

EVPN Multicast Forwarding for EVPN to EVPN Gateways

draft-rabnic-bess-evpn-mcast-eeg-00

Jorge Rabadan (Nokia)

Olivier Dornon (Nokia)

Vinod Prabhu (Nokia)

Alex Nichol (Arista)

Jeffrey Zhang (Juniper)

Wen Lin (Juniper)

IETF115, Nov 2022

London

Agenda

1. Introduction
2. Procedures
3. Next-steps

Multicast in Service Gateways

What are we missing?

EVPN Layer 2 Service Gateways

- RFC9251 defines IGMP/MLD proxy in an EVPN Broadcast Domain (single domain)
- RFC9014 defines the EVPN Layer 2 interconnect (Integrated model) where the Service Gateway connects two or more EVPN domains. Behavior is specified for unicast and BUM, but not IP multicast.
- What we are missing:

How do we interconnect two (or more) EVPN domains where RFC9014 Gateways exist and the PEs in both domains run RFC9251 procedures?

EVPN Layer 3 Service Gateways

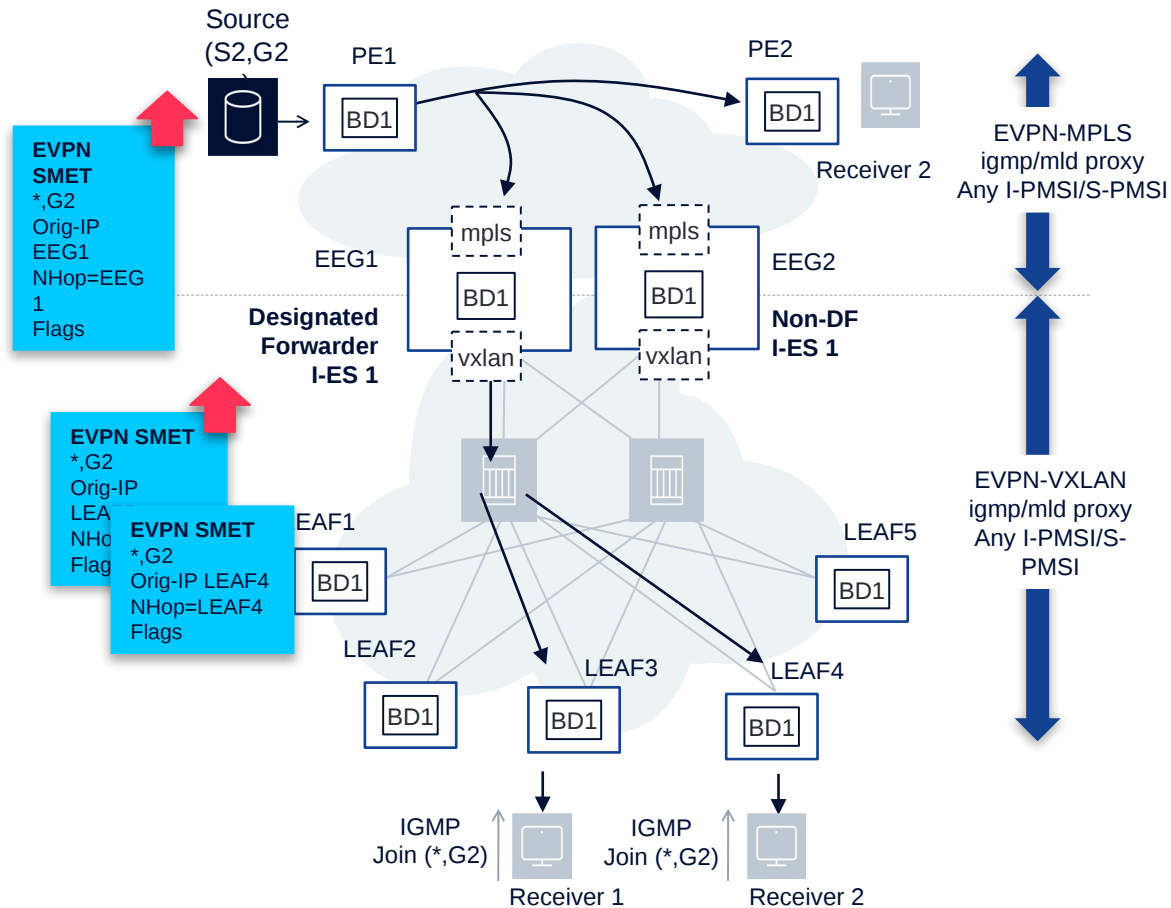
- draft-ietf-bess-evpn-ipvpn-interworking defines Service Gateway procedures to interconnect EVPN L3 domains for unicast traffic
- draft-ietf-bess-irb-mcast defines Service Gateway procedures to interconnect EVPN OISM domains to PIM or MVPN domains
- What we are missing:

How do we interconnect two (or more) EVPN L3 domains where draft-ietf-bess-evpn-ipvpn-interworking Gateways exist and the PEs in both domains run draft-ietf-bess-irb-mcast OISM procedures?

