draft-sajassi-bess-pbb-evpn-anycast-ip-tunnels-00.txt

A. Sajassi (Cisco), S. Salam(Cisco), L. Jalil (Verizon), J. Drake (Juniper), T. Singh (Juniper)

IETF 94, November 2015

Yokohama

Reference Figure

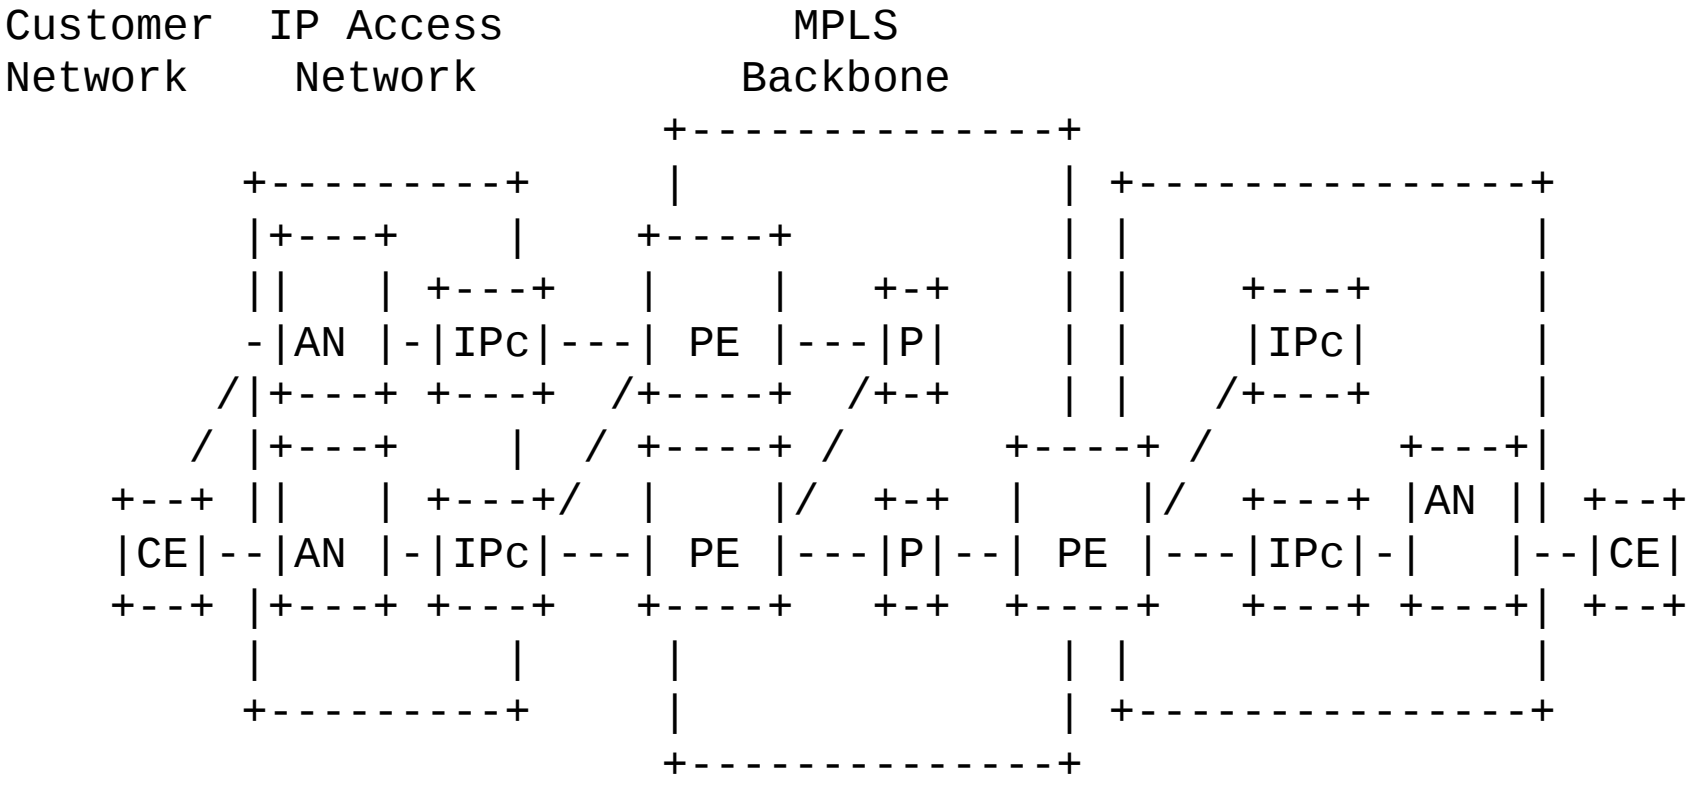

Figure 1: Target Topology

Requirements

- Support for both IP and non-IP payload – i.e., need MAC learning in data-plane
- PBB-EVPN with IP access network (both v4 and v6)
- Support QoS (policing/shaping) per tunnel/VLAN
- Simplify ANs provisioning by using Anycast IP addresses as tunnel destination on the PEs – ie, no need to provision unicast IP addresses of redundant PEs on the ANs
- Provide resilient interconnect with protection against PE node failure and IP tunnel failure
- Fast recovery from failure

Challenges with PBB-EVPN

- Challenges with All-Active redundancy mode
  - Traffic arriving from MPLS backbone gets load balanced among the PEs in the redundancy groups
  - PEs cannot perform proper policing/shaping for traffic destined toward CEs because one cannot assume traffic is evently distributed among PEs in the redundancy group (due to elephant flows)

Challenges with PBB-EVPN – Cont.

