

Extended Optimized Ingress Replication for EVPN

draft-wsv-bess-extended-evpn-optimized-ir-01

Wen Lin (Juniper)

Selvakumar Sivaraj (Juniper)

Vishal Garg (Juniper)

Jorge Rabadan (Nokia)

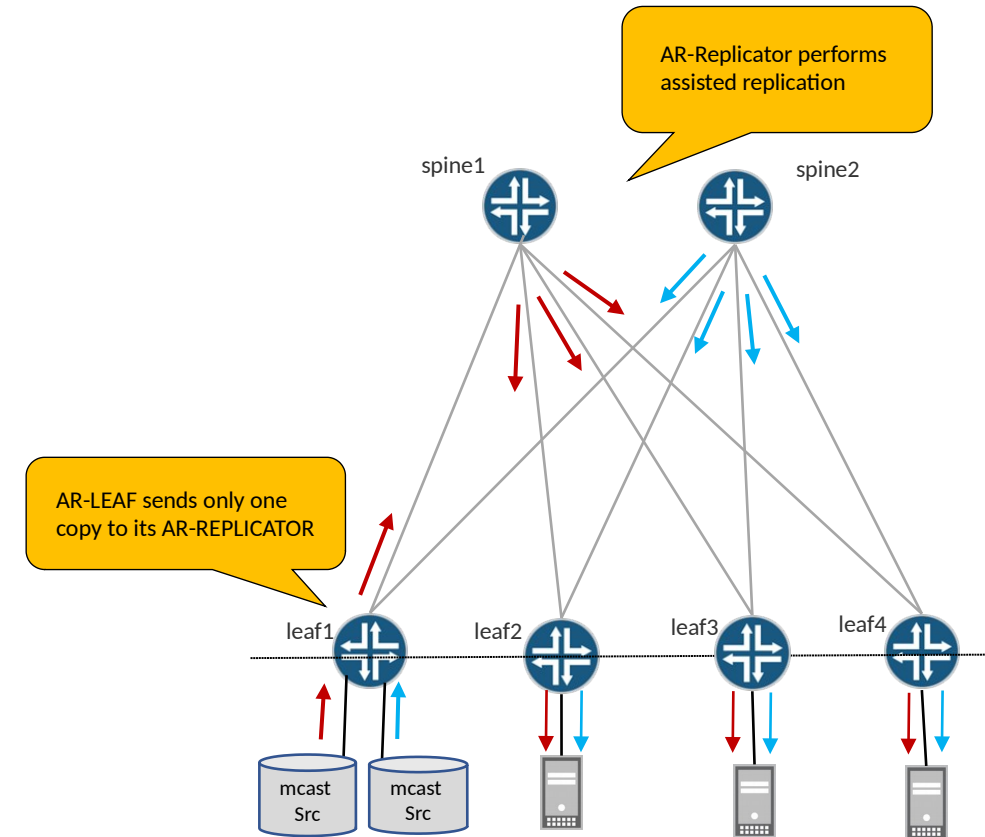
Optimized Ingress Replication for EVPN

- Assisted Replication

EVPN Optimized Ingress Replication:

draft-ietf-bess-evpn-optimized-ir-06

- An optimized ingress replication using EVPN overlay technology for NVO networks
- Efficient multicast delivery, it avoids sending multiple copies of the same multicast flow over the same uplink
- Alleviate the burden of an AR-LEAF, and save uplink bandwidth between leaf to spine



