

# Optimized Ingress Replication solution for EVPN

---

draft-rabadan-l2vpn-evpn-optimized-ir-00

Jorge Rabadan  
Senthil Sathappan  
Wim Henderickx  
Mukul Katiyar  
Ravi Shekhar  
Nischal Sheth  
Wen Lin

IETF 90, July 2014  
Toronto, Canada

# When EVPN is used as the control plane for NVO networks, what options do we have to transport BUM?

## Ingress Replication

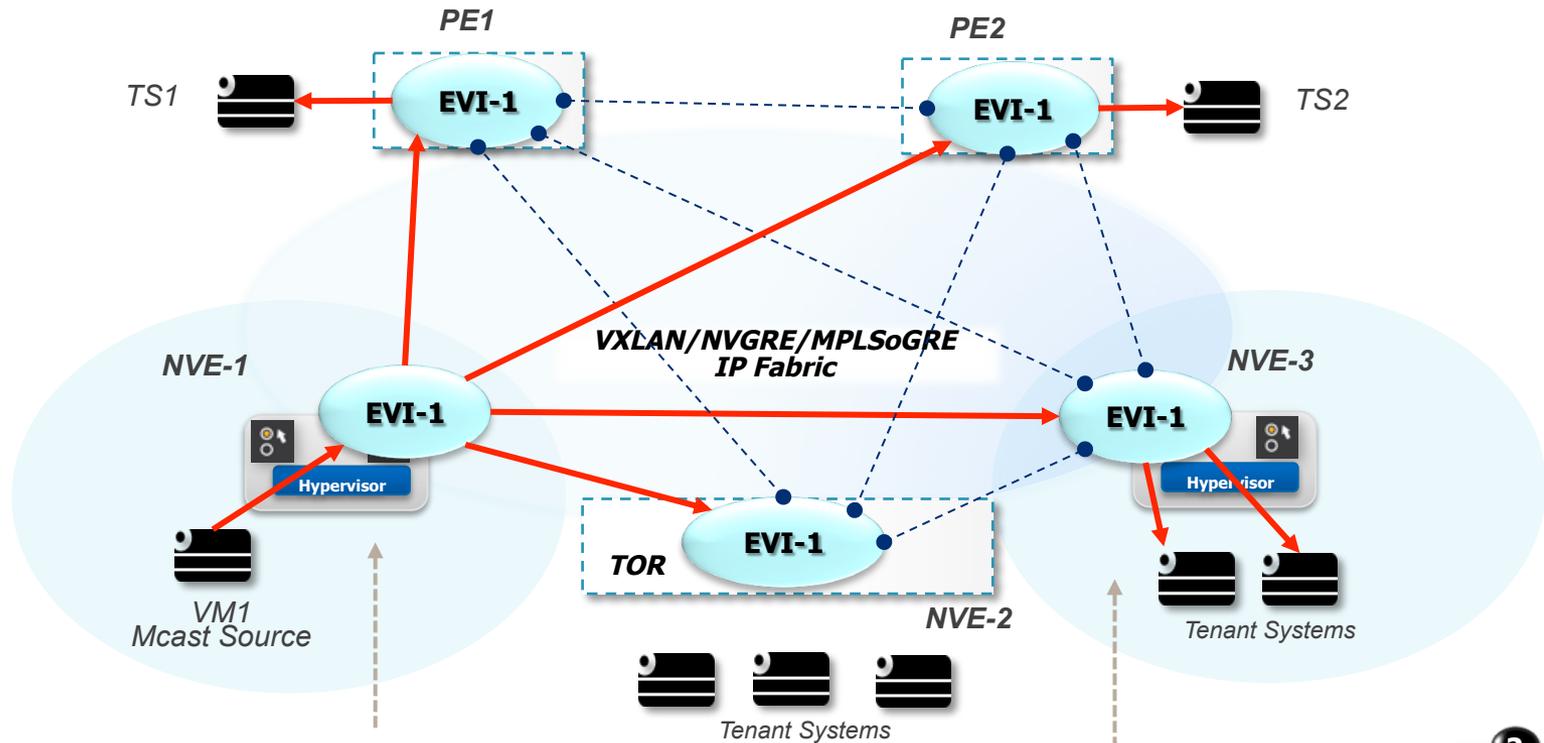
- No dependencies on the underlay network core
- Enough if the tenant does not use multicast applications (EVPN optimizes B/U)
- Inefficient for highly demanding multicast applications

## PIM-based Trees

- Dependencies on the underlay network core
- Creates states in the core
- Efficient

IR is the preferred solution when independency between overlay and underlay core is required