- Challenges with Single-Active redundancy mode
  - Based on DF election, only one PE will be forwarding traffic from access to the backbone
  - However, the DF PE may NOT be the shortest IGP path from the CE – ie, CE forwards traffic to non-DF PE where it gets dropped

Solution

- Define a new asymmetric redundancy mode for PBB-EVPN
- It behaves like All-Active in the direction of access-to-core – ie, All PEs in the redundancy group can receive traffic from ANs
- It behaves like Single-Active in the direction of core-to-access – ie, remote PEs only choose a single PE to send traffic to

Example Network

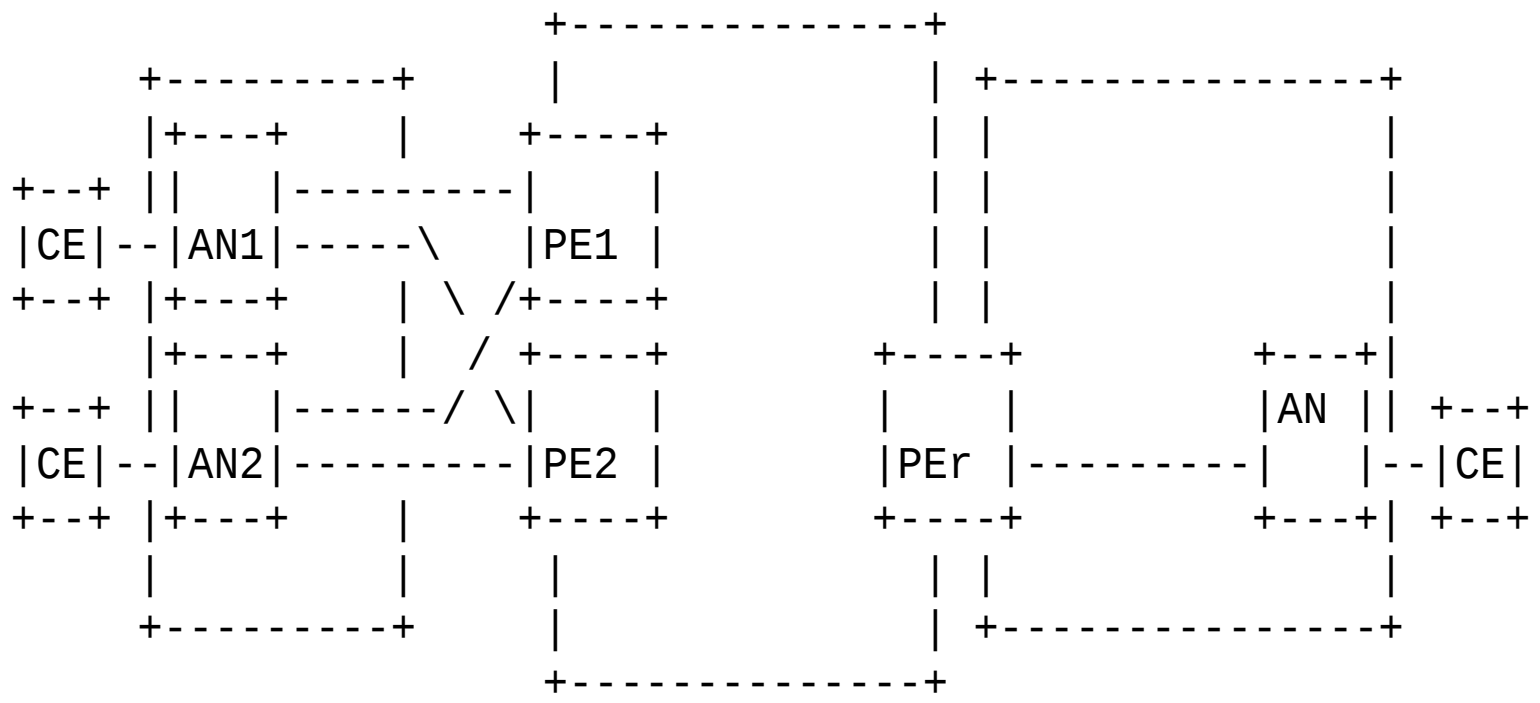

Figure 2: Example Network

Solution Overview – Known Unicast Traffic

- The AN forwards the traffic to the PE with the best IGP distance
- The PE learns MAC address against the received tunnel, performs QoS, encapsulates it with PBB-EVPN and forwards it to the destination PE
- The destination PE learns CMAC against BMAC associated with that PE/vES
- All subsequent packets destined to that CMAC are forwarded to the right PE accrodingly

Solution Overview – BUM Traffic

- From access-to-core direction, it operates just like All-Active redundancy mode in PBB-EVPN with an extension to its split-horizon filtering
- Instead of vES represented by a single BMAC address, it is represented by multiple BMAC addresses (one per PE)
- Split-horizon is performed when a PE receives a BMAC-SA that aliases to its vES
- To reduce no. of BMACs, “local bias” mechanism defined in [Overlay] is used – ie, one BMAC per PE instead of one BMAC per <PE, vES>

Next Step

- Questions ?