

**VENDOR
IMPLEMENTATION REPORT
FOR
DRAFT-IETF-IDR-BGP-BFD-STRICT-MODE**

**MERCIA ZHENG
ACEE LINDEM
JEFFREY HAAS
ALBERT FU**

Agenda

1. BGP BFD “Default” Mode
2. BGP BFD Strict-Mode
3. Nokia BGP BFD Strict-Mode Config (SRoS 23.7)
4. Juniper BGP BFD Strict-Mode Config (Junos 23.4)
5. Summary

1. BGP BFD Default Mode (most vendors)

- BFD is initiated after BGP is established
 - BGP Route Updates commence as soon as BGP is UP
- There is a finite time between when BGP is established to BFD coming up
- If there's a “break-in-middle” problem during this time, BGP session will require lengthy Hold time to timeout
 - No BFD for Fast Failure Detection
 - Traffic Blackhole for extended “Hold time” (90sec/180sec)

Sample BGP/BFD Up times on a Production router

Time from BGP Established to BFD UP = 00:51:08 – 00:50:57 ~ **11s**

```
user@router> show bgp summary group bgpGroup
```

```
Threading mode: BGP I/O
```

```
Groups: 4 Peers: 9 Down peers: 0
```

```
Table Tot Paths Act Paths Suppressed History Damp State Pending
```

```
inet.0
```

```
3880005 1190923 0 0 0 0
```

```
bgp.l3vpn.0
```

```
0 0 0 0 0 0
```

```
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn State|#Active/Received/Accepted/Damped...
```

```
10.1.1.1 65001 7635586 93761 0 42 4w1d 22:40:55 Establ
```

```
inet.0: 699605/917238/917238/0
```

```
user@router> show log messages
```

```
..
```

```
May 26 00:50:46.950 2022 router bfdd[25859]: BFDD_STATE_UP_TO_DOWN: BFD Session 10.1.1.1 (IFL 397) state Up
```

```
-> Down LD/RD(171/165) Up time:14w3d 14:36 Local diag: CtlExpire Remote diag: None Reason: Detect Timer Expiry.
```

```
May 26 00:50:46.950 2022 router bfdd[25859]: BFDD_TRAP_SHOP_STATE_DOWN: local discriminator: 171, new state: down,  
interface: xe-7/0/1.0, peer addr: 10.1.1.1
```

```
May 26 00:50:46.984 2022 router rpd[26439]: RPD_BGP_NEIGHBOR_STATE_CHANGED: BGP peer 10.1.1.1 (External AS 65001)  
changed state from Established to Idle (event BfdDown) (instance master)
```

```
May 26 00:50:47.012 2022 router rpd[26439]: bgp_bfd_callback:188: NOTIFICATION sent to 10.1.1.1 (External AS 65001):  
code 6 (Cease) subcode 9 (Hard Reset), Reason: BFD Session Down
```

```
..
```

```
May 26 00:50:57.989 2022 router rpd[26439]: RPD_BGP_NEIGHBOR_STATE_CHANGED: BGP peer 10.1.1.1 (External AS 65001)  
changed state from OpenConfirm to Established (event RecvKeepAlive) (instance master)
```

```
May 26 00:51:08.712 2022 router bfdd[25859]: BFDD_TRAP_SHOP_STATE_UP: local discriminator: 176, new state: up,  
interface: xe-7/0/1.0, peer addr: 10.1.1.1
```

2. BGP BFD Strict-Mode

- BFD is initiated before BGP is “fully” established
- Dynamic Signaling of BGP BFD Strict-Mode in BGP Open Strict-Mode Capability Code 74 (0x4a)
 - “Default” Mode is used unless both sides advertise Strict-Mode Capability
- BFD Strict-Mode Protocol Change:
After BGP OpenSent/OpenConfirm, delay sending BGP Keepalive until BFD session is Up
 - Ensure BFD Fast Failure Detection at all time
- BFD Strict-Mode with “Holddown”:
After BGP OpenSent/OpenConfirm, delay sending BGP Keepalive until BFD session is Up for “Hold time” period
 - Useful in dampening flapping links

