

EVPN Inter-subnet Multicast Forwarding

draft-lin-bess-evpn-irb-mcast-02

Wen Lin

Jeffrey Zhang

John Drake

Jorge Rabadan


IETF 95, April 2016

Buenos Aires, Argentina

Inter-subnet Unicast Traffic for Distributed GW

In the distributed GW approach, inter-subnet traffic for known unicast traffic:

- CE1 sends traffic to its local GW/IRB in the source subnet1 at L2
- On local PE1, traffic is L3 routed through IRB to the destination subnet2
- In BD2, PE1 forwards traffic to the destination CE/host at L2




Inter-subnet Multicast Traffic for Distributed GW

- Multicast Traffic is flooded to receivers in source subnet1 following the EVPN BUM procedure
- On PE3, multicast traffic is L3 routed to IRB in the subnet2 based on the regular PIM procedure
- Multicast traffic is flooded in the subnet2 following the EVPN BUM procedure


Resource Inefficient:

- Flooded in every subnet/BD; PE2 receives the same multicast flow from the core in EACH subnet.
- Receiver1 receives the multicast traffic from PE3 - hairpinning.



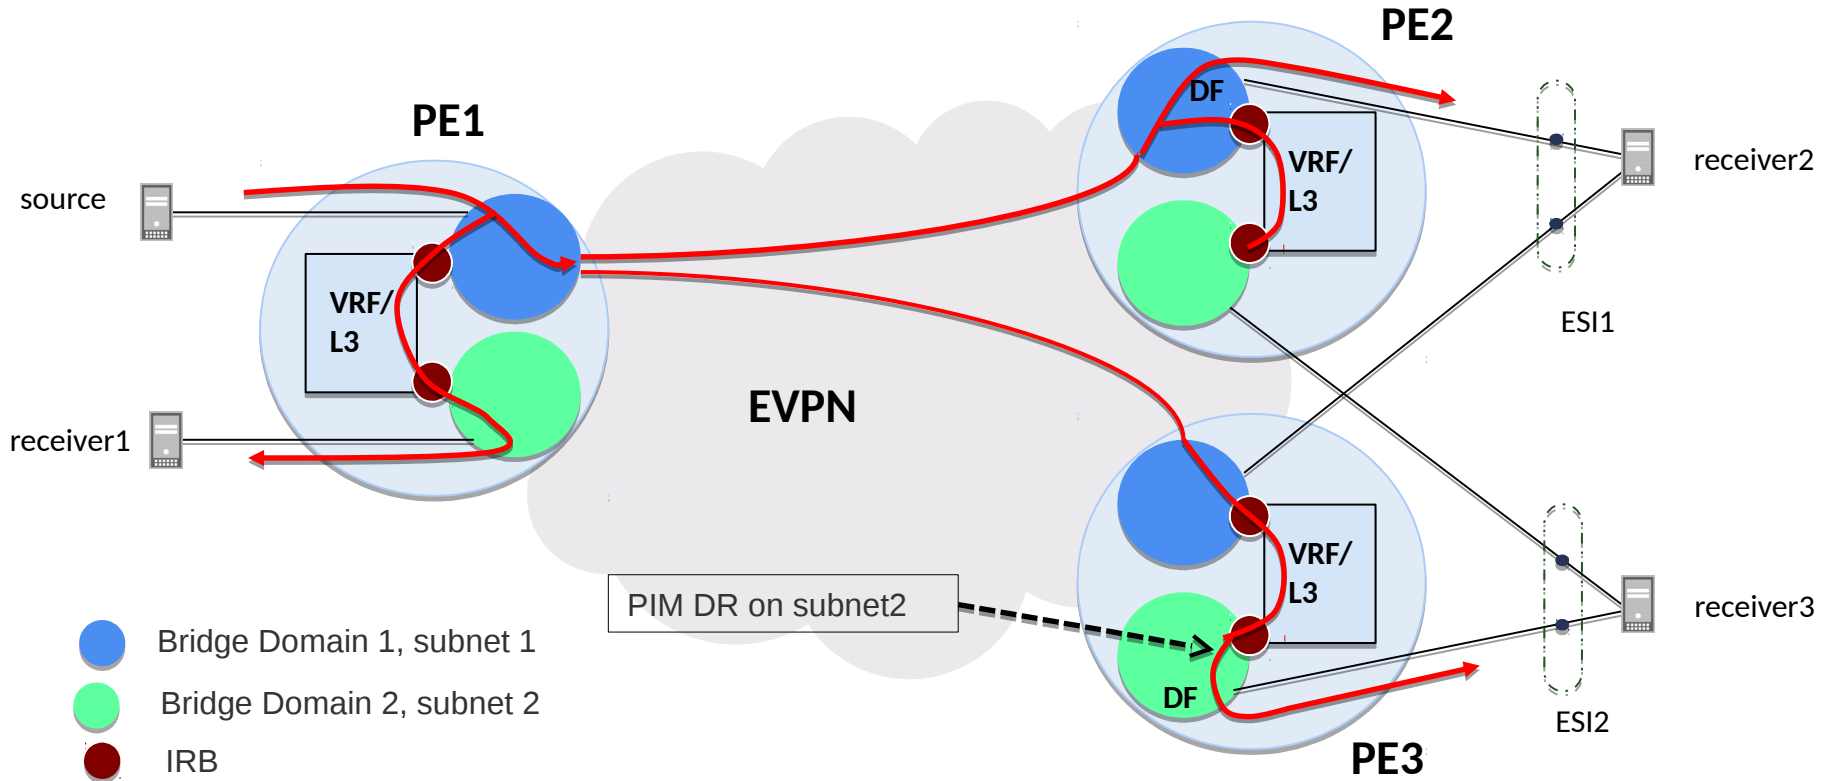
Solution to Hairpinning, Receiving Multiple Copies of the Same Multicast Traffic from the Core

- On source subnet, sends the multicast traffic to the IRB interface on other subnets that have receivers, regardless whether the gateway is DR for that subnet or not.
- On the IRB interface, if there is Membership Report from one of its ACs, sends PIM joins towards the RP or source regardless if it is DR or not.
- On non-source subnet, multicast data traffic sent out of the IRB interface is forwarded to local ACs ONLY and NOT send back to the core.



Receiver multi-homed to EVPN PE

Only DF for the ES delivers the multicast data traffic to the receiver.



Conclusion

Distributed GW approach:

- Source Subnet: no change, EVPN BUM procedure
- Non-Source subnet: L3 routed locally and forward only to local ACs regardless PIM DR state

Save bandwidth in the core:

- flooding of multicast traffic in the core only happens in the source subnet and not in other subnets.
- Avoid hairpinning of multicast traffic for local receiver on other subnets if the PE is not the DR for the subnet.

Next Step

Work in Progress for PE does not have IRB or not have the source BD.

Seek and address comments from the WG.

Will seek WG adoption afterwards