

AC-Aware Bundling Service Interface in EVPN

Ali Sajassi

Mankamana Mishra

Samir Thoria

Patrice

Requirements


- Some of the deployments require capability to have multiple subnets within single Bridge Domain.
- Each of subnet are distinguished by Vlan in Bridge Domain

None of existing service interface from RFC 7432 fulfill requirement to have single IRB (Bridge Domain) with multiple Subnets (Distinguished by VLAN)


Existing Service interface

1. **VLAN-Based Service Interface:** Each Bridge Domain has only one Vlan. It means Bridge domain to subnet mapping is 1:1
2. **VLAN Bundle Service Interface:** An EVPN instance corresponds to multiple VLANs, but single Bridge Table is maintained. MAC address is key to lookup. The MPLS-encapsulated frames MUST remain tagged with the originating VID
3. **VLAN-Aware Bundle Service Interface:** With this service interface, an EVPN instance consists of multiple broadcast domains (e.g., multiple VLANs) with each VLAN having its own bridge table. Which would map to multiple IRB interface. MAC+VLAN are learnt, and lookup is based on MAC + VLAN to get appropriate port.

Problem Statement:



- PE1, PE2 are multi-homed PE where single Bridge Domain has multiple subnets .
- When PE1 learns a MAC address from any of subnet, it Need to advertise MAC route to peers in same EVPN instance.
- Currently if PE1 originates new MAC route,
 - PE3 need to create next hop entry to reach MAC.
 - PE2 need to assign appropriate Bridge port to forward packet destined to MAC.
- Current Service Interface defined in RFC7432 does not cover this requirement, as PE2 would not have any context of which VLAN this MAC was learnt.



AC-aware Bundling Service Interface

- New Service interface type is defined to fulfill the requirement which was described earlier.
- With AC-aware bundling service interface , MAC advertisement would have Extended Community which would carry Attachments Circuit ID.
- New Extended community MUST be carried with EVPN Route Type 2 (MAC Route), 7&8 (Multicast Route)
- Attachment Circuit extended community MUST be ignored by Non Multi-homed peers.

Control plane operations

Local Route:

1. When a PE learn MAC address, it MUST attach attachment circuit ID extended community with MAC route
2. If a PE receive IGMP membership notification (join / leave) it MUST attach attachment circuit ID extended community with multicast route

Remote Route:

1. If recipient PE of MAC route is non multi-homed peer, it must process MAC route with procedure defined in RFC7432
2. If recipient PE of MAC route is multi-homed peer, It MUST get AC information from Extended community and associate MAC address with appropriate AC / Subnet / Vlan
3. Multicast route MUST also programmed with appropriate AC / Subnet / Vlan, where AC information is derived from extended community

Data Plane operations at non multi-homed peer

- This service interface does not change any procedure at non-multihomed remote peer of EVPN instance. It MUST follow procedure defined in RFC 7432

Data Plane operations at multi-homed peer

- Data received from CE MUST follow procedure defined in RFC7432. New proposal does not add any extra processing.
- Unknown unicast packet MUST follow procedure defined in RFC7432. New proposal does not add any extra processing.
- With new defined Service interface for known unicast packet MAC lookup would provide Associated Vlan and port information. Appropriate Vlan tag must be pushed and forwarded on port.
- Multicast packet MUST be forwarded based on multicast state.

Questions & Feedback