3. Nokia BGP BFD Strict-Mode Config (SRoS 23.7)

```
router "Base" {
  interface "eth-1/1/15.3009" {
    admin-state enable
    port 1/x1/1/c15/1:3009
    ipv4 {
      bfd {
        admin-state enable
        transmit-interval 150
        receive 150
        multiplier 3
      }
      primary {
        address 100.1.1.2
        prefix-length 30
      }
    }
  }
  bgp {
    admin-state enable
    group "bgpGroup" {
      admin-state enable
      peer-as 65001
      ..
    }
    neighbor "100.1.1.1" {
      admin-state enable
      bfd-liveness true
      group "bgpGroup"
      bfd-strict-mode {
        next-hop-reachability true
        advertise
      }
    }
  }
}
```

4. Juniper BGP BFD Strict-Mode Config (Junos 23.4)

```
interfaces {
  et-0/0/3 {
    vlan-tagging;
    mtu 9192;
    unit 3009 {
      vlan-id 3009;
      family inet {
        mtu 1500;
        address 100.1.1.1/30;
      }
    }
  }
  ..
}
protocols {
  bgp {
    group ebgp {
      peer-as 65002;
      neighbor 100.1.1.2 {
        bfd-liveness-detection {
          minimum-interval 150;
          multiplier 3;
          strict-bfd;
          holddown-interval 10000; /* msec */
        }
      }
    }
  }
}
```

Junos supports BFD Strict-Mode with “Holddown”, whereby BFD must be up for the configured period before BGP transitions from OpenSent/OpenConfirm to Established state.

Successful BGP Strict-Mode Interop Between Junos/SRoS

- Juniper Strict-Mode (with/without Holddown) could successfully Interop with Nokia Strict-Mode (currently no Holddown)
- BGP Strict-Mode Capability (Option 74) successfully Exchanged

Juniper:

```
user@Juniper> show bgp neighbor 100.1.1.2 | match strict
Strict BFD configured locally
Peer supports strict BFD
```

Nokia:

```
A:user@Nokia# show router bgp neighbor 100.1.1.1 | match Capability
Local Capability      : RtRefresh MPBGP 4byte ASN BfdStrictMode
Remote Capability    : RtRefresh MPBGP 4byte ASN BfdStrictMode
```

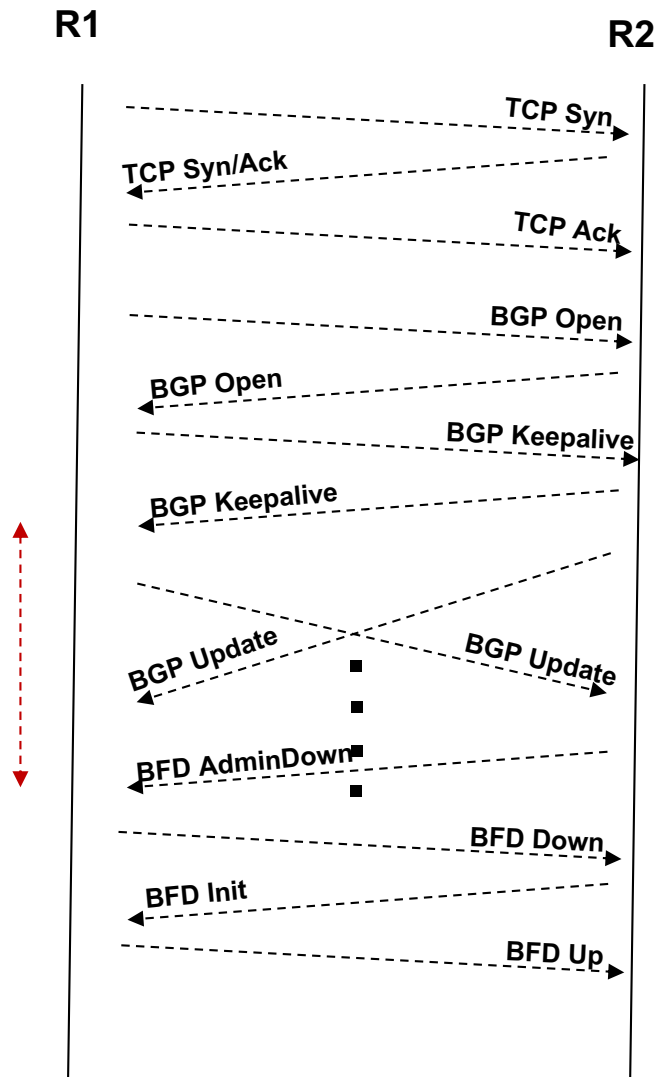

5. Summary

- BGP BFD Strict-Mode ensures that BFD is always available for Fast Failure Detection
 - Validated on Juniper/Nokia (successful Interop)
 - Support expected on 3rd Vendor soon
- Note that RFC9355 specifies similar BFD Strict-Mode for OSPF protocol, as supported by:
 - Juniper
 - Nokia
- BFD Strict Mode is recommended for all protocols

Questions?

