


draft-sajassi-l2vpn-evpn-etree-02.txt

A. Sajassi (Cisco), S. Samer (Cisco), Wim
Henderickx (Alcatel-Lucent), Jim Uttaro (AT&T)

IETF 88, November 2013

Vancouver

Scenario of Interest: A site with both Root and Leaf Acs – from IETF 85



- The packets originated from a site, will need to carry site's root or leaf indication (e.g., policy needs to be applied per site basis)
- Egress PE must use the root/leaf indication in the packet to perform appropriate filtering

➔ This scenario in E-VPN is addressed by using per-AC (per-site) policy

Discussions on the mailing list – from IETF 85

- Many exchanges on the mailing list – both public and private
- Public: Application of Split-Horizon filtering capability of EVPN for E-TREE application was not clear to some
- Private: Some argue that we should NOT mandate the use of SH filtering for all scenarios

Current Proposal

- Current proposal leverages existing split-horizon filtering mechanism and provides a good foundation to built on
 - E-VPN already supports a BGP route that identifies a site (ESI)
 - This route is used for Split-Horizon Filtering
 - Color this route with root/leaf indication using ESI Label Extended Community for both **known unicast** & BUM traffic
 - Egress filtering can be done per ESI label as before
- ⇒ no changes in data-plane !
- ⇒ very little changes in control plane (no need to define any new BGP routes or attributes) !

But it can be improved

- To work for PBB-EVPN in addition to EVPN
- To eliminate additional overhead for known unicast traffic by not requiring SH MPLS label to be sent
- To eliminate additional processing overhead on the egress PE by not requiring processing of SH label for known unicast traffic

Enhancement for EVPN

- Color the BUM traffic as before with root/leaf indication using SH label
- For known unicast traffic, advertise a root/leaf indication along with each MAC
- For known unicast traffic, use this root/leaf indication to perform the filtering on the ingress PE (instead of egress PE)
- For BUM traffic, use SH label to perform the filtering on the egress PE as before (ingress filtering cannot be done for BUM traffic)

Enhancement for EVPN – Cont.

- For Inter PE forwarding of known unicast
 - On the ingress PE, after performing a lookup on the CMAC DA, if it indicates that the CMAC DA belong to a leaf and the AC or ES is also associated with a leaf, then don't forward the packet
- For Intra PE forwarding
 - Put all the leaf ports for a given E-TREE (given VPN) in its split-horizon group and perform SH filtering internal to the box

For PBB-EVPN

- For both BUM and known unicast traffic, advertise a root/leaf indication along with each BMAC – e.g., color each BMAC with root/leaf indication
- For known unicast traffic, use this root/leaf indication on BMAC DA to perform the filtering on the ingress PE
- For BUM traffic, use this root/leaf indication on BMAC SA to perform the filtering on the egress PE (as done in baseline PBB-EVPN)
 - Filtering on egress PE is done using BMAC SA solely (no need to use any flag in DP)

For PBB-EVPN – Cont.

- For Inter PE forwarding of known unicast
 - On the ingress PE, after performing a lookup on the CMAC DA, and getting corresponding BMAC DA, if it indicates that the BMAC DA belong to a leaf and the AC or ES is also associated with a leaf, then don't forward the packet
 - Filtering on ingress PE is done using root/leaf flag in DP (just like EVPN)
- For Intra PE forwarding
 - Put all the leaf ports for a given E-TREE (given VPN) in its split-horizon group and perform SH filtering internal to the box

Next Step

- Would like to ask for WG call