

BFD Encapsulated in Large Packets

Jeffrey Haas, Juniper

jhaas@juniper.net

Albert Fu, Bloomberg

afu14@bloomberg.net

Problem Statement

- In environments where applications running across multiple hops requires a consistent MTU, blackholing may occur if the path unexpectedly changes and some link does not support that MTU.
- BFD PDUs are not large enough to provide this Path-MTU based continuity check.
 - The BFD Length field is one octet!

Proposed Solution

- BFD intentionally does not require that the encapsulating transport protocol (e.g. IP+UDP) have a length that is minimal to contain the BFD packet.
 - The original BFD authors have been here before in IGP-land. ^^
- Solution: Make the underlying packet big and BFD packet is at the beginning of it. Zero pad the rest.

Discussion

- Why zero-pad? Why not one of the usual bit-pattern tests?
 - Do you really want your line card verifying a 9kbyte pattern at BFD line rates?
 - Also leaves room to place other things at end of packet. We've avoided this for extensions... so far.
- Do we need to negotiate this on both sides?
 - I don't think so. Do you have a different opinion?