

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

Ali Sajassi (Cisco), Nabil Bitar (Verizon),
Aldrin Issac (Bloomberg), Samer Salam
(Cisco), Sami Boutros (Cisco), Florin Balus
(ALU), Wim Henderickx (ALU), Clarence
Filsfils (Cisco), Dennis Cai (Cisco), Lizhong
Jin (ZTE)

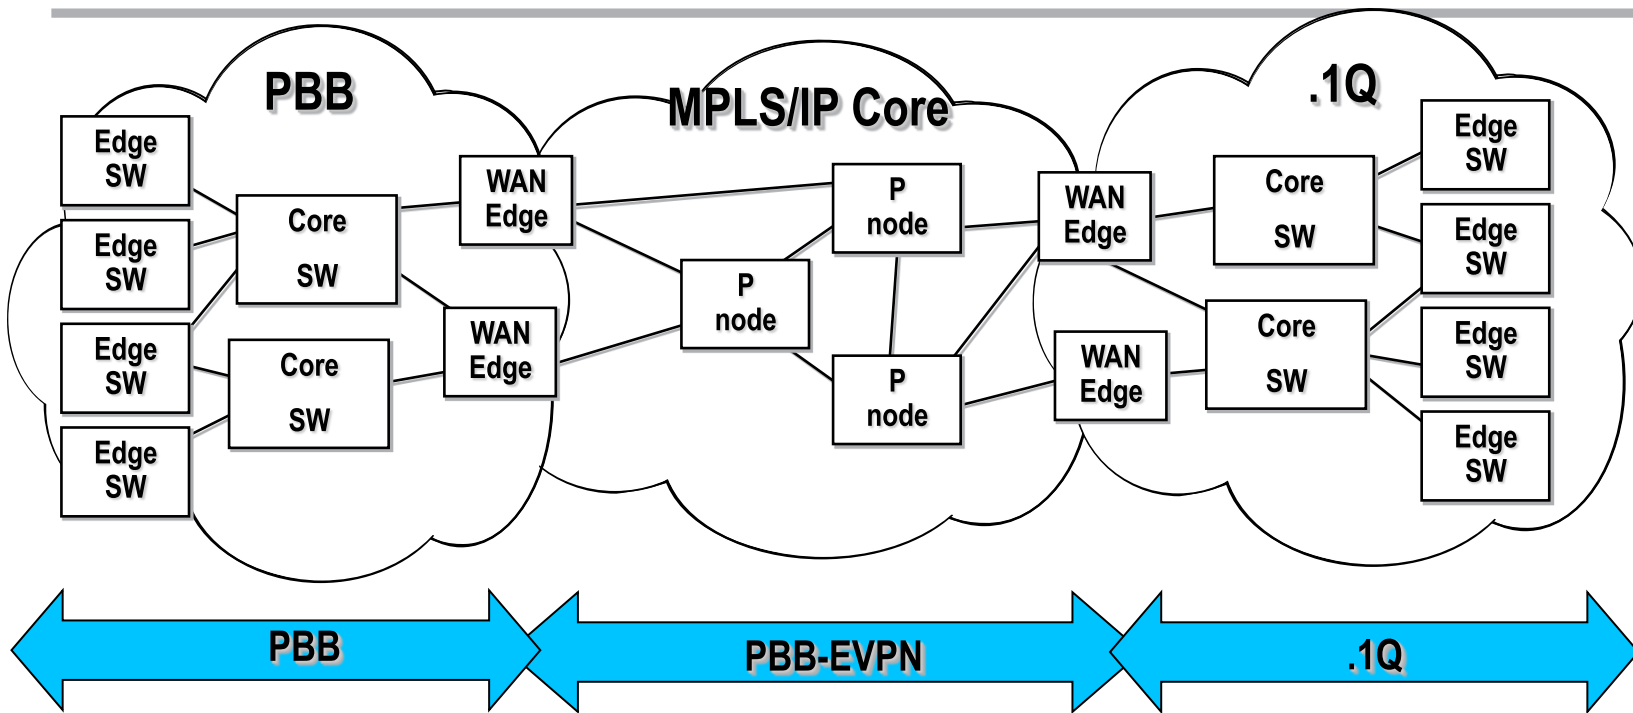
March 26th, 2012

IETF Paris

Areas of Consideration

1. Same AFI/SAFI as E-VPN ?
2. Confining the scope of B-MAC distribution
3. Multicast Pruning per I-SID
4. Multi-homing Network (Active/Standby)
5. Distributing list of I-SIDs between IGP domains via BGP
6. Modifications compare to last rev.
7. Splitting the draft into two

