

Multi-point BFD in VRRP use case

draft-mirsky-bfd-p2mp-vrrp-use-case

Greg Mirsky (ZTE)
Jeff Tantsura (Apstra)
Gyan Mishra (Verizon)

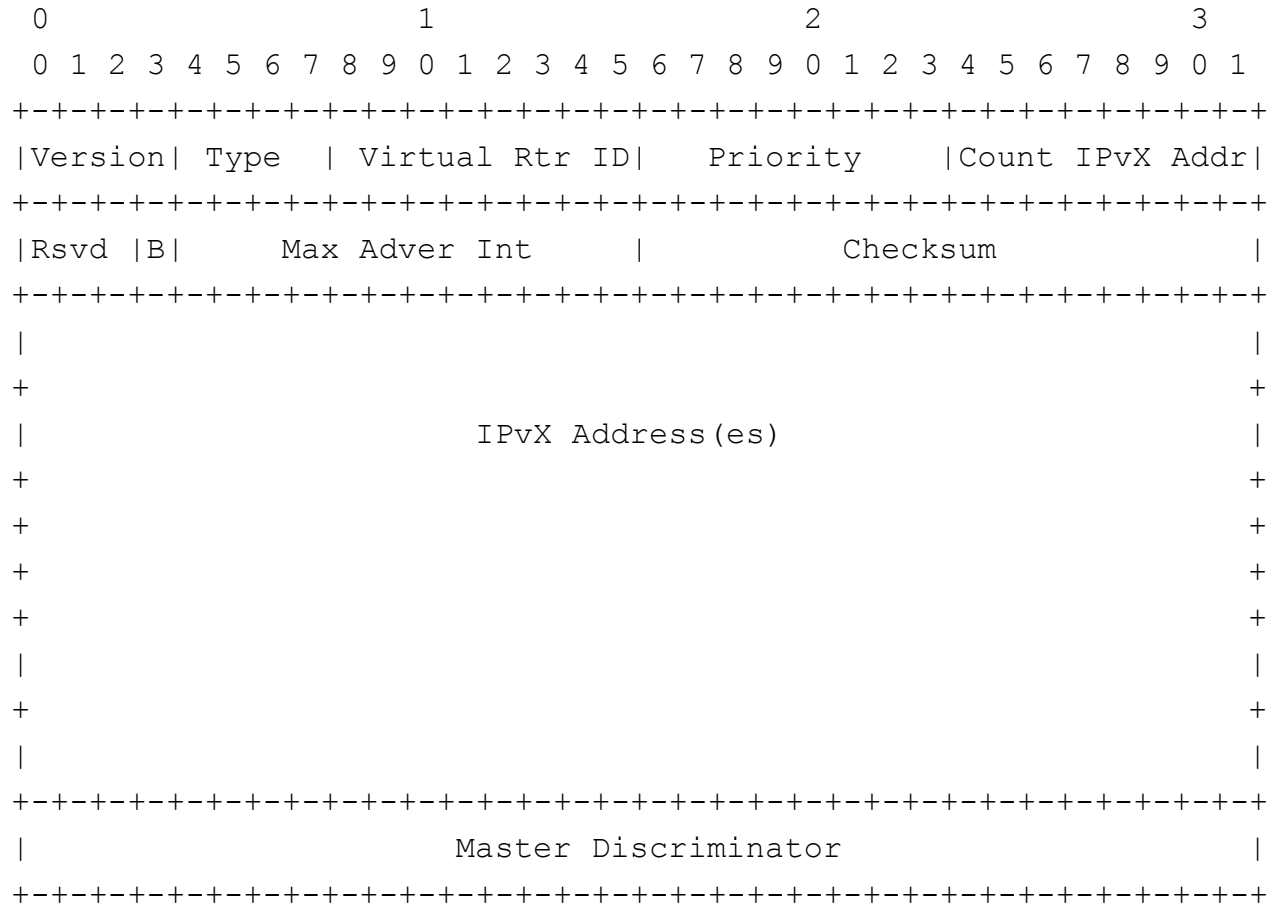
Motivation (recap)

- RFC 5798 supports 10 msec Advertisement_Interval thus enabling Master_Down_Interval to be at ~33 msec
- Backup router monitors Master
- Master is not aware of the identity of Backup router(s)
- Monitoring availability of the Master is another incarnation of Continuity Check OAM
- RFC 5880 defined BFD
- Asynchronous mode can be used by a Backup to monitor the availability of the Master
- Concerns with use of Async RFC 5880:
 - Master doesn't know the identities of Backup routers in the given VRID on the same LAN segment
 - Master is not interested in Backup state change. That can be mitigated by Master setting its `bfd.RequiredMinRxInterval` to 0

