

Multipoint BFD overview and solution space

David Ward
July 2010

Multipoint BFD Overview

- Verifies connectivity of head->tail multipoint path
- Technology independent (IP mcast, MPLS P2MP, etc.)
- Does not verify tail->head return path
- Does not verify unicast head->tail path
- Optional notification to head of tail status
- Protocol timing/scalability driven entirely from head
- Runs next to Classic Unicast BFD
- Falls out of existing Unicast BFD spec (pretty much)

Session Types

- Operation modeled as distinct session types:
 - **Point-To-Point:** Classic BFD
 - **Multipoint Head:** Session on head sending multipoint packets
 - **Multipoint Client:** Optional session on head tracking individual tail
 - **Multipoint Tail:** Session on tail tracking head

Original MP Service Definition

- Base function plus a number of options
- Options may be enabled in any combination
- **Base function: Unidirectional Transmission**
 - Head sends periodic packets along MP tree
 - based on the discriminator distributed and specific to the head
 - Tails detect BFD timeout, do "the right thing" (e.g. listen to another head)
 - Head ignorant of tails, no BFD packets sent to head
 - Simple, extremely scalable

Service Definition - option 1

- Option: **Solicit Membership**
 - Head sets Poll bit in MP transmission
 - Tails send unicast Final in reply
 - Tail transmission smeared across time specified by head
 - Head gets a Pretty Good idea of tails listening (unreliable)

Service Definition - option 2

- Option: **Tails notify head of session failure**
 - Head directs tails to send periodic packets to head when tail detects session failure
 - Upon session failure, tail sends bfd.DetectMult packets (smeared across time) and then quiesces
 - Semi-reliable (multiple packets are sent)

Service Definition - option 3

- Option: **Verify Connectivity of Specific Tail**
 - Head sends unicast Poll Sequence to specific tail (learned by solicitation or outside means)
 - Tail replies with Final (and without smear, so it's quick)
 - Head reliably learns tail state (if tail ever replies)

Service Definition - option 4

- Option: **Some Tails are More Equal Than Others**
 - Side effect of unicast Poll Sequence is that intervals carried therein override multipoint values
 - Head can thus raise transmission rate of individual tails for failure notification

Service Definition - option 5

- Option: **Silent Tails**
 - Tails may be provisioned to never reply to BFD even when head sends Polls
 - Allows for large numbers of second-class citizens in class-conscious tail population

Demultiplexing

- Multipoint (M) bit flags multipoint packets
- Packet demuxing rules select session
- Session type determines elements of procedure

Protocol Tricks and Hackery

- Multipoint packets all sent with Demand (D) bit set, **tails cannot send periodic packets while session is up**
- Required Min RX value set to zero means "no periodic transmission ever" (controls failure notification)
- Silent Tail = 1 means "no transmission ever" (no reply to polls)

Environmental Assumptions

- Tail needs to be able to differentiate between packets received on different MP trees if same head is going to be heard from on multiple trees
 - **Via discriminators specific to the head**
- Head is identified by source address

Next Steps

- Reminder M bit already exists in base spec – no need to revisit
- Could add use of p2mp procedures in p2p scenario?
- LSP-ping extensions needs to be fully addressed

<http://tools.ietf.org/html/draft-ietf-mpls-p2mp-lsp-ping-10>

- We'll reissue in the next month
 - Current version:

<http://tools.ietf.org/html/draft-katz-ward-bfd-multipoint-02>

- MIB work?