1. Same AFI/SAFI ?



- Need to be able to handle the above Interop scenario
- WAN-Edge PEs on the left, need to support just E-VPN since the received frames are already PBB encapsulated
- WAN-Edge PEs on the right, need to support PBB-EVPN

=> ALL WAN-edge PEs need to use the same AFI/SAFI

2. Confining scope of B-MAC distribution

- Distribute a given B-MAC to MESs that only need to receive it
- Two ways for doing it
 - A. Send a B-MAC along with a list of RTs associated with I-SIDs (auto-derived RTs)
 - B. Send a B-MAC along with a RT corresponding to the B-VID/EVI associated with all the I-SIDs

2. Confining scope of B-MAC distribution – Cont.

- Option-A
 - Pros
 - RTs can be auto-derived from I-SIDs
 - No need for additional config to associate I-SIDs with a B-VID
- Option-B
 - Pros
 - No need to send a long list of RTs – sometimes in tens of thousands
 - No interoperability issue when connecting a .1Q network to a PBB network

- => Option-B is chosen

3. Multicast Pruning per I-SID

- The root PE interested in establishing a P2MP LSP advertises Inclusive Multicast Route with:
 - Tag field set to I-SID
 - RT associated with B-VID/EVI
- Receiving PEs (if is configured for that I-SID), then initiate mLDP signaling for joining the tree
- If Tag field is set to zero, then it means the tree is for all I-SID (e.g., default P2MP tree)

4. Multi-Homing Network (Active/Standby)

- One B-MAC per Ethernet Segment per PE. Each PE only advertises its own B-MAC for the Segment.
- I-SID to PE mapping is learned in the data plane via B-MAC
- Operation is simple but requires C-MAC
- Scope of flushing:
 - Per B-VID/EVI: more I-SIDs flushed than necessary on port/link failure
 - Per I-SID: right scope but if done in BGP will run into issue because of it requires latched-state as opposed to last-state (because of RR route suppression). Or flushing can be done in-band via MIRP
- => Because of its simplicity, it is the recommended option

4. Multi-Homing Network (Active/Standby) - Cont.

- Flushing can be triggered based on either B-MAC withdraw or MAC Mobility Attribute (in segment-route draft)
- It is recommended to use MAC mobility attribute upon link failure or when receiving TCN from access network
- It is recommended to use MAC withdraw upon link decommissioning

4. Multi-Homing Network – Cont.

- One virtual B-MAC per Ethernet Segment, all PEs in RG advertise the same (virtual) B-MAC.
- B-MAC can be auto-derived from the MAC address portion of the ESI.
- Requires I-SID/VLAN list to be advertised with MAC route (e.g. vectorized list).
- If protection done using Backup DF, then C-MAC flushing can be avoided for all failure scenarios except Split MHN.
- Enhanced MAC learning required to avoid extra I-SID lookup in imposition (PE adjacency selection based on I-SID).