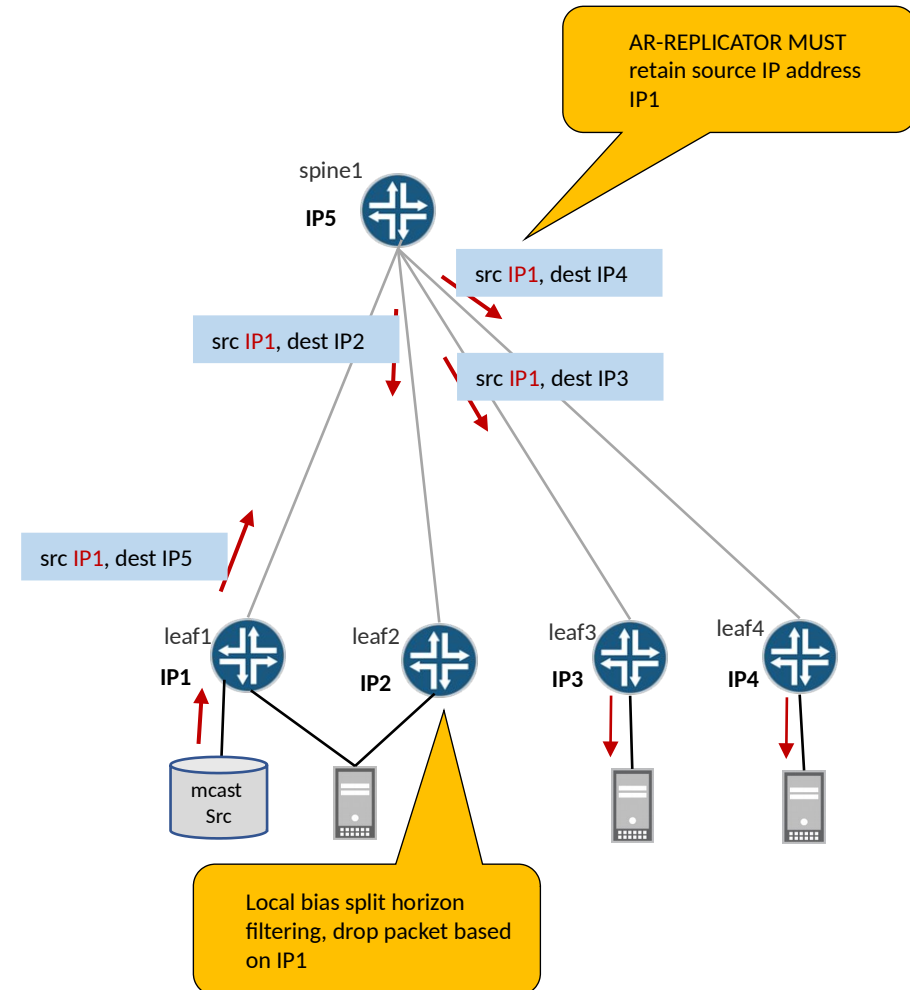
Multihoming Support for Assisted Replication

AR-LEAF Multihoming support requires

- AR-LEAF uses local-bias based split horizon filtering method
- For assisted replication, an AR-Replicator must retain the source IP address for the traffic it receives from its AR-replication tunnel – i.e. an AR-Replicator must use its ingress AR-LEAF's IP address as its source IP address

Challenges

- AR-Replicator may be incapable to retain the source IP address
- Or implementation complexity



Extended Optimized Ingress Replication for EVPN

- draft-wsv-bess-extended-evpn-optimized-ir-01


Extended EVPN Optimized Ingress Replication

- Support both multihoming split horizon methods
 - local-bias/source IP per RFC8365
 - ESI label per RFC 7432
- Simplify the process to support AR-LEAF multi-homing for EVPN in NVO networks
 - No need for AR-REPLICATOR to retain its ingress AR-LEAF's IP address or include ESI label

Extended Procedure for AR-REPLICATOR

An AR-REPLICATOR supports the following extended procedure - Extended-MH AR-REPLICATOR:

- Originates REPLICATOR-AR route as usual, and attach an EVPN Multicast Flag extended community with Extended-MH-AR flag set:



