Applications of Bidirectional Forwarding Detection (BFD)

Rahul Aggarwal

rahul@juniper.net

Overview

- Goals
- Protocol Overview
- Applications
- BFD's applicability for Voice over IP
- IETF status
- Conclusion

Goals

- Detection of forwarding plane-to-forwarding plane connectivity (including links, interfaces, tunnels etc.)
- A **single** mechanism that is independent of media, routing protocol, and data protocol
- No changes to existing protocols

Goals

- Faster convergence of routing protocols, particularly on shared media (Ethernet)
- Detection of one-way link failures
- Semantic separation of forwarding plane connectivity vs. control plane connectivity

BFD Protocol Overview

- At its heart, Yet Another Hello Protocol
- Packets sent at regular intervals; neighbor failure detected when packets stop showing up
- Intended to be implemented in the forwarding plane to the extent possible (avoiding fate sharing with the control plane)
- Context defined by encapsulating protocol (sending inside IPv4 packets signals IPv4 connectivity; also could be sent over IPv6, directly over the datalink, or whatever)
- Always unicast, even on shared media

BFD Protocol Overview

- Not just for direct links; could also be used over MPLS LSPs, or other unidirectional links (where the return path is potentially routed over multiple hops)
- Sufficient context in the protocol to keep track of multiple parallel paths between systems
- Timing is adjustable on-the-fly, allowing for adaptivity to avoid catastrophic collapse due to false failure detection
- Fully predictive (identical packets sent and received when all is well), making implementation in firmware or hardware more practical

BFD Operation

- BFD Control Packets sent in both directions, providing basic connectivity check and continuous parameter negotiation
- Optional Echo Mode can be negotiated
  - BFD Echo Packets transmitted addressed to originating system; other system forwards them back through regular forwarding path
  - Exercises entire forwarding path in destination system
  - May not always be possible or desirable, thus it is negotiated

BFD Async Mode

- Control packets flow in each direction

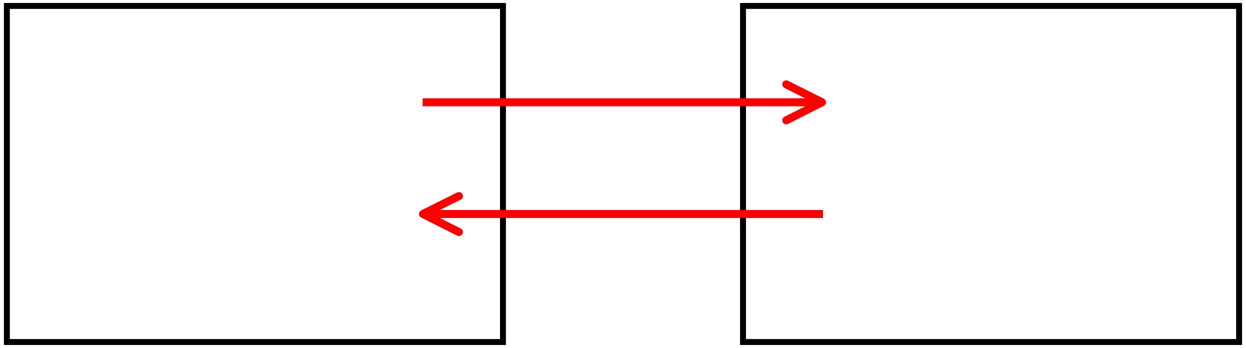

BFD Echo Mode

- Control packets flow in each direction
- Echo packets loop through remote system
- Control packet flow more sedate

