

BFD for 1 + 1 protection schemes
with point 2 point adjacencies:
post-MPLS_WC meeting

David Ward

March 2010

Unidirectional LSP solution space

Unidirectional LSP solution

- **Use BFD for unidirectional LSPs with no or minimal change**
 - draft-ietf-bfd-mpls (will you allow it by config as well?)
- **4 sessions: 1 for each of the 1 + 1 links (rx and tx)**
 - **Each session is responsible for monitoring the live-ness of one unidirectional link**

The multi-hop “control response” channel should be over corresponding return path but, in a failure situation; the “control response” channel would come over paired path

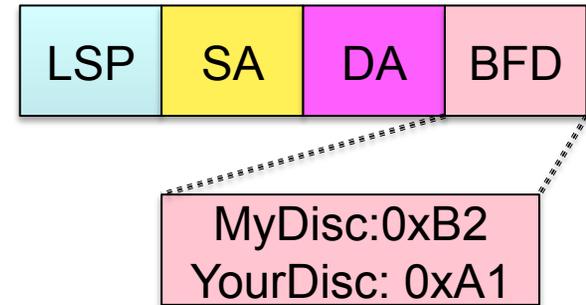
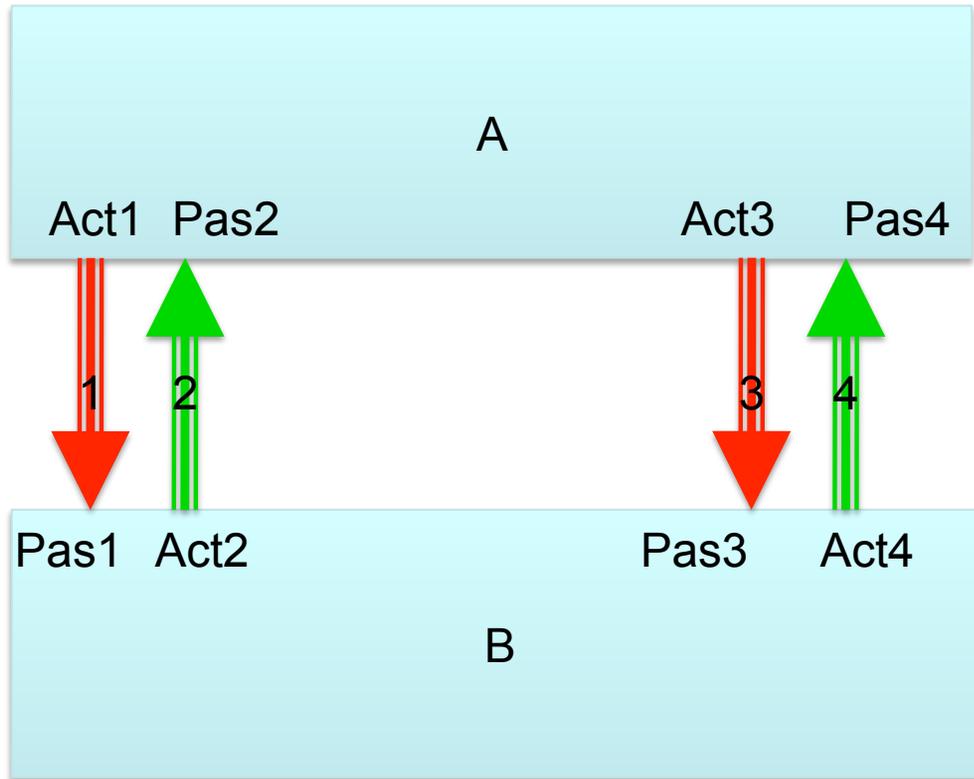
Pros:

- **Enables slow start easily**
- **Enables ability to transmit diags**
- **Completely BW compat and widely deployed**
- **Technique works in all MPLS environments**

Cons:

- **requires at least one return path to restart a BFD session**

Unidirectional LSP solution space



Summary



- Tx for session B->A over link 2
- Rx for session A->B over link 1 or link 3 in case of failure

- BFD control packets are encapsulated in the MPLS label stack that corresponds to the FEC under fault detection.
- Egress LSR is single hop from BFD perspective, TTL is set to 1.
- Egress LSR routes BFD packet based on the destination IP address.

Multipoint BFD overview and solution space

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)

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

Session Types

- Operation modeled as distinct session types:
 - PointToPoint: Classic BFD
 - MultipointHead: Session on head sending multipoint packets
 - MultipointClient: Optional session on head tracking individual tail
 - MultipointTail: Session on tail tracking head

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

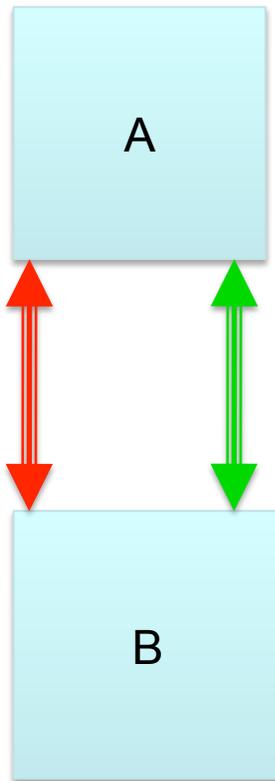
- Tail needs to be able to differentiate between packets received on different MP trees if same head is going to be heard from on multiple trees
 - Via discriminators specific to the head
- Head is identified by source address (which cannot change)

Solution overview

- In 1 +1 environment create 4 independent sessions: Head end (tx) driven
 - No fundamental savings
- Enables receiver to have CC/CV independent of bidirectional connectivity
- If use any functionality of options 1-4 then it is directly analogous to BFD for LSP solution though return control not required continuously

NOTE: Need to make changes to enable slow start

Summary for MP for P2P links



Two MP sessions over p2p links

One from A to B

Another (separate) from B to A

Options:

Use 'option 5'

Set D bit

Set rx trans to 0 or make silent tail

Summary

- Existing MPLS-BFD is the choice that requires no extensions or modifications if any return path exists
 - Gives full functionality of slow start and messaging of tail to head
- Case for MP functionality over P2P links is a small, corner case of P2P functionality
 - No return path
 - Functionality of MP option 5 and any other functionality is never required
 - Slow start without return path requires modifications
- Recommendation: use MPLS-BFD procedures