

draft-sajassi-bess-evpn-l3vpn- multihoming-00.txt

A. Sajassi (Cisco), S. Salam(Cisco), D. Cai
(Cisco)

IETF 94, November 2015

Yokohama

Requirements

- S-PEs in the redundancy group provide Single-Active redundancy the CEs
- The SPEs in a redundancy group must appear as a single IP peer to the CE, and a single eBGP session will be established between a given CE and its associated S-PEs
- In the case of S-PE failure, pseudowire failure or S-PE isolation from access network, the fail-over time should be minimized by optimizing both the backup pseudowire establishment as well as the BGP convergence time.

Reference Figure

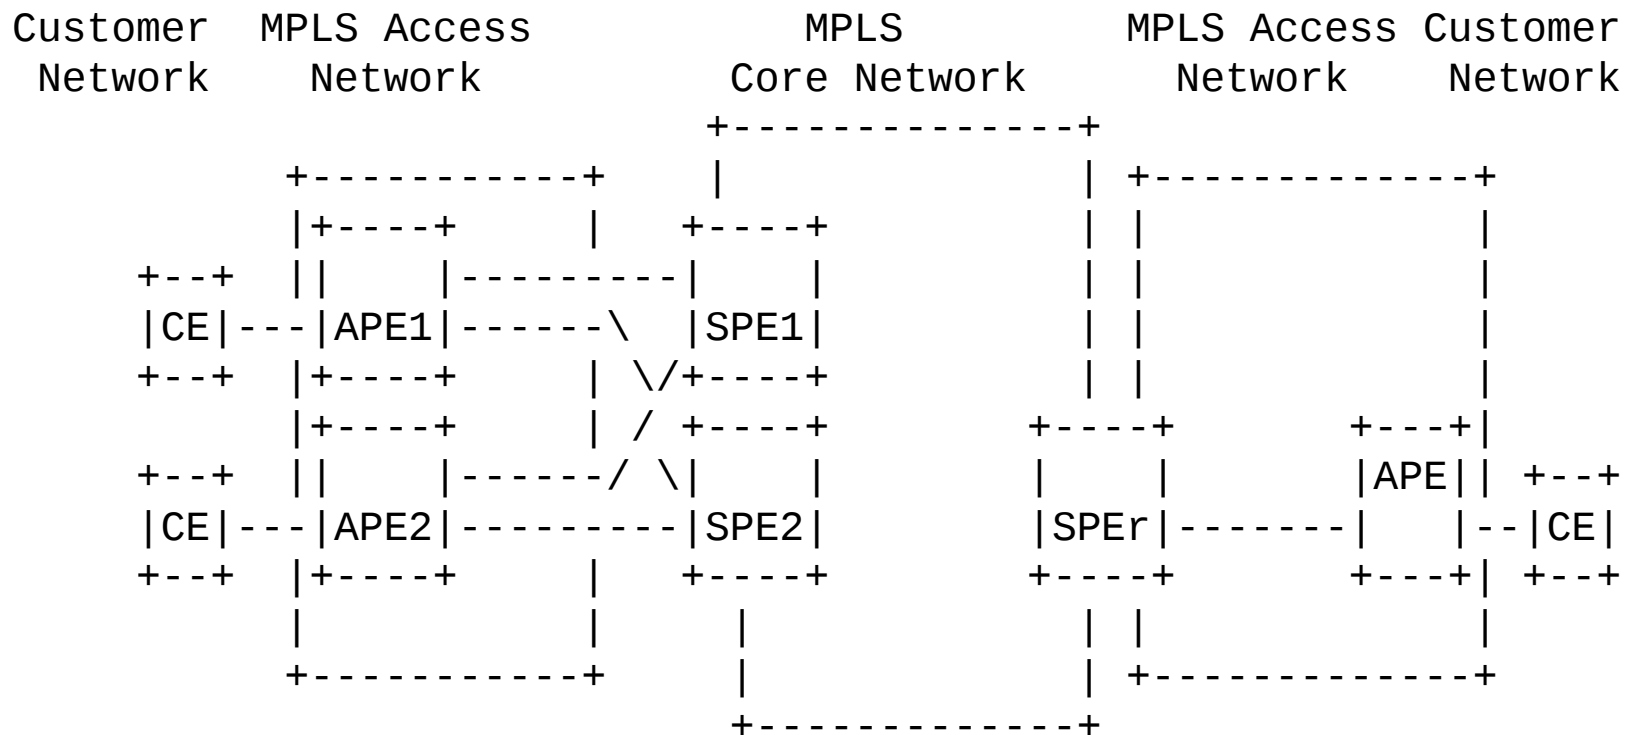


Figure 1: Network Topology

Challenges with L3VPN Multihoming

- No Synch-up mechanism for IP-VRF and ARP tables among PEs in the redundancy group
- No Synch-up mechanism for IP-VRF tables between CE and backup PE

Solution Summary

- Run EVPN Single-Active redundancy mode for vES represented by a setup of PWs
- Primary PE (and backup PE) advertises Active (and Standby) in Preferential Forwarding Status bit over corresponding PW
- All PEs in the same redundancy group are configured with IP and MAC anycast addresses on virtual interfaces corresponding to termination of PWs onto the VRFs
- Primary PE advertises prefixes learned over its AC to other PEs

Solution Summary – Cont.(2)

- Remote PE install primary and backup paths (via ESI) for each received prefix based on RFC 7432 procedures
- Backup PEs synchronize their IP-VRF tables by installing prefixes pointing to their local ESI when the received ESI is local
- Backup PE synchronize their ARP tables by installing MAC addresses pointing to their local ESI when the received ESI is local

Solution Summary – Cont. (3)

- When primary PW failure is detected, the primary PE withdraws Eth Segment Route and performs Mass Withdraw
- Remote PE upon receiving Mass Withdraw, switch to backup path
- Backup PE upon receiving Eth Segment Route withdrawal message, becomes DF, activates backup PW, and re-establishes BGP session to synch-up its VRF table with the CE
- Data can be forwarded via backup PE to/from the CE since FIBs are populated a priori

Next Step

- Questions ?