EVPN multi-homing port-active load-balancing

IETF-101 [London]

Patrice Brissette
Samir Thoria, Ali Sajassi

March, 2018

Active/Standby MC-LAG using ICCP

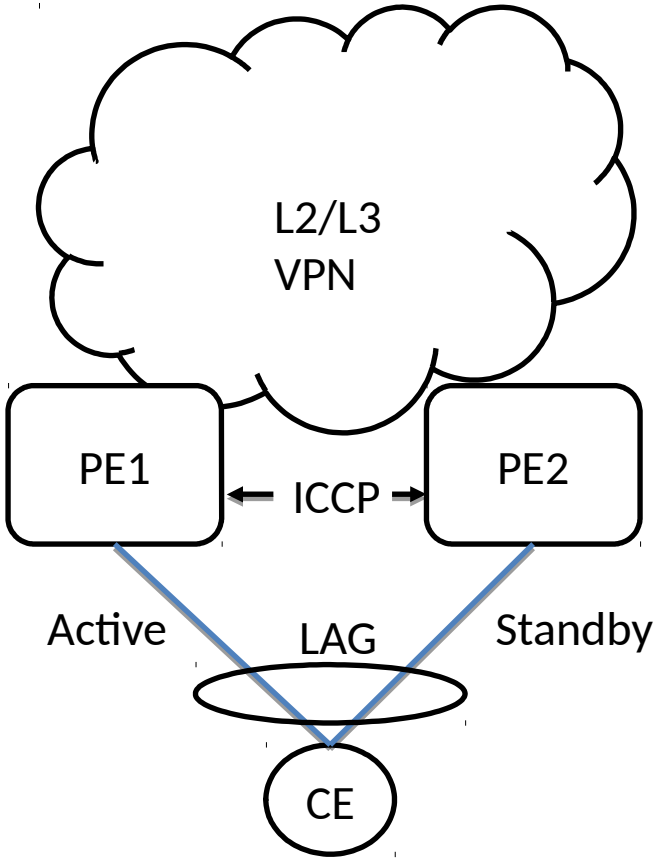

Following picture shows a typical MC-LAG topology using ICCP to drive the active-standby state on peering routers.

That legacy approach requires to run ICCP protocol over LDP between peering routers to elect the active interface and the standby interface.

Per interface active/standby in EVPN

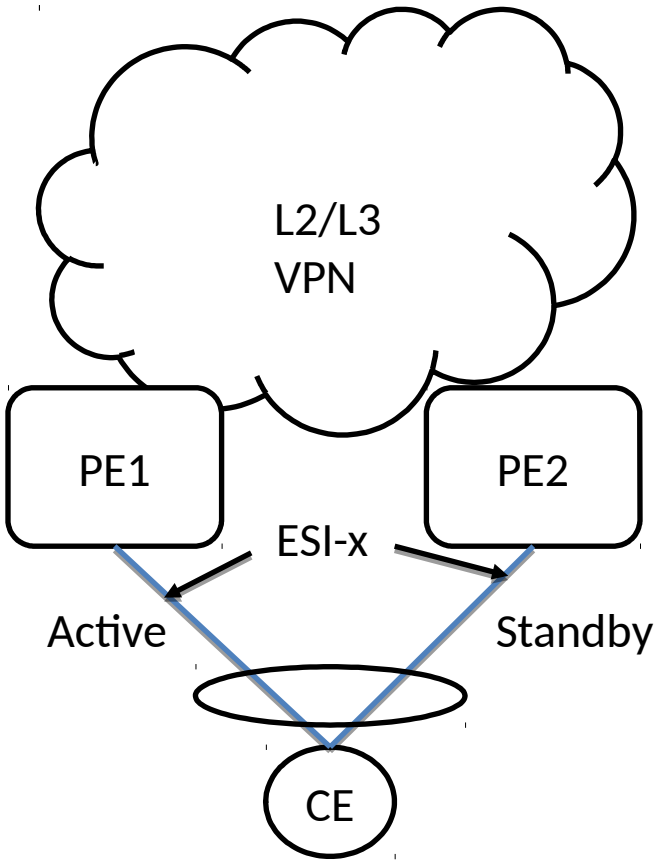

Following steps describe the proposed design with EVPN:

1- ESI is assigned per bundle interface. This could be auto-derived or manually configured.
2- No L2 service is required e.g. neither bridge-domain, neither xconnect (therefore, no L2VPN configuration)
3- Ethernet-Segment is configured in **per-port** load-balancing mode
4- Router exchanges only Ethernet-Segment route (RT-4). No other EVPN routes are used for redundancy.
5- Each router perform Designated Forwarder Election.
6- DF router keeps the entire access interface active
7- Non-DF router brings the entire access interface in down state. If the interface is running LACP protocol, then the non-DF may also set the LACP state to OOS (Out of Sync) as opposed to interface state down, this allows for better convergence on standby to active transition.

Advantages

- Open standards based per interface single-active redundancy mechanism that eliminates the need to run ICCP and LDP.
- Agnostic of overlay technology (MPLS, VXLAN, SRv6) and associated services (L2, L3, Bridging, Xconnect).
- Provides a way to enable deterministic QOS over MC-LAG attachment circuits
- Ease of configuration, no need to configure ICCP/LDP.
- Fully compliant with RFC-7432, does not require any new protocol enhancements to existing EVPN RFCs.
- Can leverage various DF election algorithms e.g. modulo, HRW, etc.
- Replaces legacy MC-LAG ICCP-based solution, and offers following additional benefits
  - Efficiently supports 1+N redundancy mode (with EVPN using BGP RR), ICCP requires full mesh of LDP sessions among PEs in redundancy group
  - Fast convergence with mass-withdraw is possible with EVPN, no equivalent in ICCP
- Customers want per interface single-active redundancy, but don’t want to enable LDP (e.g. they may be running VXLAN or SRv6 in the network). Currently there is no alternative to this. Customer is looking at deceasing the number of protocol being used on their network.