BFD Proxy Connections over Monitored Links

Brian Snyder
bsnyder@idirect.net

http://tools.ietf.org/html/draft-snyder-bfd-proxy-connections-monitored-links-00

IETF90 Toronto

July 20-25, 2014

Problem

- Offering a layer 2 solution to the Satellite market exposed new challenges.
- L3 routing KA timer settings would result in tradeoff between convergence time versus bandwidth overhead.
- Satellite bandwidth is very expensive, overhead needs to be minimized. (There could be thousands of modems in one network).
- Customers are used to very quick convergence as we were previously directly attached L3 devices.

Possible Solutions 1

- DLEP
  - Designed for MESH environments
  - Multicast complications
  - Not ubiquitous
  - Not a ratified standard
  - Can take time for vendors to implement
    - Limited L3 application support (ex: no BGP)
    - Scalability (<100 connections)
    - VLAN Trunking

Possible Solutions 2

- BFD
  - Not as informative as DLEP (no link state).
  - Proxy design would break security extensions
  - Very “chatty”. Even more so then L3 hello timers. (*)
- (*) No “deal breakers”… because…..

Overview

- The radio link state is monitored for station keeping purposes
  - “Chattiness” could be minimized by spoofing BFD messaging in the radio devices.
- L3 KA timers can be pushed as high as possible. (Ideally they could be turned off).
- Satellite outroutes are a broadcast channel, so IGP routing is a natural choice.
- Asynchronous timer routing protocols are ideal as outroute and inroute have very different characteristics.
  - Inroute design decisions must scale to the thousands.
  - Therefore, one could rely on outroute HELLO packets to drive convergence and tune down inroute as low as possible.
- Hence, EIGRP and IS-IS are ideal.

Proxy Design Goals

- The proxy can ‘sniff’ the traffic to auto-learn about BFD sessions. (0 config)
- “Eat” all the KA BFD packets from external devices. This keeps all the BFD overhead off the monitored (and expensive) link.
- Reply to all the KA packets to keep sessions alive. (If DUM is up)
- Inject BFD control packets (state change events) to connected network equipment in order to communicate DUM status events.

Proxy Details

- Proxy must keep an OAM object per reachable neighbor. Demux by destination MAC.
- Upon intercepting further BFD packets and locating OAM object (Plus internal checks)
  - If DUM is down, drop.
  - If state is :
    - ADMIN_DOWN: Forward on monitored link
    - DOWN: Reply with constructed BFD Packet (Clear your discriminator field)
    - INIT: Reply with constructed BFD Packet (Set state UP)
    - UP: Reply with constructed BFD Packet

Constructed Reply Packets

- Swap source and destination MAC addresses
- Swap source and destination IP addresses
- Swap discriminator fields
- Set UDP checksum to 0 (optimization)

Integration Improvements

- **BFD Timers**: Allowing for connected equipment to configure a very high BFD interval value. This use case puts forth a useful situation where sub-second failure is not needed but where BFD is still very useful. Relaxing timer configuration strictness would help scale.
- **BFD Demand Mode implementation**: Alternative to above, this would allow all the KA processing to disappear -> event driven is more scalable.
- **BFD Protocol:** Adding the notion of a proxier could assist with enabling security support in this use case.