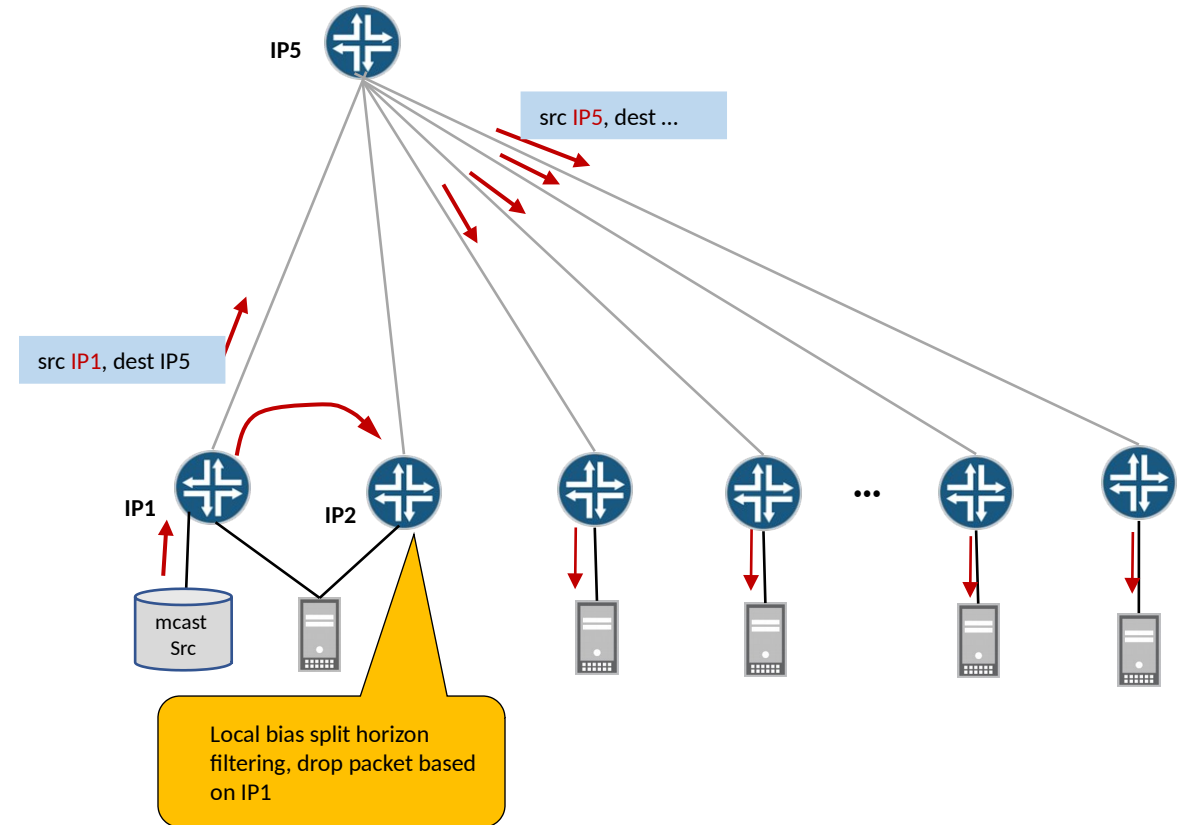
Flags Octet Bit3: 3 -> Extended-MH-AR

- No need to retain source IP address or include ESI label for its AR-LEAF
- AR-REPLICATOR only performs assisted replication to other NVEs that is not multihomed with the ingress AR-LEAF
 - Per aliasing function, an AR-REPLICATOR has the knowledge of an AR-LEAF's multihoming peer(s).

Extended Procedure for AR-LEAF

An AR-LEAF supports the following extended procedure

- Detect its AR-REPLICATOR's capability based on EVPN Multicast Flags extended community
- Send one copy to AR-REPLICATOR as usual
- Does normal IR to its multihomed peer if AR-REPLICATOR is unable to retain its source IP or include its ESI label



Conclusion and Next Steps

- Extended Optimized Ingress Replication support AR-LEAF multihoming and overcome hardware limitation or implementation complexity for AR-REPLICATOR
- Allows split horizon filtering method based on ESI label for MPLSoX encapsulation
- Alleviate the burden of an AR-LEAF, and save uplink bandwidth
- Would like to seek feedback from WG