


draft-jdurand-auto-bfd-00

Path validation toward BGP next-hop  
with AUTO-BFD

Jérôme Durand, Cisco

David Freedman, Claranet

# Problem we want to solve (sec 1)



2 EBGP peerings Established  
Connectivity between peers down  
➔ BLACK HOLE

# Solution requirements (sec 3)

- Solution independent of IXP. IXP members MUST be able to detect and remedy to such issues without anything on IXP
- ➔ Main difference from **draft-ymbk-idr-rs-bfd**
- Other requirements detailed in section 3

# Auto-BFD Solution (sec 4)

- AUTO-BFD is configured on the BGP peering to the BGP RS.
- Every time a new BGP next-hop is received from this peering, AUTO- BFD triggers a new BFD session with this next-hop
  - Asynchronous mode
  - Timers and security configuration can be locally added
- Routes coming from the AUTO-BFD enabled BGP neighbor are then checked based on the BGP next-hop and its BFD session state.
- Acceptance of routes is then subject to the administrative policy based on BFD session state (discard route, change LP...)

**➔ IXP member is in control**

# Session ageing 1/2 (sec 5)

- Important: we don't want sessions to stay for ever
- The tricky part: it must work with asymmetric policies
  - A stops sending routes to B (through BGP RS)
  - B still sends routes to A (through BGP RS)
  - ➔ We don't want B to tear down the BFD session as A would then believe B is down

# Session ageing 2/2 (sec 5)

- IXP members implementing AutoBFD signal they still need the session (ie. that they still receive routes) using bfd.LocalDiag in BFD control packets
- Members do not tear down a session when they receive this flag.
- If there is no flag, members tear down the session after a give timer

**➔ No change to BFD protocol**

# Ask for the WG

- Questions ? Comments ?
- Adopt as WG document

Thank you !