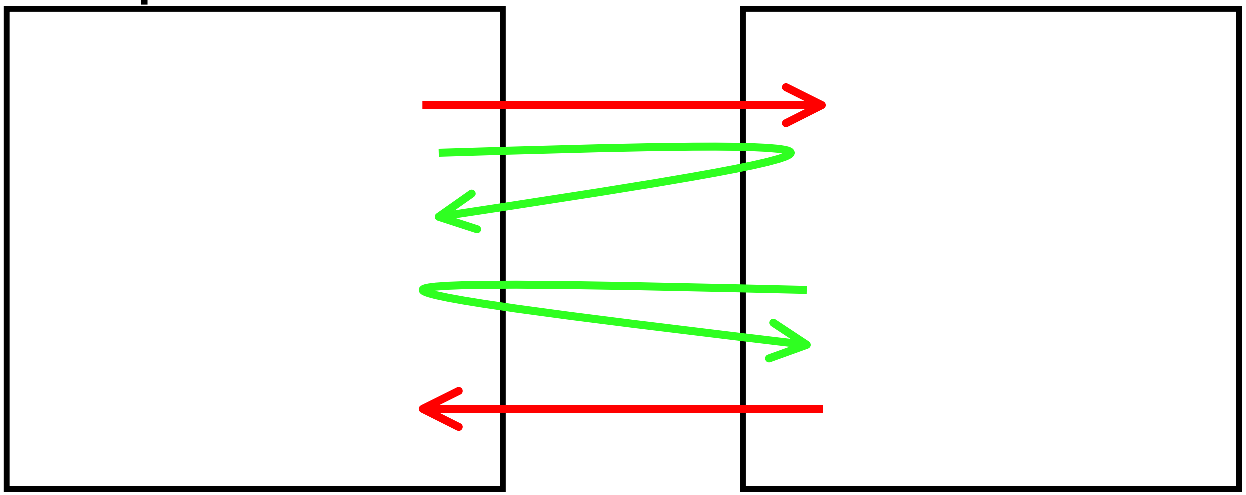

BFD Operation

- No discovery protocol; there are already enough of those (IGPs et al)
- Semantics of BFD session failure are contextual
  - BFD-over-IP implies neighbor failure; IGP neighbor should be torn down
  - BFD-over-Ethernet implies switch failure; subnet should be withdrawn from routing protocols
- Works over a wide range of time constants
  - Timers specified in microseconds, allowing very fast or very slow detection
  - Systems specify how quickly they can receive BFD packets so that boxes of differing abilities can interoperate

BFD Applications

- IGP liveliness detection
- Tunnel liveliness detection
  - MPLS LSPs
  - IP-in-IP/GRE tunnels
- BFD for edge network availability
- BFD over ethernet

BFD for IGP Liveliness Detection

- One of the first motivations for BFD
- Faster convergence particularly on shared media
  - Sub-second IGP adjacency failure detection
- IGP hellos can be set to higher intervals
  - Can improve IGP adjacency scaling

BFD for IGP Liveliness Detection..

- BFD session boot-strapped by OSPF/IS-IS
  - When the adjacency is established
  - First BFD control packet is demuxed based on the incoming IGP interface
- State of the IGP adjacency is tied to the BFD session
  - Adjacency is brought down if the BFD session fails

BFD for MPLS LSPs Motivation

- BFD can be used to detect the liveliness of MPLS LSPs
- Light-weight compared to LSP-Ping
- Sub-second failure detection times
- Periodic fault detection possible
- Fault detection for a larger number of LSPs

BFD for MPLS LSPs Operation

- BFD provides LSP data plane verification
- However BFD doesn't provide data plane – control plane verification
  - LSP-Ping does
- Use BFD in conjunction with LSP-Ping
  - Establish a BFD session for the MPLS LSP
  - Boot strap BFD using LSP-Ping
  - Periodically use LSP-Ping to verify data plane against the control plane
  - draft-raggarwa-mpls-bfd-00.txt

BFD for MPLS LSPs Operation…

- Health of the MPLS LSP tied to the BFD session
- Failure of the BFD session results in
  - MPLS LSP being declared down
  - Alarms can be issued
- Fast-reroute considerations
  - Guard against ‘spurious’ failures
  - BFD detection time should be more than fast-reroute failover time

BFD for MPLS LSPs Applicability

- Transport LSPs
  - Periodic fault detection to meet SLAs
- MPLS PWs
  - In the absence of end-to-end Layer 2 OAM, periodic fault detection may be desired
- Fast-reroute bypass LSPs
  - Protect a number of primary LSPs
  - Fast failure detection of a bypass LSP may be desirable
  - BFD session for the bypass LSP