# Problem statement: IR is inefficient and may be too expensive for low-performance NVEs



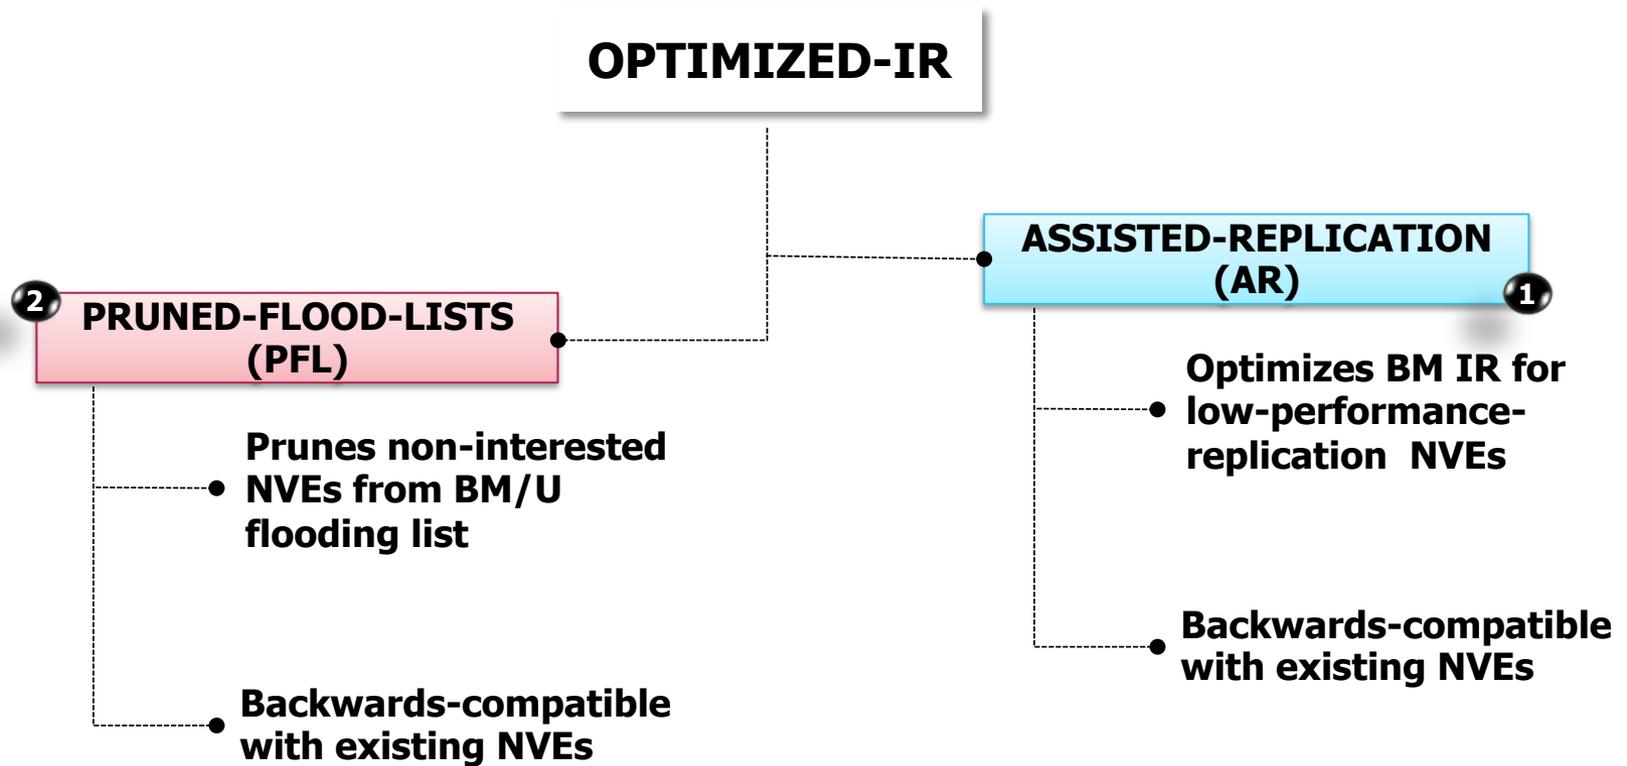
1

**Applications can be seriously affected in low-performance-replication NVEs**

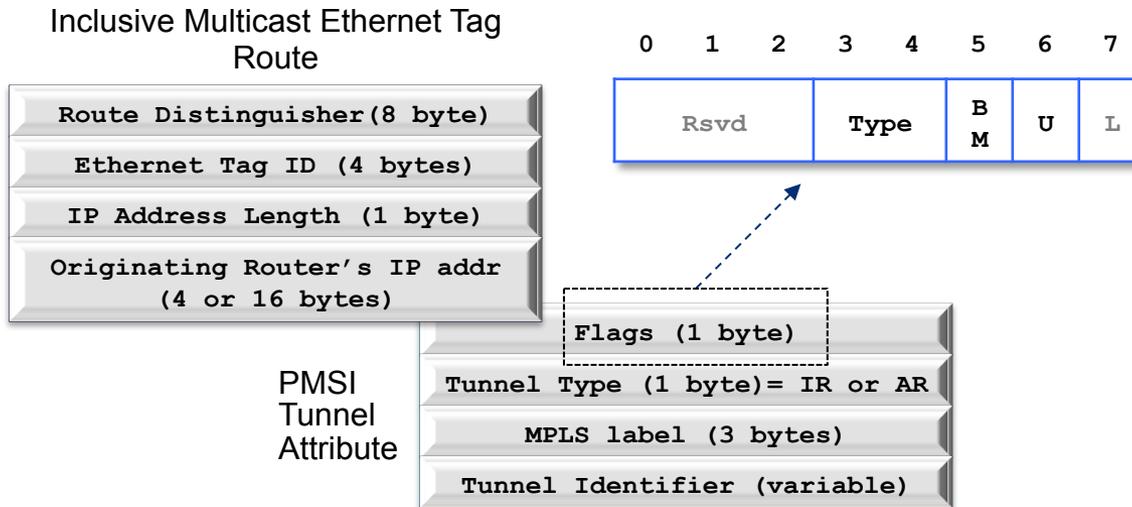
2

**NVEs might not be interested in all B/U/M traffic**

# Solution: Optimized-IR for EVPN



# EVPN BGP attributes for optimized-IR



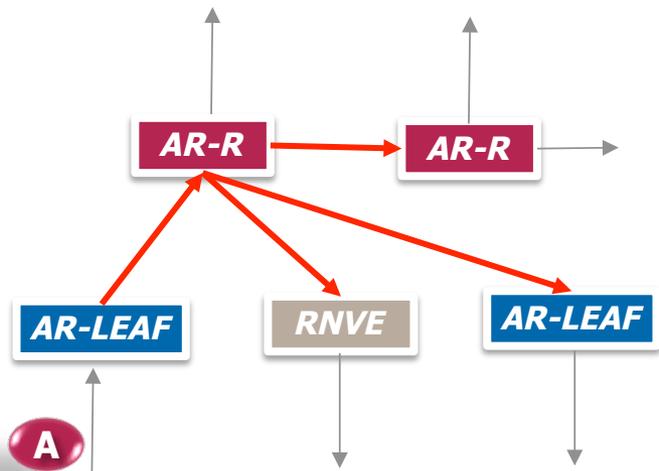
## New proposal

- AR type field
  - 00 = RNVE (non-AR support)
  - 01 = AR-Replicator
  - 10 = AR-Leaf
- PFL flags
  - BM (Broadcast and Multicast)
  - U (Unknown)
    - 1 = "prune me"
    - 0 = regular behavior

<b>AR ROLE</b>	<b>FUNCTION</b>	<b>Inc. Mcast Route information</b>
AR-REPLICATOR	Assists AR-leafs	IR incl. mcast route (IR IP) AR incl. mcast route (AR IP, tunnel=AR, T=1)
AR-LEAF	Sends BM only to AR-Rep	IR incl. mcast route (IR IP, T=2)
RNVE	Non-AR support	IR incl. mcast route (IR IP)

