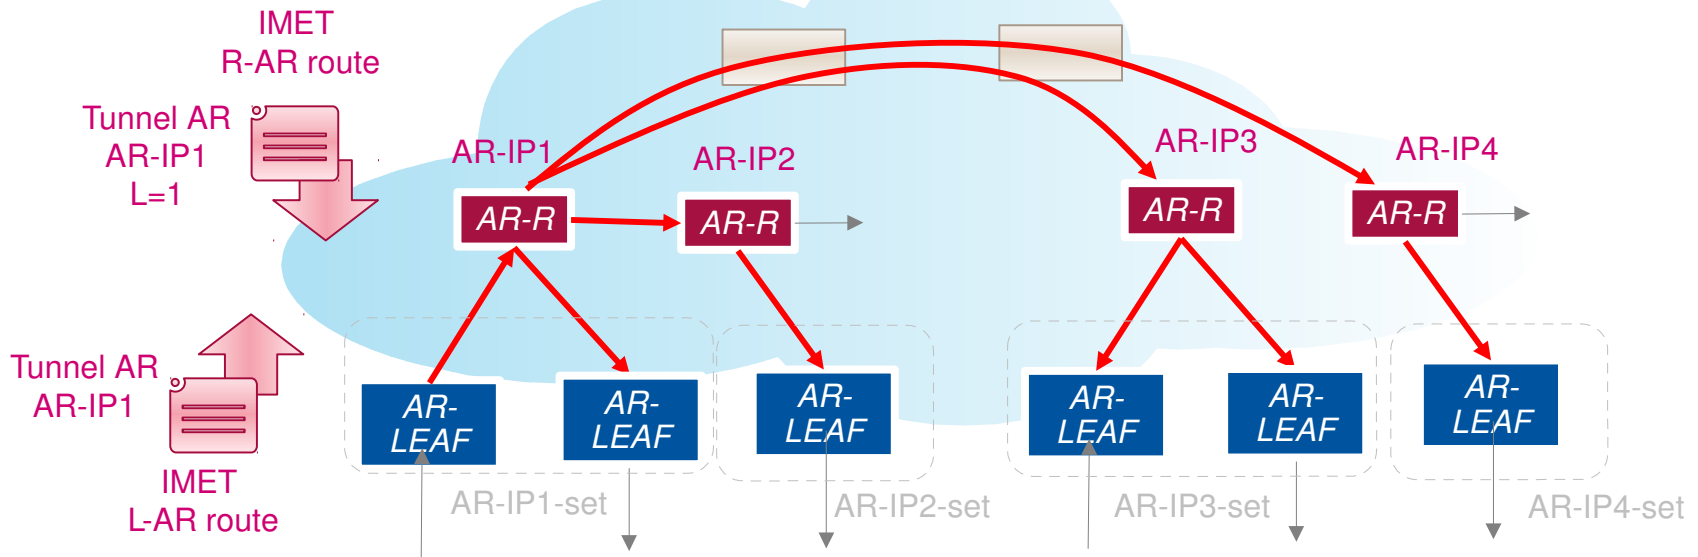
Backwards compatible: RNVE does IR
Unicast traffic follows regular MP2P tunnels

→ *BM on AC*
→ *BM on overlay tunnel*

Selective AR (Assisted Replication)

- 1 AR-R advertise R-AR routes with tunnel type = AR, Tunnel id = AR-IP L=1

- 3 Each AR-R builds its own AR-L-set and a remote AR-R-set
- Received BM with IP DA=AR-IP1 and IP SA=(IP_in_AR-L-set) is forwarded to AR-L-set AND AR-R-set



- 2 AR-L picks up one AR-R and send:
- L-AR route for the AR-IP1
 - BM packets with IP DA=AR-IP1

- 4 Received BM with IP DA=AR-IP3 and IP SA != (IP_in_AR-L-set) is forwarded to AR-L-set only

Changes in this rev

- AR procedures for single-IP AR-Replicators
 - For AR-R that support a single-IP and still need IR and AR forwarding
- AR procedures compatible with EVPN split-horizon for non-MPLS tunnels
 - Non-MPLS EVPN multi-homing split-horizon relies on tunnel IP SA checking
 - The AR-R may decide not to modify the tunnel IP SA

Conclusions and next steps

- Optimized-IR allows the efficient distribution of BM without dependencies between the underlay and overlay networks
- The authors would like to request WG adoption