BFD for MPLS LSPs
Layer 2 Transport over MPLS

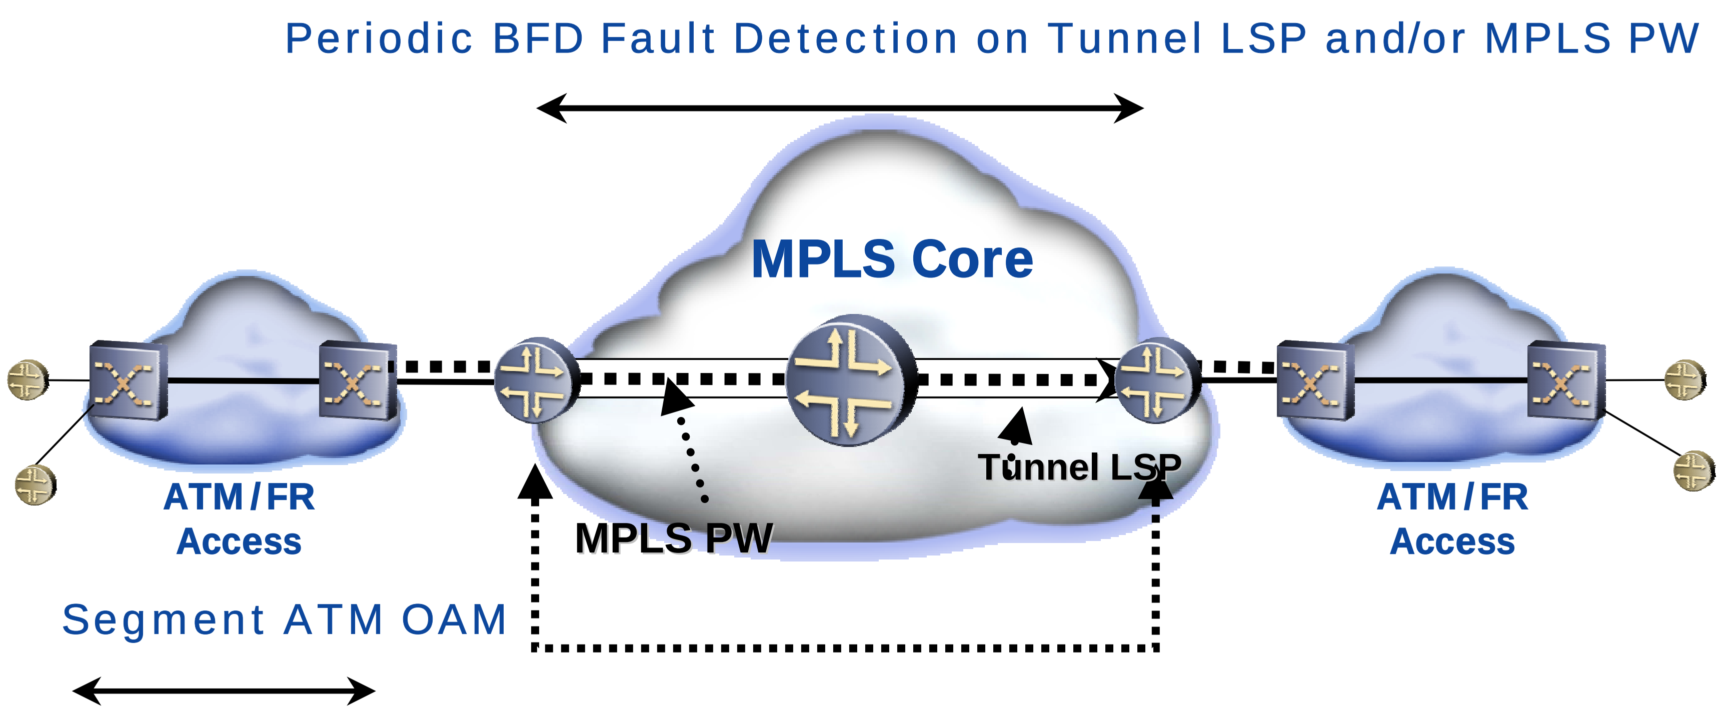

BFD For GRE Tunnels Motivation

- GRE tunnels do not have a liveliness mechanism
- GRE tunnel failure detection depends on IGP
  - Doesn't provide a means to test the tunnel egress's GRE forwarding path
  - Sub-second failure detection is a challenge

BFD For GRE Tunnels Operation

- BFD session is established between the tunnel endpoints
- BFD session boot-strapped using:
  - Configuration or
  - Auto-discovery
- Health of the GRE tunnel tied to the BFD session
  - BFD session failure results in declaring the GRE tunnel down

BFD For GRE Tunnels Applicability

- IP applications over GRE tunnels
- 2547 VPNs over GRE tunnels
- MPLS PWs over GRE tunnels
  - Layer 2 VPNs
  - VPLS

BFD For Edge Network Availability

- Last Mile i.e. CE to PE fast failover is a missing piece
  - Particularly for an ethernet infrastructure
- Often relies on application timeout
  - For eg. EBGP session timeout
- Failure detection in the order of seconds
  - Not acceptable for applications such as Voice over IP
- BFD can be used to fill the Last Mile void

BFD Fills the Last Mile Void

- A big opportunity
  - Driven primarily by Voice over IP
- Router-to-host liveliness detection
  - For eg. Host is a media gateway

Router to Host Liveness Detection

- BFD between router and host
  - The host must support BFD
- Host can be the Voice over IP gateway
- Host can also be a web server
- Support for common routing paradigms between router and host
  - Static routes
  - EBGP peering

BFD for Static Routes

- Host support
  - Dual homed host
  - BFD session to each router
    - Should not be an echo session
  - BFD session failure triggers failover to the backup router

BFD for Static Routes..

