

BFD for Multipoint Networks over Point-to-Multi-Point MPLS LSP

draft-mirsky-mpls-p2mp-bfd-01

Greg Mirsky

IETF-99 July 2017, Prague

BFD for multipoint networks

- BFD for multi-point network uses Demand mode, defined in RFC 5880, from the very start – no three-way handshake
- Only root transmits BFD control messages with its My Discriminator and Your Discriminator set to 0
- Leaf cannot demultiplex BFD sessions by Your Discriminator as in Asynchronous mode that is used in MPLS LSP per RFC 5884
- Leaf uses 3 tuple <My Discriminator, source address, identity of the multipoint tree> to demultiplex BFD sessions

IP/UDP Encapsulation

- Destination IP address randomly chosen:
 - from 127/8 range for IPv4
 - from 0:0:0:0:0:FFFF:7F00/104 range for IPv6
- Destination UDP port number 3784
- Source UDP port number from 49152 through 65535 range

Non-IP encapsulation

- Overhead of IP/UDP encapsulation, especially with IPv6, is significant
- Cannot use G-ACh type BFD Control, PW-ACH encapsulation (without IP/UDP Headers) 0x0007 defined in RFC 5885
- Use G-ACh type MPLS-TP CV message 0x0023 defined in RFC 6428
- Define Source MEP ID IP Address TLV
- Assign new type from CC/CV MEP-ID TLV registry

Bootstrapping BFD session over p2mp MPLS LSP

- LSP Ping with BFD Discriminator with Target FEC TLV from sub-TLVs defined in Section 3.1 RFC 6425
- BGP-BFD Attribute as defined in "Multicast VPN fast upstream failover"

Next steps

- Your comments, suggestions, questions always welcome and greatly appreciated
- WG adoption