Example Flow

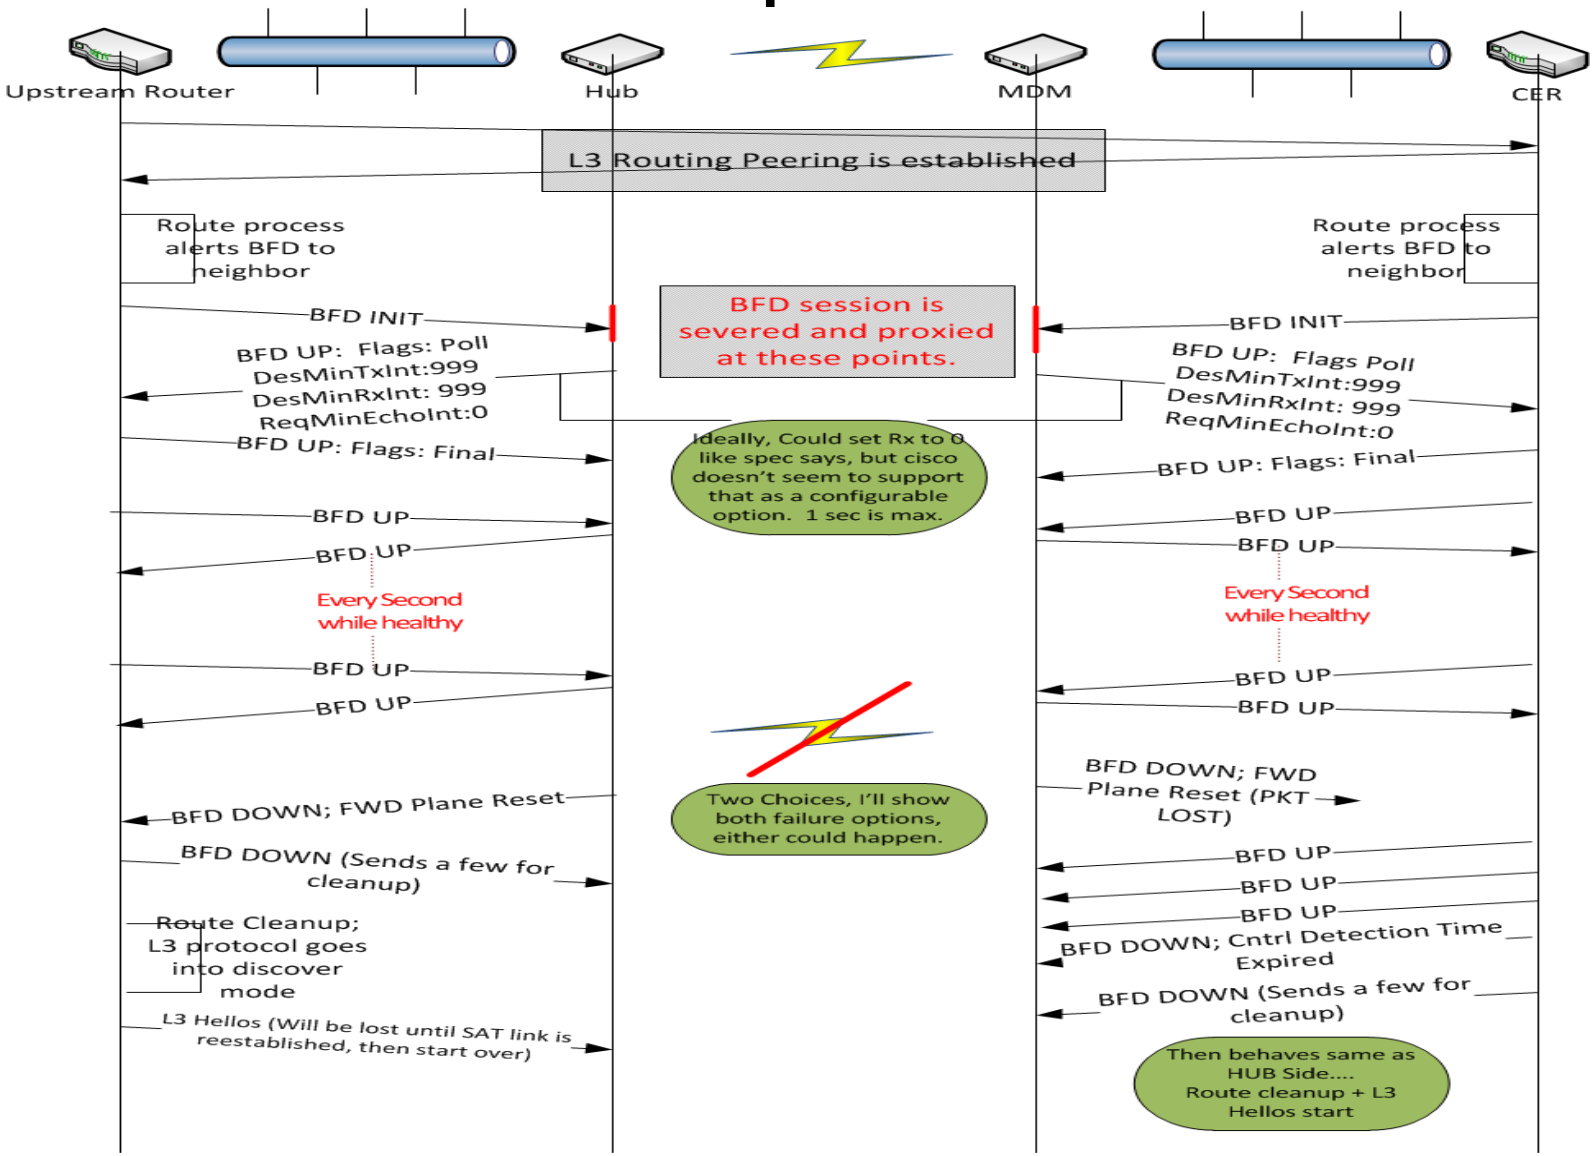

Questions??