

# PIM VRRP Interoperability

draft-zhou-pim-vrrp-00


Wei Zhou

[weizho2@cisco.com](mailto:weizho2@cisco.com)

# Rationale

- PIM has no inherent redundancy capability.
- PIM operation is completely independent of First Hop Redundancy Protocols (FHRP) such as VRRP and HSRP.
- There is a need to enable IP multicast forwarding resiliency in redundant network with VRRP enabled.
  - Make Master router the DR after VRRP failover such that Master router is the one responsible for maintaining mroute states and multicast forwarding.
  - Make VRRP Master router the one responsible for processing PIM J/P messages targeting VRRP virtual IP address.

# Use Case



Downstream router X has a static route with C as next-hop.

- A and B are routers running VRRP
- C is the virtual router address
- A and B need to adjust PIM DR priority depending on who is Master
- Need interaction between PIM and VRRP

# Proposed Solution

- Allow PIM to track VRRP group on an interface
  - Leverage VRRP's capability to track uplink status
- PIM DR priority adjustment
  - Adjust DR priority to a configured value when becomes VRRP Master, allows DR be predictable before and after a switchover
  - Make VRRP Master the DR, process IGMP Join and start forwarding traffic

# Tracking & Failover

- PIM keeps tracking VRRP state
  - Upon VRRP switchover, new Master sends PIM Hello with new GenID using Virtual IP as source address
  - Trigger downstream routers to send PIM Join to virtual IP
  - Only Master router will process PIM Join, create mroute state and pull traffic from upstream and start forwarding
  - Backup routers ignore PIM Join/Prune messages targeting the Virtual IP

# PIM Assert

- If only one VRRP group, Backup routers will send a large penalty metric preference (`PIM_ASSERT_INFINITY - 1`) and make MR the Assert winner.
- If there are multiples VRRP groups configured on an interface, Assert metric preference will be (`PIM_ASSERT_INFINITY - 1`) if and only if all VRRP groups are in Backup.
- If there is at least one VRRP group in Master state on an interface, then original Assert metric preference will be used.

# BiDir Group

- If only one VRRP group, Backup routers will send a large penalty metric preference in Offer (`PIM_BIDIR_INFINITY_PREF- 1`) and make MR the DF winner.
- If there are multiples VRRP groups configured on an interface, Offer metric preference will be (`PIM_BIDIR_INFINITY_PREF- 1`) if and only if all VRRP groups are in Backup.
- If there is at least one VRRP group in Master state on an interface, then original Offer metric preference to RP will be used.

# Further Considerations

- Support of HSRP
  - The proposed scheme can also enable HSRP aware PIM with similar switchover and tracking mechanism described here