- Router support
  - Static route to the subnet with the host as the next-hop
  - BFD session associated with the static route's next-hop.
  - BFD session failure results in de-activating the static route
    - A backup route to reach the host takes over

BFD for EBGP Peers

- CPE support
  - Dual homed CPE
  - BFD session to each EBGP peering router
  - BFD session failure triggers failover to the backup router
- Router support
  - BFD session to the EBGP peer
  - BFD session failure results in declaring the EBGP peer down

Voice Over IP
Meeting SLAs

- End to end failure recovery times of ~300ms
- MPLS backbone
  - Fast reroute provides failure recovery times of ~50ms
- What about the edge ?
  - BFD on the edge allows fast failure detection
  - Pre-setup alternate path allows fast failure recovery

BFD for Edge Availability Voice over IP

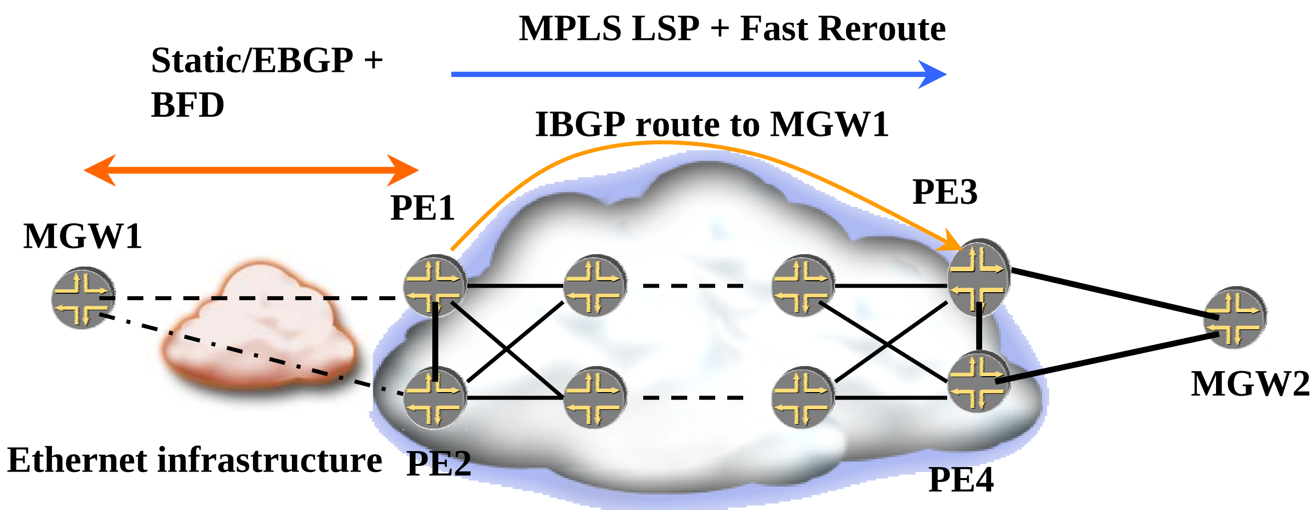

- MGW Media Gateway
