

MVPN Upstream DF Selection

draft-wang-bess-mvpn-upstream-df-selection-01

Heng Wang @Huawei

Fanghong Duan @Huawei

IETF 114

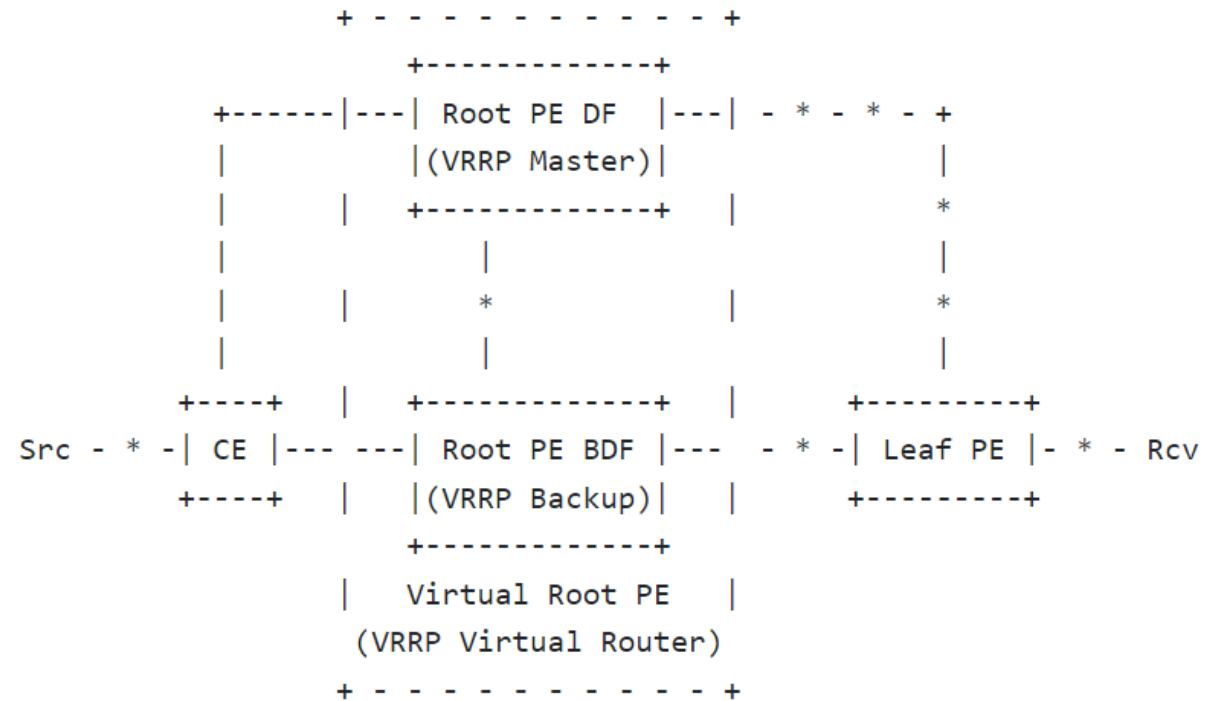
Jul. 2022

Background

- [RFC9026] defines UMH Selection Based on Provider-Tunnel Status for MVPN fast failover.
 - “Hot root standby” will result in the traffic redundancy throughout the backbone network.
 - It is somewhat complicated for the downstream PEs to find an efficient and accurate method to determine the "status" of a P-tunnel.
 - ✓ Discontinuous multicast flow
 - ✓ Lack of effective BFD capability
 - ✓ Remote upstream fault need to be perceived faster and more directly
- The idea: an upstream fast failover consideration.
 - Upstream PEs perform a Designated-Forwarder selection to prevent the traffic redundancy in the P-tunnel.
 - Downstream PEs perform “ANYCAST” RPF checking.
 - Upstream DF status described here is per-VPN, and could be more refined (per-leaf or per-flow) for further use case with stateless multicast underlay (IR/BIER) adopted.

Upstream Designated Forwarder Selection

- Map the role of the VRRP routers to that of the upstream PEs in MVPN dual homing upstream PEs.
 - Virtual Router -- pair of dual homing upstream PEs
 - Virtual Router Master -- the primary upstream PE
 - Virtual Router Backup -- the standby upstream PE



Upstream Designated Forwarder Selection

- Both the primary and standby PEs install VRF PIM state corresponding to BGP Source Tree Join route and send C-Join messages to the CE toward C-S.
- (C-S,C-G) flow arrive at both the primary and the standby upstream PEs.
- Only the primary upstream PE (Virtual Router Master according to VRRP) forwards (C-S,C-G) flow to downstream PEs through a P-tunnel.
- Other private implementations for DF selection (which should be deployed per VRF) could also be optional.

Downstream PE Behavior

- Standby C-multicast route advertising described in [RFC9026] is still necessary.
- Standby PE Community is no longer necessary.
- Downstream PEs recognize the Upstream DF Selection behavior (by using some provisioning methods) and execute the “ANYCAST” RPF checking.
- Downstream PEs accept the C-flow from any of candidate upstream PEs and forward it to CEs, the upstream DF selection prevents the C-flow duplication in backbone.

Next Steps

- Seek for comments and discussions.

Thanks