

EVPN SAI WG

Discussion on SAI PR 1025

March 2020



Agenda

- Overview of SAI PR 1025
- Tunnel with dest-ip approach
- Tunnel without dest-ip approach

Overview of SAI PR 1025

- Consensus of both the approaches
 - L3 VxLAN will continue to operate with existing P2MP approach. No change from the existing approach
 - L2 VxLAN has requirement to operate in P2P approach. Each remote peer (static or discovered through EVPN) will be represented by the corresponding SAI tunnel object.
 - For more information: <https://github.com/opencomputeproject/SAI/pull/1025>

Approach 1 - Tunnel with dest-ip

- When a bridge is created (802.1d or 802.1q), the corresponding flooding group is created inherently. When the members are added/removed, it adds/removes the corresponding member from the bridge.
- Create tunnel bridgeport per dest-ip tunnel.
- Similar to any other interface, add/remove the bridge member (tunnel bridgeport) to the corresponding bridge (802.1d or 802.1q).

Approach 2 - Tunnel without dest-ip

- When a bridge is created (802.1d or 802.1q), the corresponding flooding group is created inherently. When the members are added/removed, it adds/removes the corresponding member from the bridge.
- However, the above approach is followed only for the physical interface
- For logical interface (like vxlan tunnel or in future mpls pwe), I2 multicast group needs to be created.
- Associate the above multicast group to the corresponding vlan/bridge.
- Create tunnel as members of the I2mc group
- When last tunnel is removed from the corresponding I2mc group, it needs to be dissociated with the bridge and destroy the I2mc group

EVPN deployments

- With EVPN deployments, the tunnel membership for the corresponding bridge can change at a very rapid pace (Consider 4k vlans with 8 evpn remote peers).
- Creation/deletion of l2mc group with association/dissociation of the l2mc group to/from vlan/bridge and keeping track of the tunnel membership per vlan/bridge needs to be considered from both application as well as SAI perspective. (Especially, when it has been agreed that for L2 VxLAN there will be SAI tunnel objects for each of the remote EVPN peer).