- BFD between MGW and PEs
- Enables fast detection/failover

BFD for Edge Availability
Voice over IP

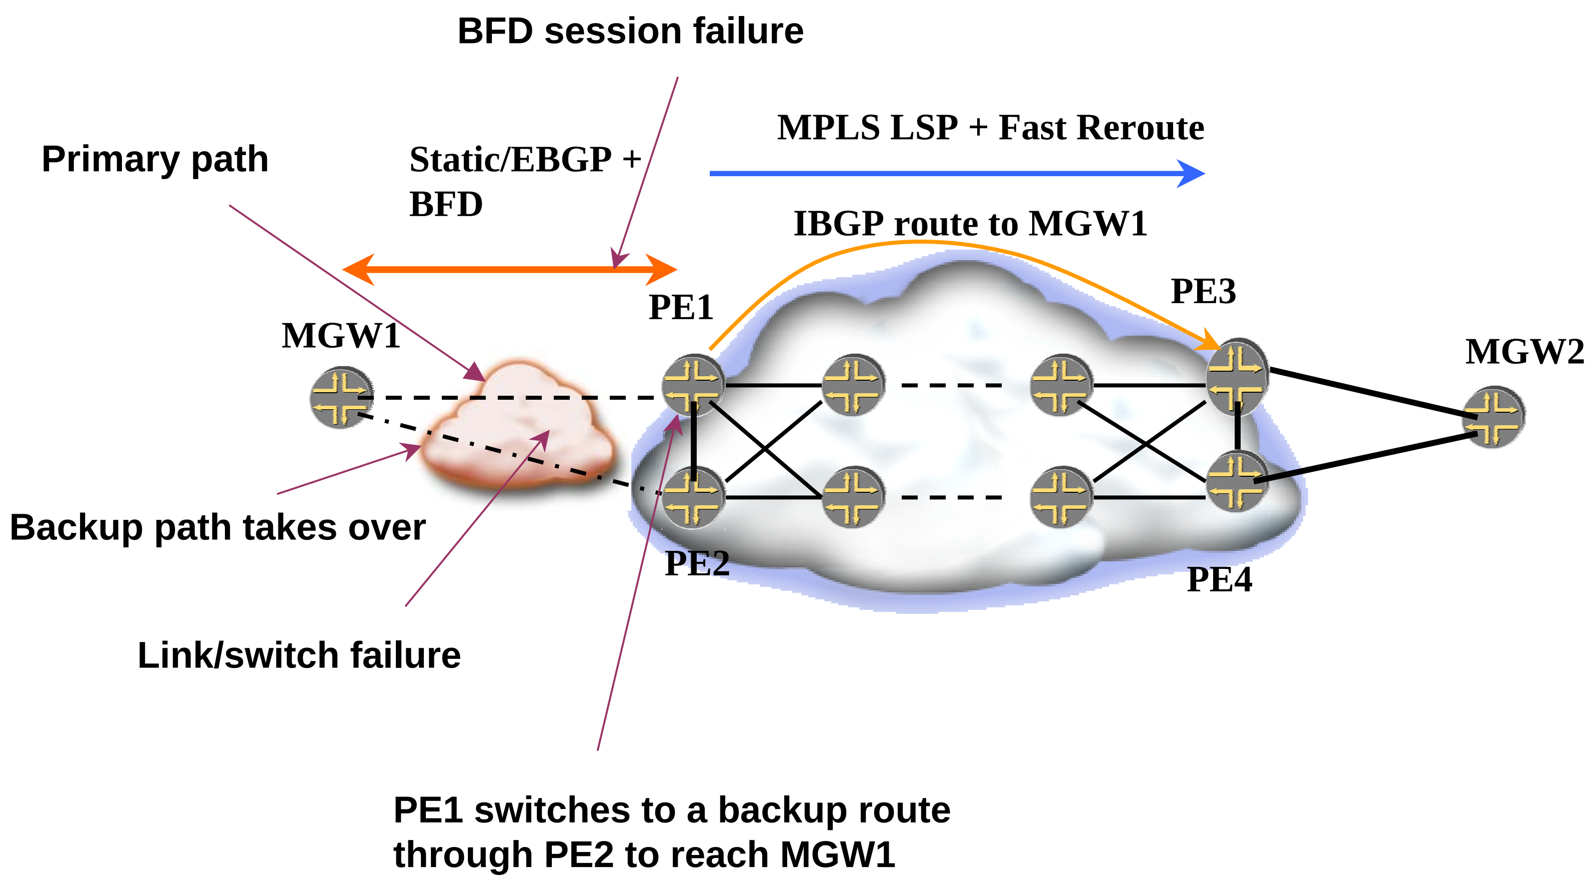

BFD Over Ethernet Motivation

- OAM mechanism for ethernet links is needed
  - Fast failure detection over ethernet not possible currently
- Router to a CE
- Router to an ethernet switch

BFD Over Ethernet Operation

- BFD directly over ethernet i.e. 802.3
- BFD session established between the router and a CE or an ethernet switch
  - Driven by configuration
- BFD control packets can be sent directly over ethernet
  - A new ether type to identify BFD packets is needed

BFD IETF Status

- Protocol jointly developed by Juniper and Cisco
- Base spec: draft-katz-ward-bfd-01.txt
- Over IP: draft-katz-ipv4-ipv6-01.txt
- Over MPLS: draft-raggarwa-mpls-bfd-00.txt
- A BFD WG has been formed

Conclusion

- Increasing interest in the service provider community
- Shipping in JunOS for IGP since 6.1