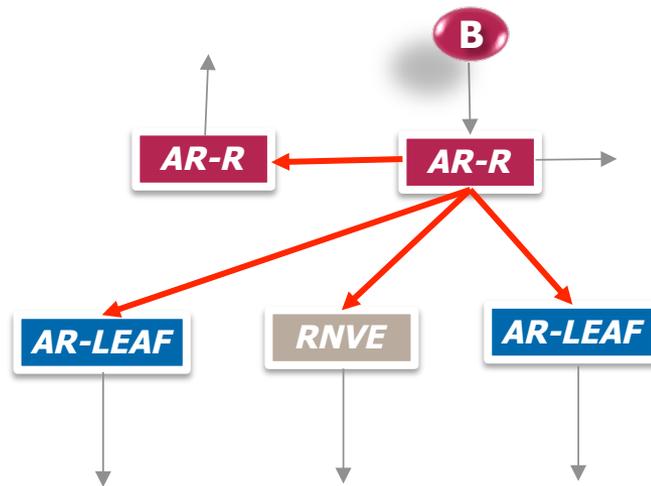
# Broadcast/Multicast forwarding

## ASSISTED-REPLICATION (AR)



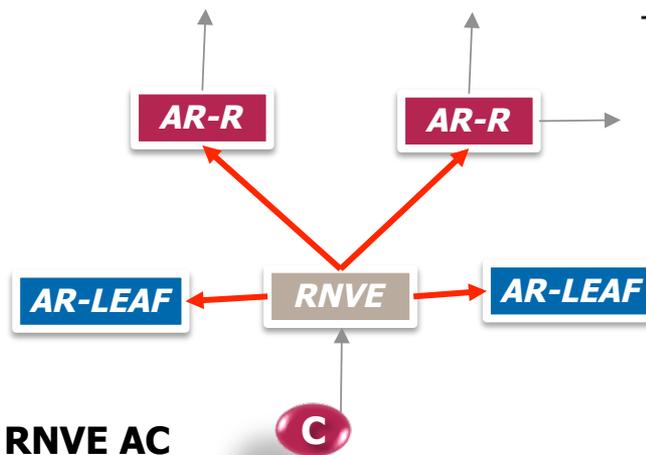
### BM received on AR-LEAF AC

- AR-R selected by AR-LEAF
- AR IP used as IP DA
- AR-R forwards based on IP lookup
- Unknown unicast follows IR to avoid packet re-ordering



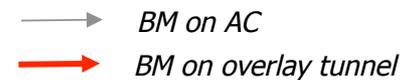
### BM received on AR-R AC

- Normal IR



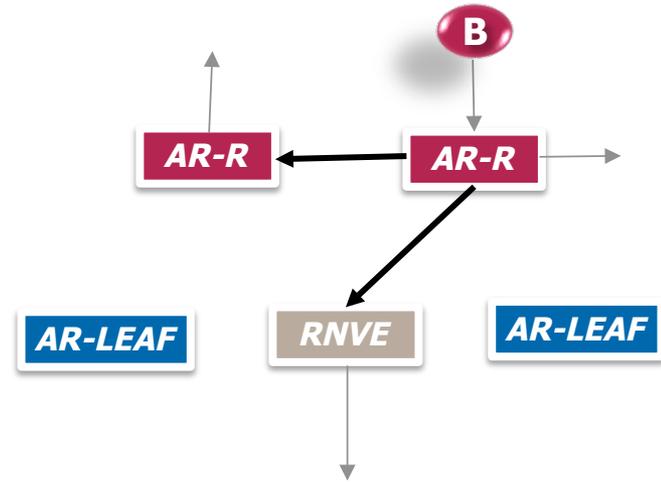
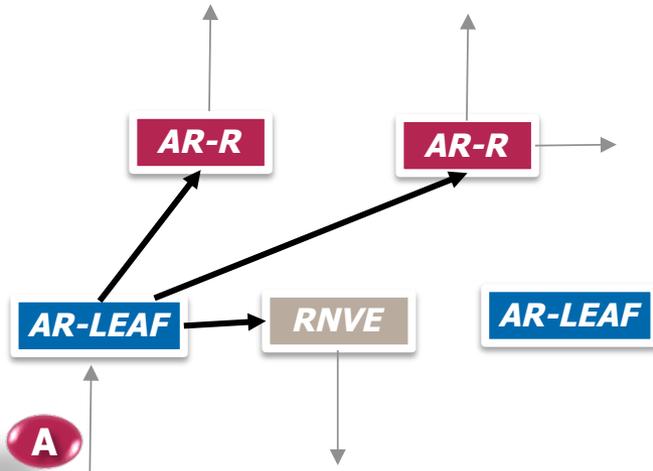
### BM received on RNVE AC

- Normal IR

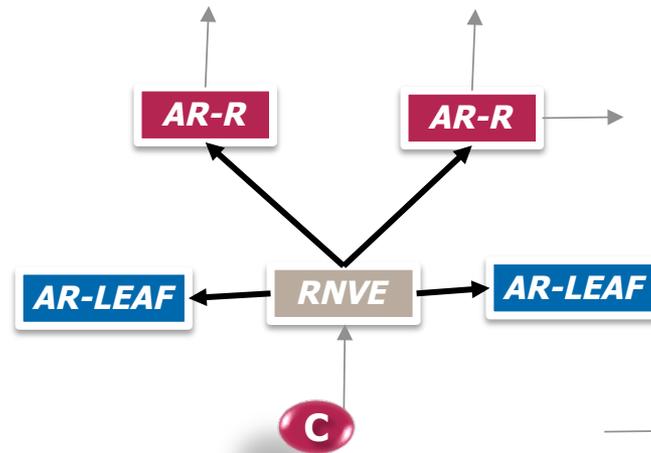


# Use of PFL in optimized-IR

## PRUNED-FLOOD-LISTS (PFL)



- PFL can be used independently of AR
- PFL flags are used in IR incl. Mcast routes
- In this example NVEs in hypervisors are not interested in receiving unknown unicast



 *Unknown on AC*  
 *Unknown on overlay tunnel*

# Conclusions and next steps

- Ingress replication is the preferred way of sending BUM in an EVPN-overlay network.
- Optimized-IR proposes a solution to improve the performance in NVEs with poor replication capabilities (AR) and avoids the replication of BUM to uninterested NVEs (PFL).
- The authors request feedback from the WG