Proposed solution with update

- BFD for Multipoint Networks (RFC 8562) extends BFD for use in multipoint and multicast networks
- The Master advertises its BFD My Discriminator in VRRP control packet (extension to RFC 5798) to bootstrap p2mp BFD session
- A Backup creates MultipointTail session and uses My Discriminator value from VRRP control packet and Master IP address to demultiplex BFD Control packets
- As a result, the p2mp BFD session per VRID has a Master router as root and Backup routers as tails. The Master uses one of the IP addresses associated with the VRID as the source IP address for p2mp BFD Control packets
- In p2mp BFD, there's no three-way handshake to initialize the BFD session, and thus the Your Discriminator field is always zeroed

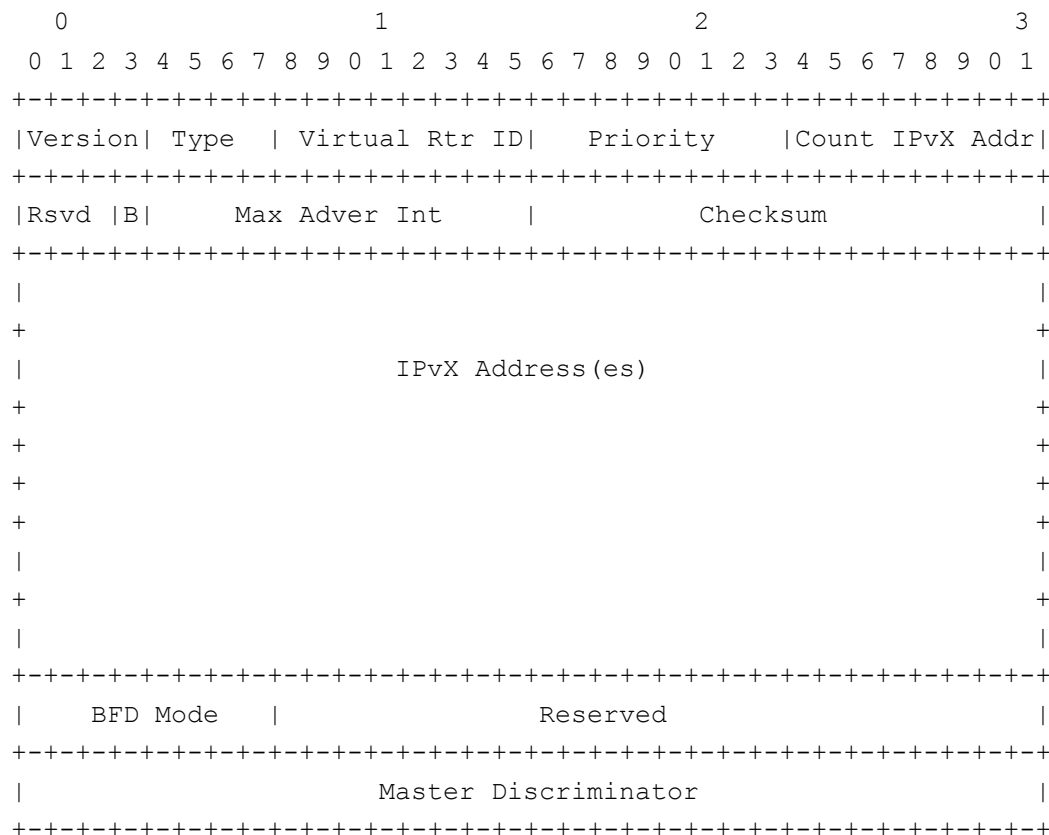
VRRP Control packet to support BFD



How it suppose to work (with the update)?

- Master router
 - creates BFD session of MultipointHead type
 - allocates My Discriminator:
 - begins transmitting VRRP control packets with My Discriminator;
 - begins transmitting BFD Control packets with:
 - Virtual Router MAC address, as specified in RFC 5798, as source MAC address
 - one of Virtual router IPv4/IPv6 addresses as the source IP address
- a Backup router receives VRRP control packet with My Discriminator and creates a new BFD session of MultipointTail type
- the tail BFD node on the Backup router demultiplexes BFD Control packets to particular VRID based on source IP address and My Discriminator
- when value of the bfd.SessionState on MultipointTail is Down, which can be interpreted as the Master router is not Up.
- As a result, the Backup router may elect itself as a new Master router

VRRP Revertive behavior



BFD Mode – one octet-long field. It defines the mode of a tail. Optional values:

- Silent (RFC 8562)
- Active, Polling Multicast/Composite (RFC 8563)
- Active, Unsolicited Notification (draft-mirsky-mpls-p2mp-bfd)

Next steps

- Comments are welcome
- Ask for WG adoption