draft-ietf-bess-rfc7432bis-00.txt

A. Sajassi (Cisco)

IETF 109, November 2020

Online

Areas of Improvement

- Errata already logged for RFC 7432
  - Need to ensure that all the route label field descriptions say the label is encoded in the high-order 20 bits (bits 8 to 24)
  - Editorial fixes
- Clarifying relationship among MAC-VRF, bridge table, EVI, bridge domain, VLAN, and VID
- Introduction of L2 attributes EC with CW, MTU and P/B bits
  - Update of the DF Election procedure, so that, for single-active, PEs elect not only a DF but also a BDF and signal P/B respectively
- EVPN routes priority handling
- Clarification for best path selection for default GW
- Support of DCB allocated labels for ESI-labels, which then adds support for MP2MP tunnels in EVPN

MAC-VRF, BT, EVI, BD, VLAN, VIDs

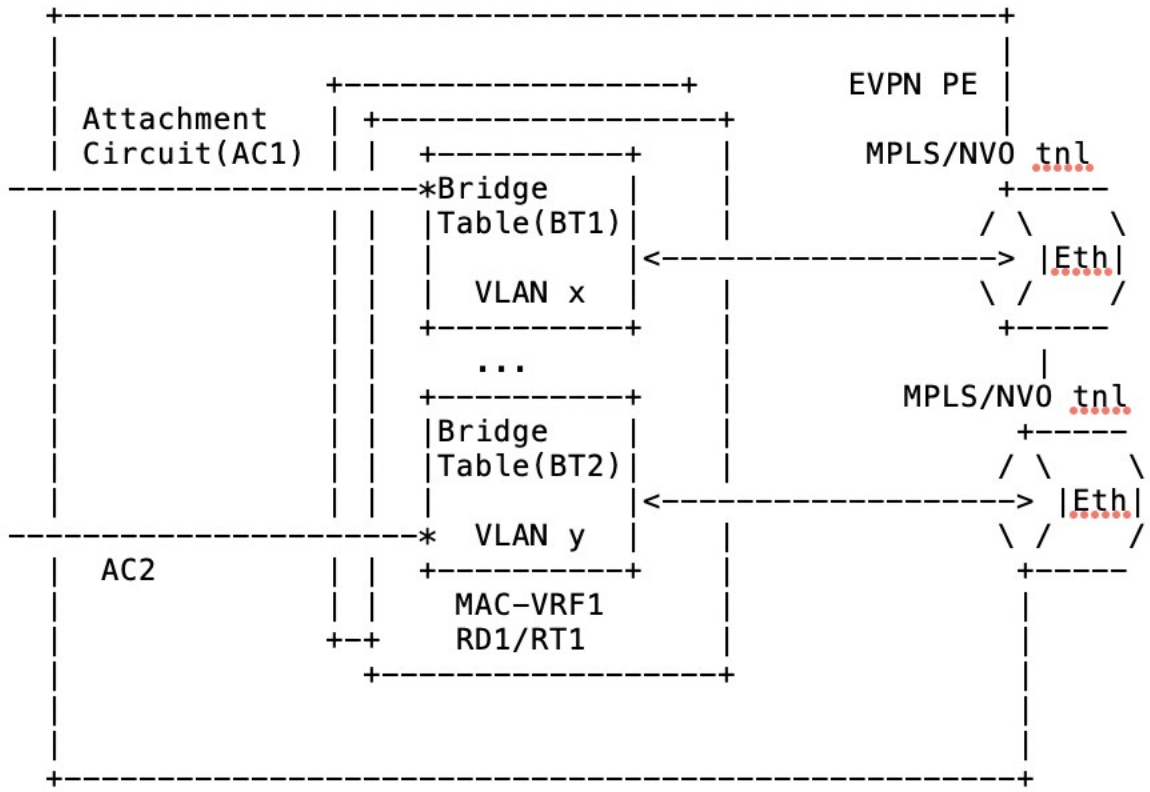

Figure 1: EVPN IRB PE Model

PE Model – Continue

- - An EVI can consists of one or more BDs. Furthermore, a MAC-VRF can consists of one or more BTs. A BD is typically represented by a single VLAN ID (VID); however, it can be represented by multiple VIDs (i.e., Shared VLAN Learning (SVL) mode in 802.1Q).
- In VLAN-based mode, there is one EVI per VLAN and thus one BD/ BT per VLAN. Furthermore, there is one BT per MAC-VRF.
- In VLAN-bundle service, it can be considered as analogous to SVL mode in 802.1Q i.e., one BD per EVI and one BT per MAC-VRF with multiple VIDs representing that BD.
- In VLAN-aware bundle service, there is one EVI with multiple BDs where each BD is represented by a VLAN. Furthermore, there are multiple BTs in a single MAC-VRF.

