

draft-sajassi-l2vpn-pbb-evpn-01.txt

Ali Sajassi (Cisco), Nabil Bitar (Verizon), Aldrin
Issac (Bloomberg), Samer Salam (Cisco), Sami
Boutros (Cisco)

March 30th, 2011

IETF Prague

Baseline Requirements (E-VPN)

1. All-Active Multi-Homing
 - a) Flow-based Load Balancing
 - b) Flow-based Multi-Pathing
 - c) Geo-redundant PE nodes
 - d) Optimal Traffic Forwarding
 - e) Flexible Redundancy Grouping Support
2. Multi-homed Network
3. Multicast Optimization with MP2MP LSP (in addition to P2MP LSP)

Baseline Requirements (E-VPN)

4. Ease of Provisioning
5. New Service Interface
6. Fast Convergence
7. Flood Suppression
8. Finer control over MAC address learning
 - Control which nodes learn which MAC
 - Support of hub-and-spoke and extranet topologies

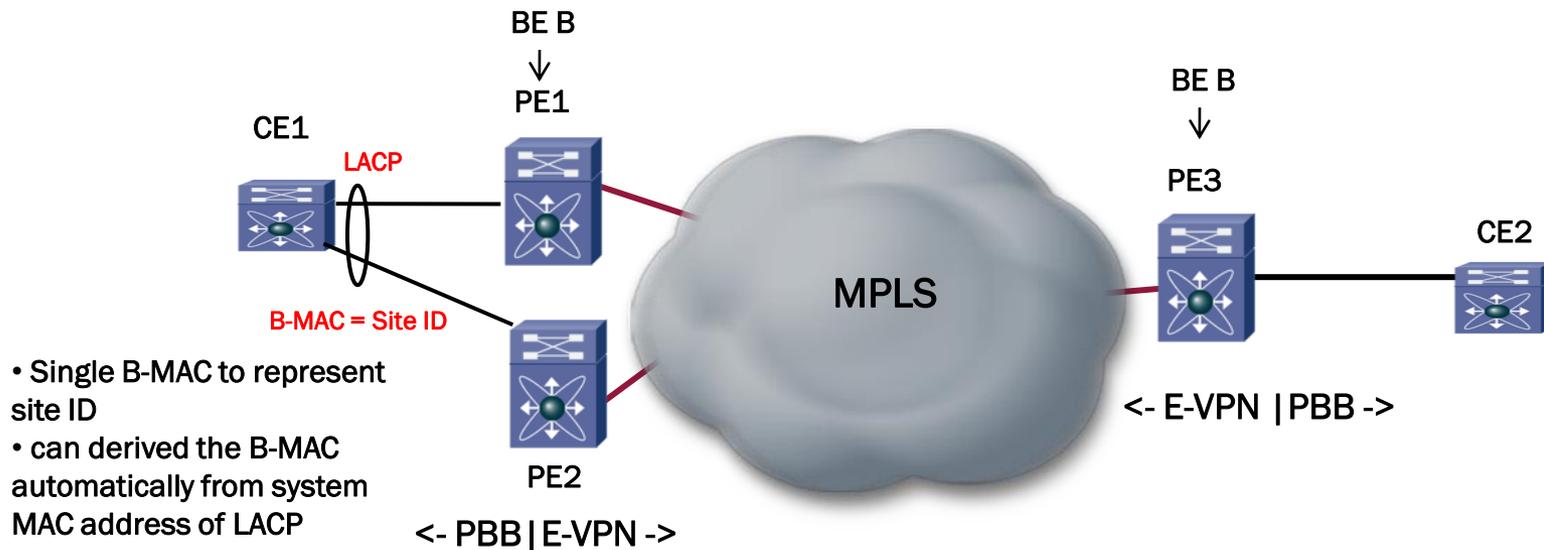
Additional Requirements (PBB-EVPN)

1. MAC Advertisement Route Scalability
 - To support millions of C-MAC addresses (million of VMs)
2. C-MAC Mobility with MAC sub-netting
 - Support C-MAC address mobility while retaining the scalability benefits of MAC sub-netting
3. C-MAC Address Learning and Confinement
 - MES nodes to maintain C-MAC addresses in their RIB & FIB for active flows ONLY

Additional Requirements (PBB-EVPN)

4. Interworking with TRILL & 802.1aq/.1bp networks and C-MAC Transparency
 - To avoid learning of C-MACs by DC WAN Edge PE
5. Per Site Policy
 - To support connectivity policy rules at the granularity of a site (or segment)
6. Avoiding C-MAC flushing
 - To avoid C-MAC flushing upon link, port, or node failure for multi-homed devices
7. Avoid transient loop for known unicast when doing egress MAC lookup

Solution Overview



- Advertise local B-MAC addresses in BGP to all other PEs that have at least one VPN in common just like E-VPN
- Build a forwarding table from remote BGP advertisements just like E-VPN (e.g., association of B-MAC to MPLS labels)
- PEs perform PBB functionality just like PBB-VPLS
- C-MAC learning for traffic received from ACs and C-MAC/B-MAC association for traffic received from core

BGP Encoding

- MAC Advertisement Route (to advertise B-MACs)
 - RD (type1) = IP-addr : 0
 - Eth Tag = 0
 - RT set = RTs corresponding to all EVIs associated with the B-MAC
- Segment Route (would be needed for mLACP state synchronization just like E-VPN)
- Ethernet A-D Route (not needed because site and service ID are identified by B-MAC SA and I-SID) - this simplifies the stuff considerably

Advantages

1. MAC Advertisement Route Scalability

- A single B-MAC represents a multi-homed site
- A single B-MAC can represent all single-homed sites
- A single C-MAC represents a single VM
- => several order of magnitude difference between C-MAC & B-MAC

2. C-MAC Mobility with MAC sub-netting

- Typically C-MACs are not managed and thus can not be sub-netted
- B-MACs on the other hand are always managed and can easily be sub-netted
- Even when C-MACs are sub-netted, VM mobility contradicts the effect of sub-netting

Advantages – Cont.

3. C-MAC Address Learning and Confinement

- With C-MAC learning in control plane, C-MACs are always in RIBs and maybe also in FIBs
- With B-MAC learning in data plane, C-MACs are not never in RIBs and they are only present in FIBs for active flow

4. Interworking with TRILL & 802.1aq/.1bp networks and C-MAC Transparency

- PBB encapsulation enables end-to-end tunneling of C-MAC addresses for the access networks thus avoiding termination and learning by DC WAN Edge PE

Advantages – Cont.

5. Per Site Policy

- Since B-MAC addresses are per site, BGP policy per MAC gives us very nice set of per-site policy

6. Avoiding C-MAC flushing

- Since B-MAC represent a site, a link, port, or node failure doesn't change the B-MAC address – it only changes number of next hop for that B-MAC

7. Avoid transient loop for known unicast when doing egress MAC lookup

- Since B-MAC SA is always transmitted with every frame, checking of every frame against its source MAC SA for known unicast frame is already provided by PBB

Conclusion

- PBB functionality has been incorporated in many VPLS PEs and enjoys wide deployment (draft-ietf-l2vpn-pbb-vpls-interop)
- PBB-EVPN can be supported rather easily by the PEs that support PBB functionality
- PBB-EVPN offers many benefits as described above

Next Steps

- Take it to the mailing list for more discussions