draft-sajassi-bess-evpn-rfc8317bis-01.txt

A. Sajassi (Cisco), J. Rabadan (Nokia), J. Drake (Juniper), A. Appachi gounder (Google), A. Bamberger (Arista)

IETF 119, March 2023

Brisbane

History

- Presented Rev00 at IETF 117 in July 2023
- No major changes for Rev01 !!
  - Expanded terminology section
  - Fixed typos in figures 1, 2, and 3
  - Further clarified flag usage in section 6

- Following is just a review of rev00 with updated figures

Scenarios in RFC 8317

1. Leaf or Root Site(s) per PE
2. Leaf or Root Site(s) per AC Scenario (superset of scenario-1)
3. Leaf or Root Site(s) per MAC Address (no changes)

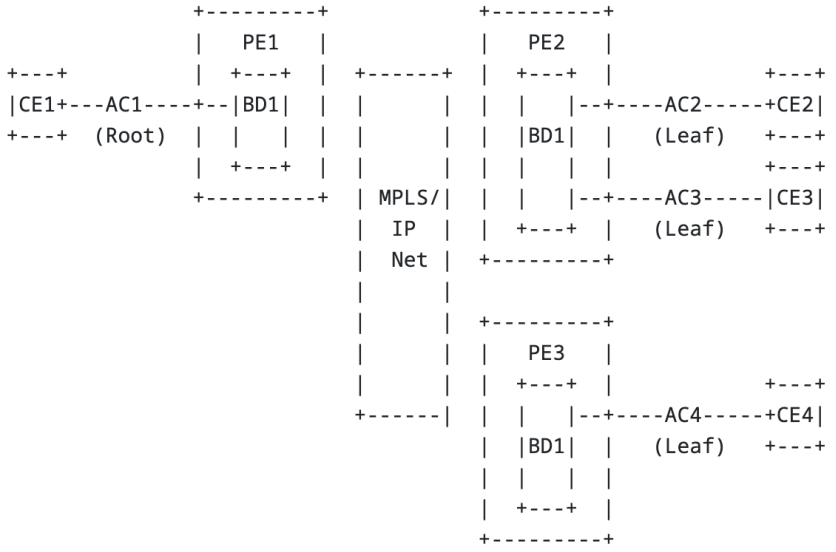

Figure 1: Scenario 1

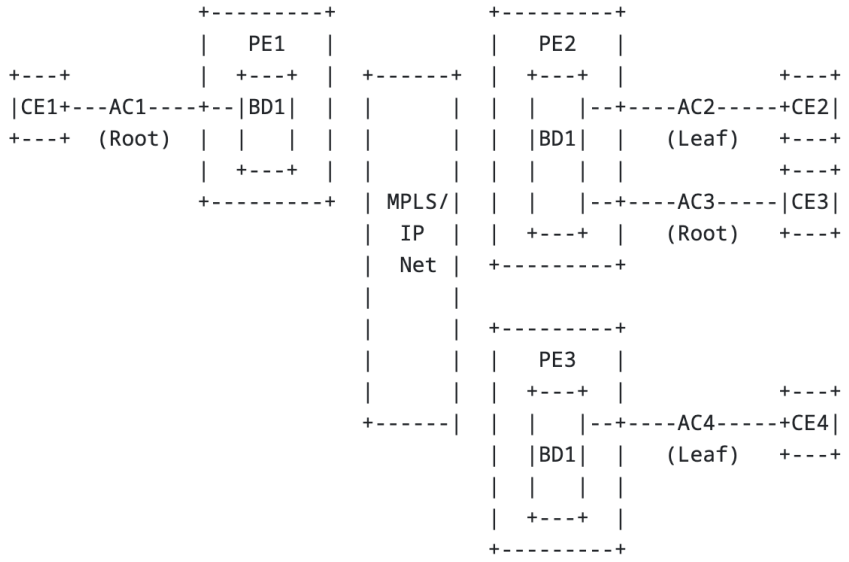

Figure 2: Scenario 2

Clarifications

1. Clarified when talking about about root or leaf in context of an EVI, it is basically root or leaf for a BD where BD = EVI (for VLAN-based or VLAN-bundle services) and BD= EVI+VLAN for VLAN-aware service
2. Changed the diagrams to show 3 PEs instead of 2 for better demonstration of root/leaf ACs

Changes Relative to RFC 8317

1. Changes to scenario-1 - Leaf or Root sites per PE:
   a) Removed 2-RT option because it doesn’t address EVPN host mobility
   b) 1-RT option which was the recommended option, is now the only option
   c) Ingress filtering of known unicast traffic is done at ingress PE by coloring MAC/IP routes as before
   d) For this simple scenario-1, the BUM filtering can also be done at ingress PE for ingress rep

Changes Relative to RFC 8317 - II

1. Changes to scenario-1 - Leaf or Root sites per PE:

   e) BUM filtering is achieved by coloring IMET routes via existing tailored BGP route import/export policy or via signaling extensions

   f) The transmit policy matches the IMET route and colors it with BGP standard community, and the receive policy checks IMET routes with this policy and discard them

   g) When a PE can start with scenario-1 and move to scenario-2 (or vise versa), then

Changes Relative to RFC 8317 - III

1. Changes to scenario-2 - Leaf or Root sites per AC:
   a) Ingress filtering for this scenario when IR used, can be adaptive – i.e., to dynamically decide toward what PE to perform ingress filtering and toward what PEs not to.
   b) Need to convey ingress source type (leaf or root) in data-plane. MPLS was already covered in 8317. This draft extended it for VxLAN.

Changes Relative to RFC 8317 - IV

- Section 5.3 is added to describe coloring of non-MPLS encapsulated user traffic (e.g., VxLAN) for egress filtering of BUM traffic
- Need to indicate whether user traffic is from root or leaf
- Approach: to send a IMET route along with a Tunnel Encap EC and E-Tree EC. If the Tunnel Encap is VxLAN (or GENEVE), and leaf-indication flag is set, then VNI identifies both EVI and leaf as a source of traffic.
- The receiving PE, when receives traffic with this leaf VNI, it drops the packet at leaf ACs and passes it at root ACs.

Changes Relative to RFC 8317 - V

- Added section 5.5 on Adaptive Filtering

Condition	Root AC	Leaf AC	E-Tree EC Fields
1	No	No	No E-Tree EC
2	Yes	No	No E-Tree EC
3	No	Yes	Root-Flag=0, Leaf-Flag=1 Leaf-VNI= valid or Leaf-VNI= 0xFFFFFF
4	Yes	Yes	Root-Flag=1, Leaf-Flag=1 Leaf-VNI= valid or Leaf-VNI= 0xFFFFFF

Figure-4: E-Tree EC Setting by an Ingress PE per BD

Changes Relative to RFC 8317 - VI

Condition	Ingress PE role	Ingress AC role	Flood List
1	Root	Root	Use All-PEs flood list with Base VNI for BUM traffic
2	Leaf	Leaf	Use non-Leaf-PEs flood list with Leaf VNI for BUM traffic
3	Root+ Leaf	Root	Use All-PEs flood list with Base VNI for BUM traffic
4	Root+ Leaf	Leaf	Use non-Leaf-PEs flood list with Leaf VNI for BUM traffic

Figure 5: When to use "non-Leaf PEs" flood list

Next Step

- The draft is very stable
- Requesting WG adoption

THANK YOU!