L2 attributes EC

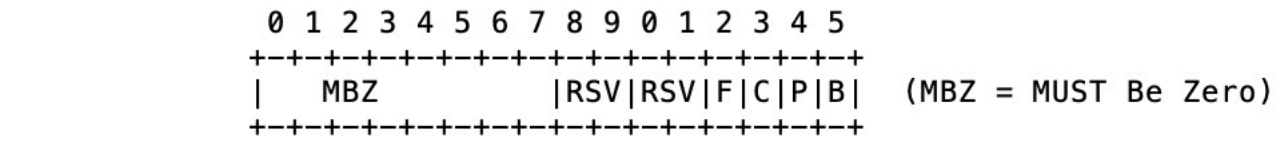

The bits in Control Flags are extended by the following defined bits:

Name	Meaning
F	If set to 1, a Flow Label MUST be present when sending EVPN packets to this PE.

1. **The EVPN Layer 2 Attributes Extended Community, when added to Inclusive Multicast route:**
   - **Bridge/MAC-VRF attributes MTU, Control Word and Flow Label are conveyed**
2. **The EVPN Layer 2 Attributes Extended Community is included on Ethernet A-D per EVI route and:**
   - **per-ESI attributes P, B are conveyed,**

CW Signaling & Backward Compatibility

PE2 \ PE1	Signals CW=0	Signals CW=1	Does Not Signal
Signals CW = 0	Both PEs MUST NOT use Control-Word (agreement)	Both PEs MUST NOT use Control-Word (mismatch)	PE1 SHOULD use (RFC7432, s.18) PE2 raises alarm (mismatch)
Signals CW = 1	Both PEs MUST NOT use Control-Word (mismatch)	Both PEs MUST use Control-Word (agreement)	PE1 SHOULD use (RFC7432, s.18)
Does Not Signal	PE2 SHOULD use (RFC7432, s.18) PE1 raises alarm (mismatch)	PE2 SHOULD use (RFC7432, s.18)	Both PEs SHOULD use Control-Word (RFC7432, s.18)

EVPN routes priority handling

1. Ethernet AD per ES (Mass-Withdraw Route Type 1) and Ethernet Segment (Route Type 4) are lower scale and highly convergence affecting, and SHOULD be handled in first order of priority
2. Ethernet AD per EVI, Inclusive Multicast Ethernet Tag route, and IP Prefix r oute defined in [I-D.ietf-bess-evpn-prefix-advertisement] are sent for each Bridge or AC at medium scale and may be convergence affecting, and SHOULD be handled in second order of priority
3. MAC advertisement route (zero and nonzero IP portion), Multicast Join Sync and Multicast Leave Sync routes defined in [I-D.ietf-bess-evpn-igmp-mld-proxy] are considered 'individual routes' and very-high scale or of relatively low importance for fast convergence and SHOULD be handled in last order of priority.

Best Path Selection for Default GW

- When comparing two routes, the route having Gateway GMAC EC is preferred over the route that doesn't have GMAC EC. The PE that has advertised the MAC route without GMAC EC upon receiving the route with GMAC EC, SHALL withdraw its route and raise an alarm.
- When comparing two routes where both routes having GMAC EC, normal BGP best path processing will be applied
- When comparing local and remote route having Gateway GMAC EC, the local route is always preferred.
- MAC Mobility EC SHALL not be attached to routes having GMAC EC on the sending side and SHALL be ignored on the receiving side.

Additional work for Rev01

- Loop protection (from the other draft that we decided to leave)
- DCB labels that would allow MP2MP LSPs – although this has implications in the references
- Also improving the terminology similar to RFC8584
- Fast MAC mobility handling

Next Step

- Publish draft-ietf-bess-rfc7432bis-00.txt after the meeting
- Solicit input
- No rush for WG LC – i.e., expect several IETF meetings to ensure everything is clear