

VXLAN DCI Using EVPN
draft-boutros-l2vpn-vxlan-evpn-01.txt

Sami Boutros

Ali Sajassi

Samer Salam

Dennis Cai

IETF 86, March 2013

Orlando, Florida

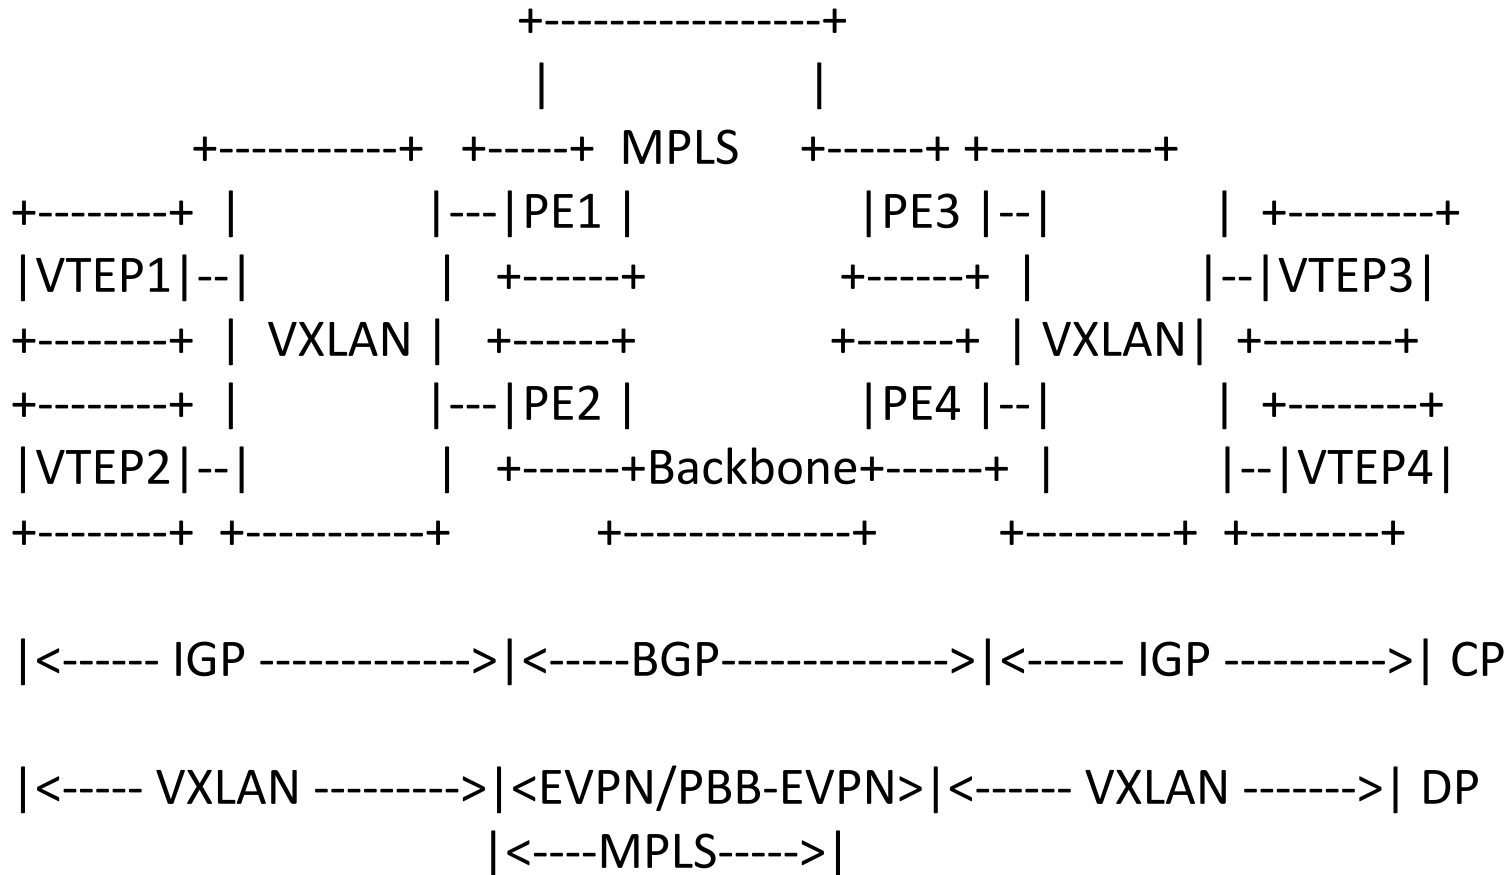
VXLAN DCI Using EVPN

- This is to provide intra-subnet connectivity at Layer 2 and control-plane separation among the interconnected VXLAN or NVGRE networks over the EVPN MPLS/IP network.
 - The scope is limited to data plane learning in this document.

Requirements

- Control Plane Separation among VXLAN/NVGRE Networks.
- Layer 2 Extension over the MPLS/IP Network for intra-subnet connectivity.
- Support for IRB for both inter-subnet routing and intra-subnet bridging for a given VNI/VSID.

Solution overview



Legend: CP = Control Plane View

DP = Data Plane View

EVPN Routes

- BGP MAC/B-MAC Advertisement Route for EVPN/PBB-EVPN
- Ethernet Auto-Discovery Route for EVPN
- Per VPN Route Targets
- Inclusive Multicast Route to distribute the VNI information over the MPLS network.
 - Discovery of the PEs participating in a given VNI.
 - Stitching of the IP multicast trees, local to each VXLAN site, with the Label Switched Multicast (LSM) trees of the MPLS network.

Handling Unicast traffic

- Host MAC addresses will be learnt in data plane from the VXLAN network.
- Host MAC addresses will be learnt over the MPLS/IP core:-
 - In control plane for EVPN.
 - In data plane [in case of PBB-EVPN].
- L2 Unicast traffic destined to the VXLAN network will be encapsulated with the IP/UDP header+Customer bridge VNI.
- L2 Unicast traffic destined to the MPLS/IP network will be encapsulated with the MPLS label.

Handling Mcast traffic

- Each VXLAN network independently builds its P2MP or MP2MP shared multicast trees for one or more VNIs.
- In the MPLS/IP network, multiple options are available for the delivery of multicast traffic:
 - Ingress replication
 - LSM Options.
- The stitching must ensure for MH VXLAN network:
 - No Packet Duplication
 - No Forwarding Loops

Multicast Stitching with Per-VNI Load Balancing

- PE nodes connected to multi-homed VXLAN network, perform BGP DF election to decide which PE node is responsible for forwarding multicast traffic associated with a given VNI.

Next steps

- Comments are appreciated.

Thank you