Layer 2 Multicast EVPN-to-EVPN Gateway

Or a Layer 2 EEG

Example for EVPN-MPLS to EVPN-VXLAN gateways
Valid for any combination of transport at both domains



EEGs are Gateways as in RFC9014 (integrated model) plus additional extensions for IP Multicast

Control Plane

- SMETs imported by all EEGs (no need for synch routes)
- SMETs are proxy'ed by EEGs between EVPN domains
- A minimum set of proxy SMET routes to attract the traffic for a group and version is distributed between domains
- I-ES DF proxies SMET routes (optionally non-DF as well, based on fast convergence vs BW utilization trade-off)
- EEGs add D-PATH on SMET routes to avoid loops

Data Plane

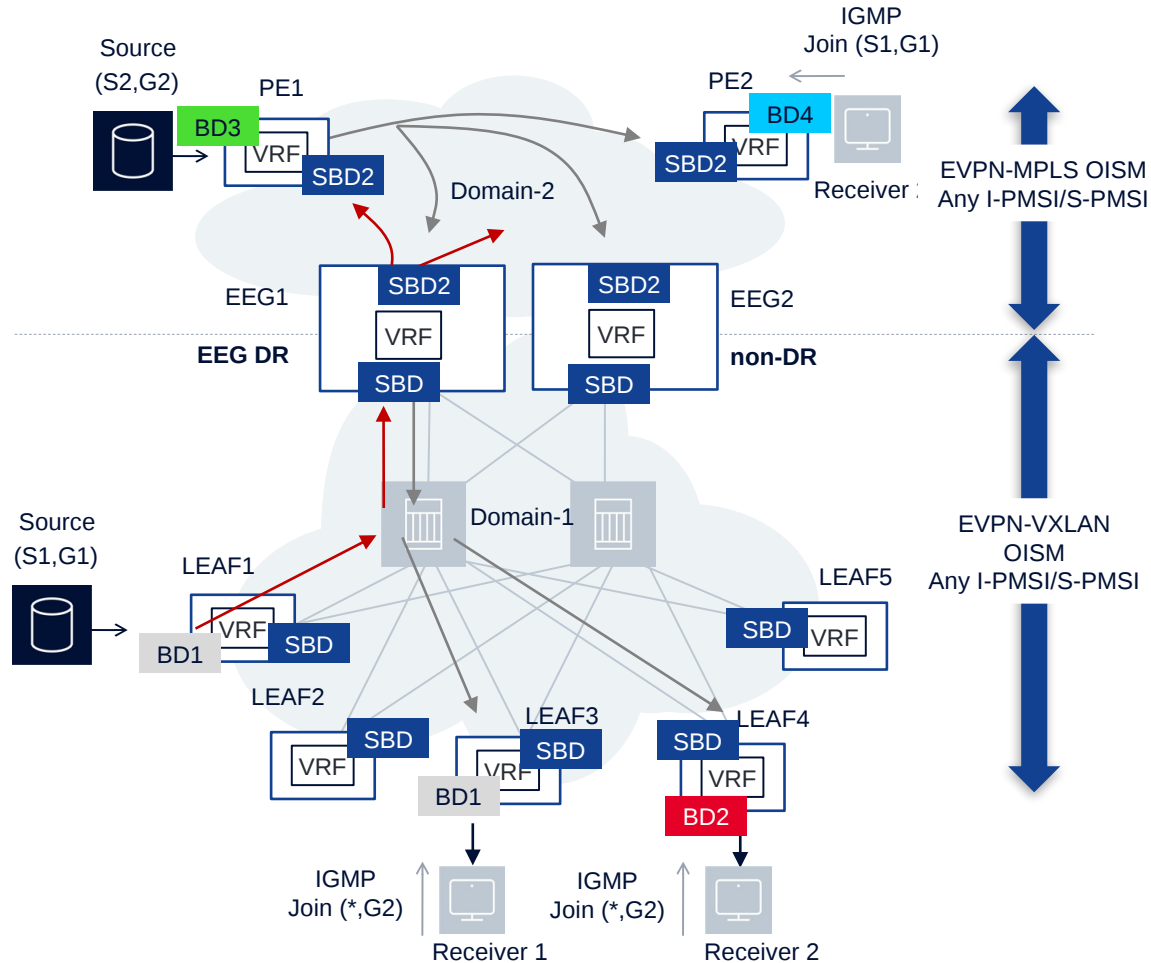
- Non-DF blocks ingress/egress mcast (as per RFC9014)

EEG local receivers and sources are supported

Layer 3 Multicast EVPN-to-EVPN Gateway

Or a Layer 3 EEG

Example for EVPN-MPLS to EVPN-VXLAN gateways
Valid for any combination of transport at both domains



EEGs are Gateways as in I-D.ietf-bess-evpn-ipvpn-interworking, plus additional extensions for IP Multicast

Control Plane

- SMETs imported by all EEGs and proxy'ed / redistributed
- A minimum set of proxy SMET routes to attract the traffic for a group and version is distributed between domains
- A new flag "EEG" added to the SBD-IMET routes so that a DR election is run among all EEGs
- The DR is responsible of redistributing SMET routes between domains (optionally the non-DR as well, based on fast convergence vs BW utilization trade-off)
- The non-DR removes the SBD IRB from the L3 OIF list

EEG local receivers and sources are supported

Next Steps

draft-rabnic-bess-evpn-mcast-eeg-00

The authors would like to request feedback from the Working Group

Thank you