Backup Slides..

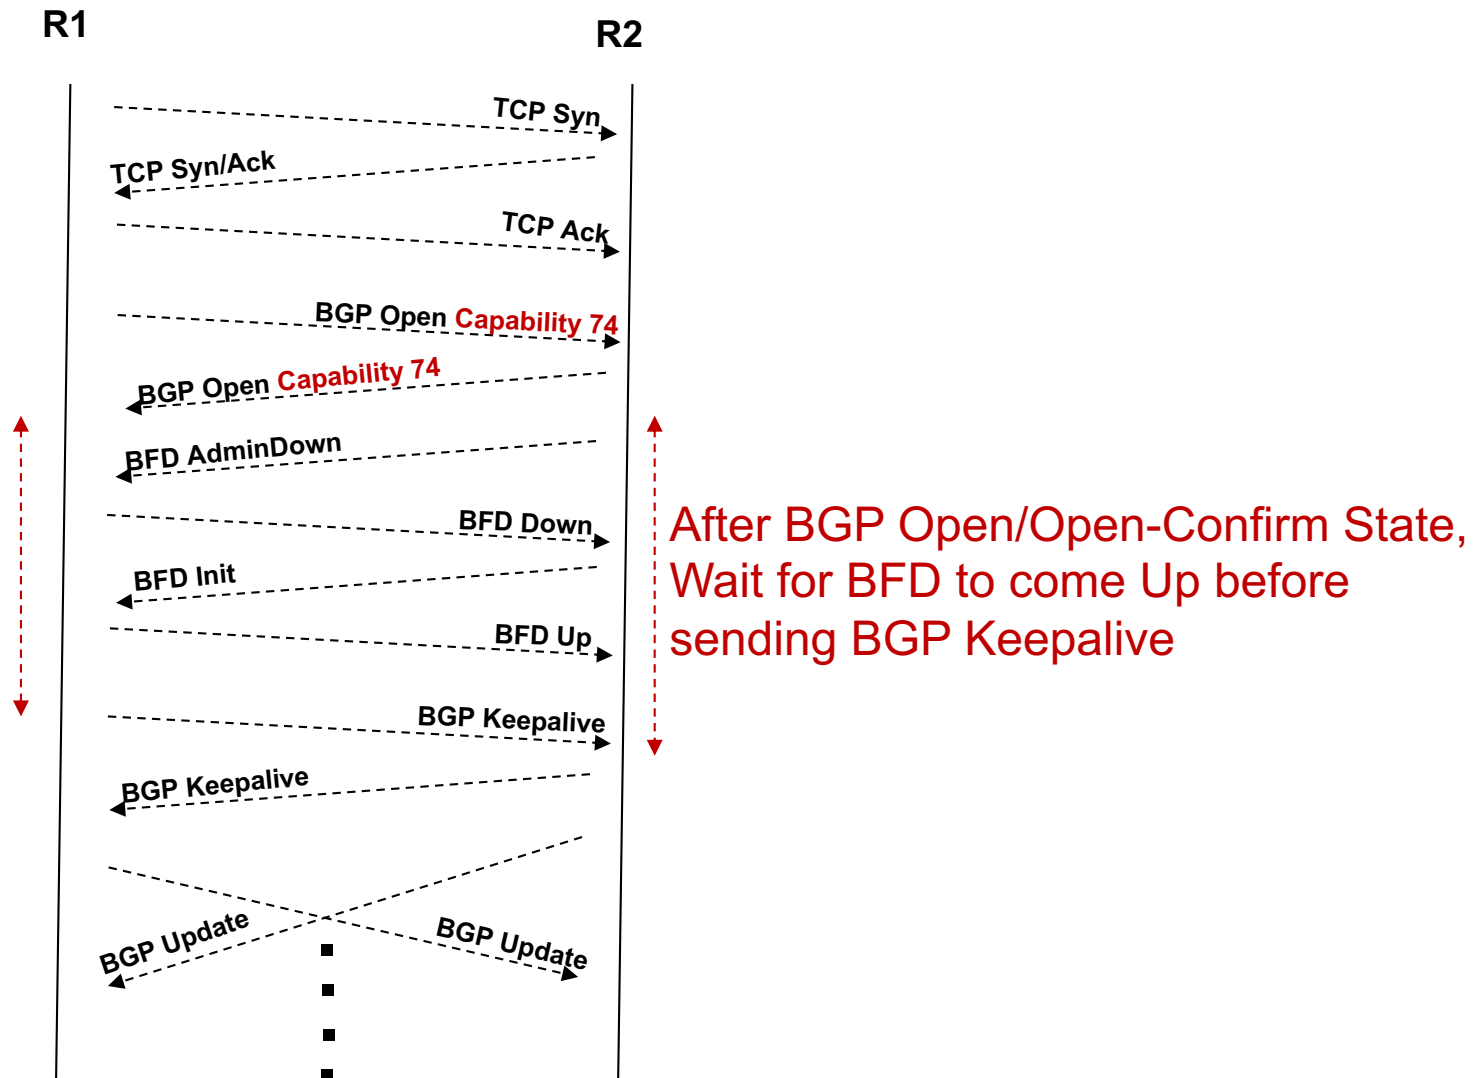
BGP BFD "Default" Mode



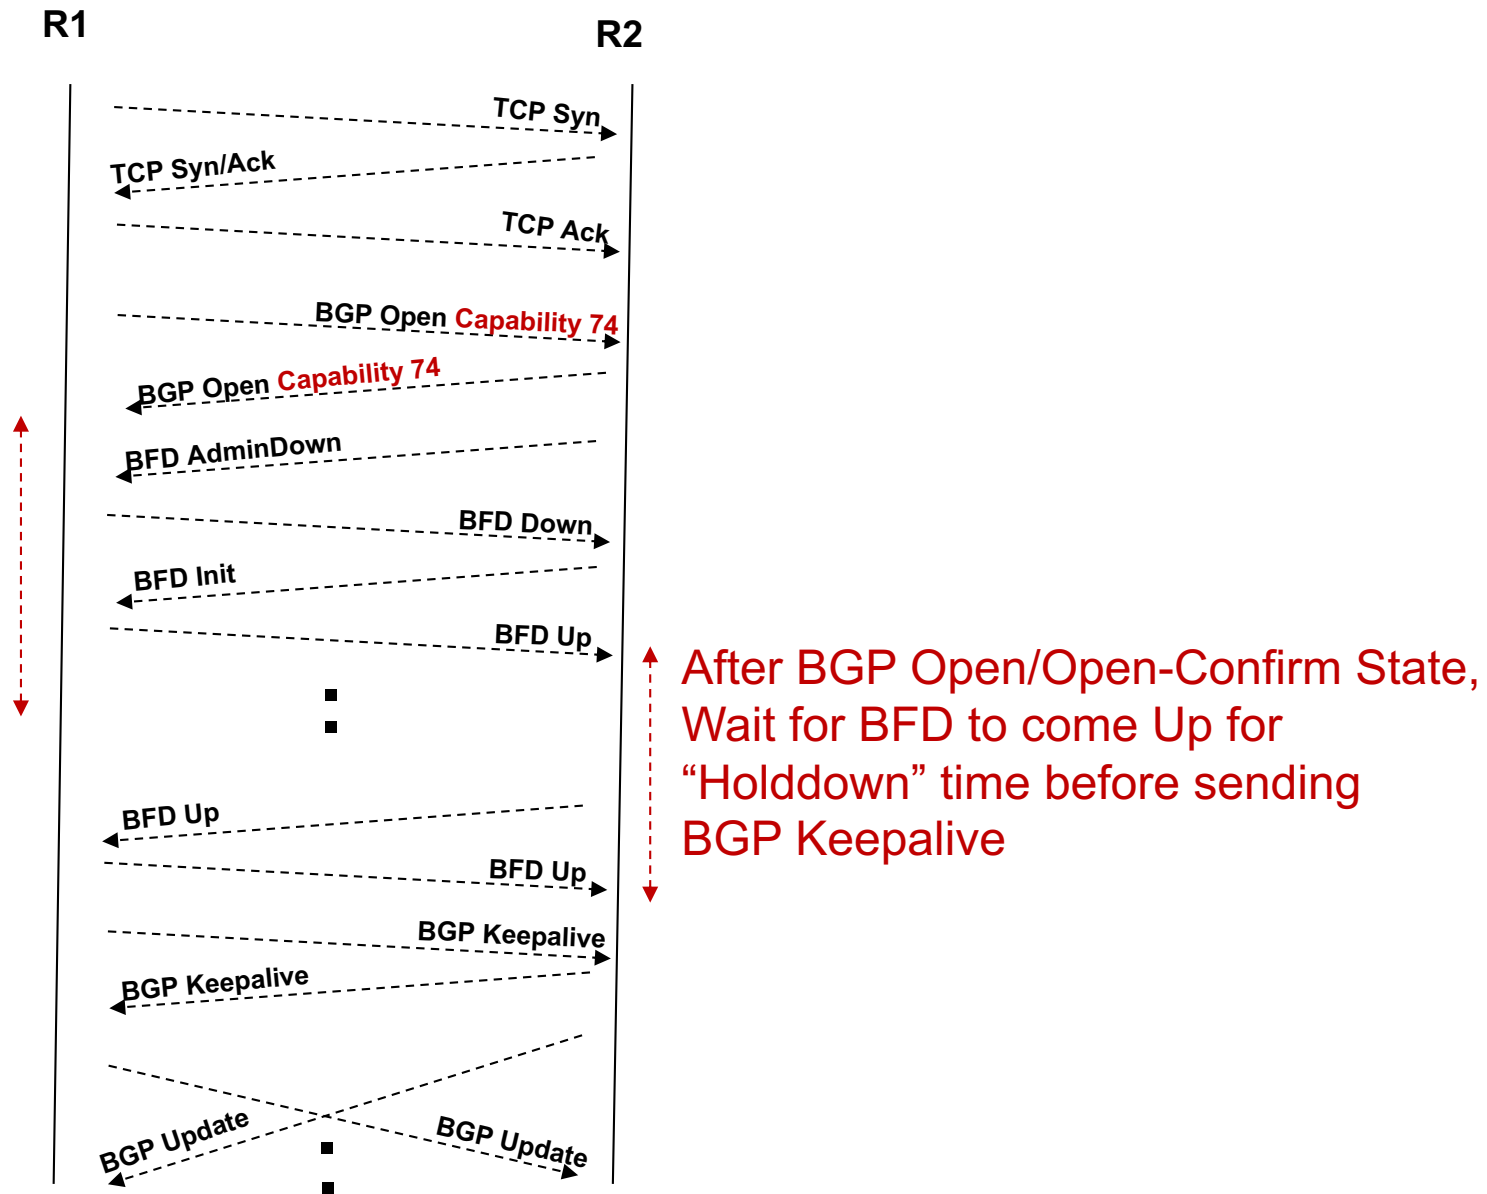
After BGP is Established,
Some delay before BFD
Is initiated.

“Break-In-Middle” failure during
this time will incur lengthy
service outage.

BGP BFD Strict-Mode



BGP BFD Strict-Mode with Holddown



Juniper Proprietary BGP BFD Strict-Mode Config (Available since Junos 15.1)

```
interfaces {
  et-0/0/3 {
    vlan-tagging;
    mtu 9192;
    unit 3009 {
      vlan-id 3009;
      family inet {
        mtu 1500;
        address 100.1.1.1/30;
      }
    }
    ..
  }
}
protocols {
  bgp {
    group ebgp {
      peer-as 65002;
      bfd-liveness-detection {
        minimum-interval 1000;
        multiplier 3;
      }
      neighbor 100.1.1.2 {
        bfd-liveness-detection {
          minimum-interval 500;
          multiplier 3;
          holddown-interval 10000;
        }
      }
    }
  }
}
```

Proprietary implementation does not support Dynamic Strict-Mode Signaling

- Both sides must be configured with Strict-Mode
- BFD must be Up before TCP 3-Way Handshake begins
- Complicates deployment