

draft-boutros-l2vpn-vxlan-evpn-02.txt

Sami Boutros (Cisco), Ali Sajassi (Cisco), Samer Salam (Cisco), Dennis Cai (Cisco), Samir Thoria (Cisco), John Drake (Juniper)

July 2013

Berlin

Changes from Rev01

- Added All-active multi-homing scenario
 - Sec. 2.2 – Requirements
 - Sec. 3.1 – Redundancy procedure
 - Sec. 4.6.2 – Multicast trees stitching
- Ether Tag setting in EVPN MAC route
 - VNI-based service: Ethernet tag = 0
 - VNI-aware bundle mode: Ethernet tag = VNI

Interconnecting VXLAN Networks w/ EVPN

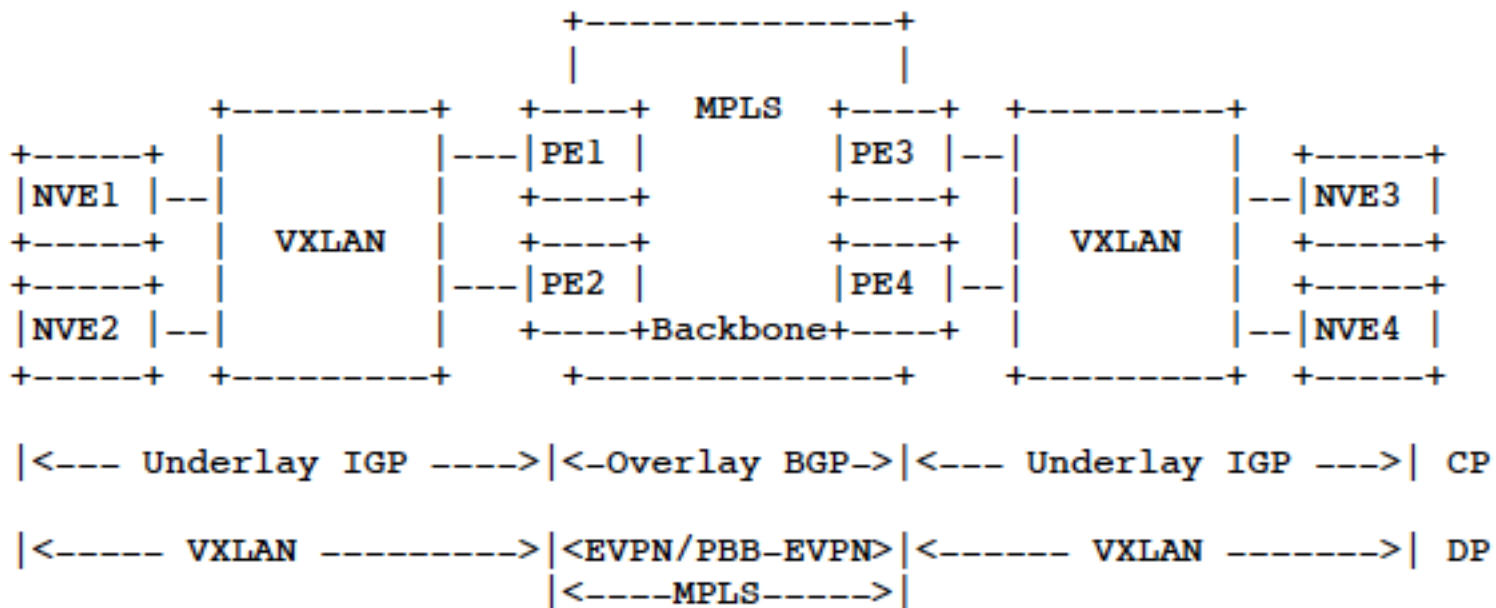


Figure 1: Interconnecting VXLAN Networks with VXLAN-EVPN

Redundancy & All-Active Multi-homing

- If each EVPN PE uses a unique VTEP address, then all-active multi-homing results in MAC flip/flopping at the NVEs
- In order to address this issue, anycast address is used for all PE devices multi-homed to a given VXLAN network (which is treated as an multi-homed site)
- Two advantages:
 - Prevents any flip/flopping at the NVEs
 - Enables load-balancing via ECMP for DCI traffic among multi-homed PEs

Multicast Tree Stitching w/ per VNI Load Balancing

- PIM Bidir is used in this rev
 - For both its scalability advantage in multicast states and
 - Alleviating any special processing for RPF check
- There will be number of trees built in the core – e.g., one per spine node
- NVEs can select any of the trees for the purpose of load balancing
- Each PE joins all the trees

Multicast Tree Stitching w/ per VNI Load Balancing – Cont.

- EVPN DF election procedure is used to decide which PE should be the DF for a given VNI
- Only a single PE allows for BUM traffic to pass through between VXLAN and EVPN networks for the VNIs for which it is a DF
- VNI is used to identify BD
- Host MAC learning is performed in data-plane along with its associated VTEP address

Next Steps

- Solicit more inputs