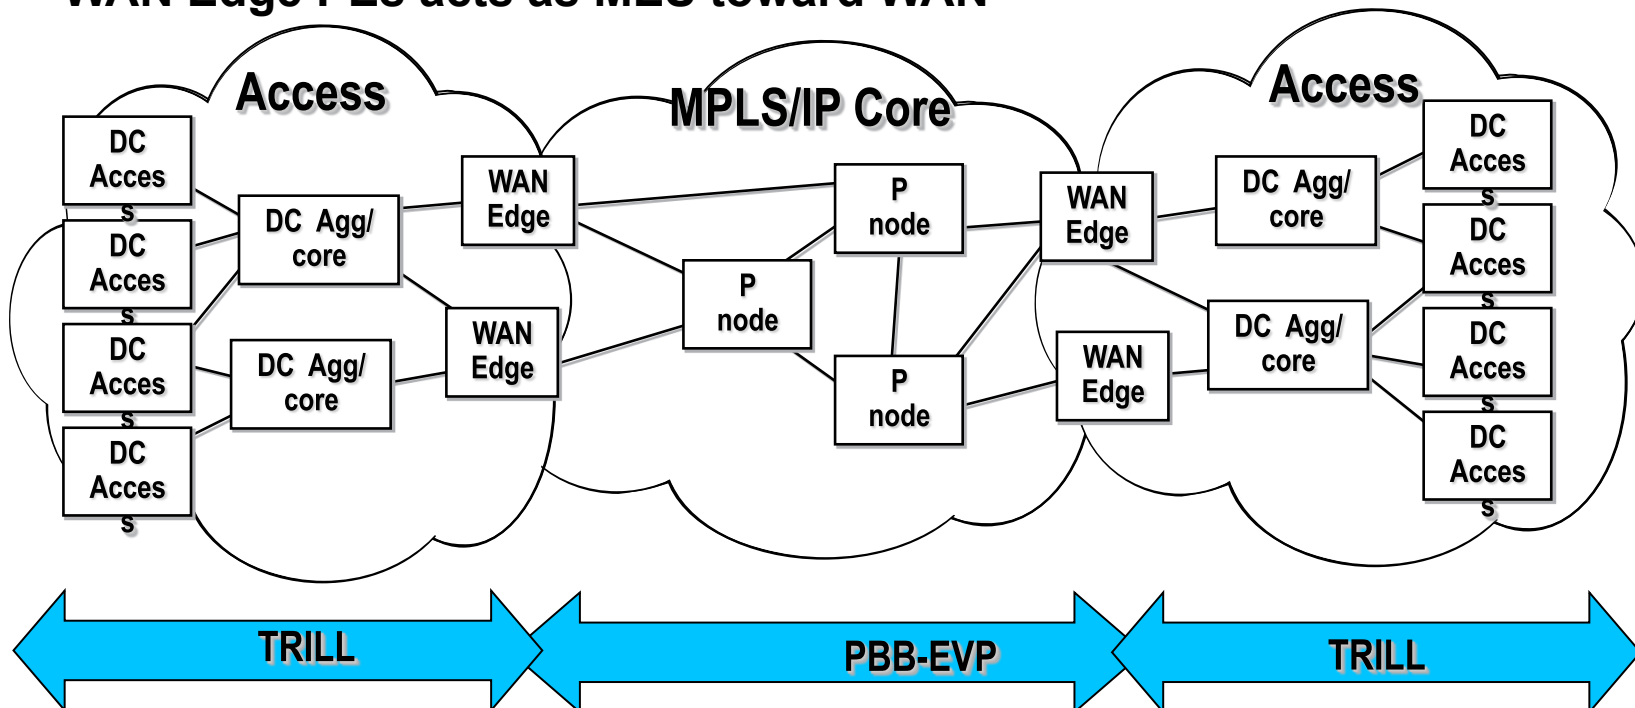
- => for future study

5. Distributing list of I-SIDs between IGP domains via BGP

- Use a vectorized list of I-SID (encoded efficiently) as a BGP attribute
- This would work for I-SID registration/dereg because only last-state is of interest.
- One may ask why not use the same attribute for flushing? It cannot be used for flushing because of RR route suppression (e.g., we need latched state instead of last state)

6. Modifications compare to last rev.

- Added additional text to further clarify DCI operation for TRILL
 - No need to terminate TRILL encapsulation at WAN Edge PEs
 - WAN Edge PEs act as Edge Rbridge wrt control plane and transit Rbridge wrt data-plane
 - WAN Edge PEs acts as MES toward WAN



7. Splitting the draft into two

- Co-authors of this draft would like to separate PBB/SPB parts of the draft from TRILL parts
- Thus, we would like to publish TRILL parts of this draft as `draft-ietf-l2vpn-trill-evpn-00.txt`