

# Failure and Reachability Detection


shim6 wg, IETF 63 Paris  
Iljitsch van Beijnum

# Reachability Detection

- Multihoming:
  1. detect failures
  2. route around them
- Management of address candidate important but out of scope
- Terminology: address pair = unidirectional path

# Simple Approach

- Simple approach: "ping" all source/dest pairs, pick best round trip time
- Doesn't detect unidirectional reachability
- Doesn't scale well:  $n_a * n_b$



# Complexity++

- Use semi-reliable procedure to indicate probe was received
  - finds unidirectional connectivity
- Use heuristics to pick good candidates, stop probing when results are reasonable
  - avoids (now  $2n^2$ ) exploration explosion
  - but: may not find "best" link

# This is hard work...

- So split off a separate light-weight protocol to confirm previously existing reachability
- Two suggestions for this:
  - "correspondent unreachability detection" (CUD)
  - "forced bidirectional communication" (FBD)

# CUD

- Similar to IPv6 neighbor unreachability detection (NUD):
  - upper layer protocols indicate that there is progress
  - if no ULP progress, do request/response connectivity test
  - if this fails, full reachability exploration

# FBD

- Observation: most protocols generate traffic in both directions
- Make this a rule, generate "keepalive" when otherwise traffic in one direction
- When only packets in one direction: suspected failure, do full reachability exploration

# Tradeoffs

- CUD needs positive ULP feedback to be efficient
- f-bidi needs no cooperation from ULP but may be able to use negative ULP feedback
- FBD only adds packets in the otherwise idle direction



# More Tradeoffs

- CUD: sender detects failures in either direction
- FBD: receiver detects failure in incoming direction
  - depend on other side to detect failures in our outgoing
- Uncooperating ULPs: CUD always probes

# Granularity

- Detect host-to-host reachability
- Detect ULID-to-ULID reachability
  - all traffic between hosts/ULIDs over same locator pair!
- Detect locator-to-locator reachability
- Sweet spot probably at ULID-to-ULID but reuse locator-to-locator info

# Heuristics and Optimizations

- When other side initiates full reachability exploration, we should probably do it too
  - They may know something we don't
- Three way check confirms both directions, less packets but rely more on other side
- What's happening to other sessions?
- ICMP, layer 2 and 3 events...

# NATs and Firewalls

- NATs: protocol may be reused for IPv4
- Want fate sharing with real traffic
  - Can't be done: we'd have to be TCP or IP option, but IP options == firewall trouble
- Fate sharing with shim (or other?)  
negotiation doable, need to encapsulate
- Do we want to sneak by firewalls?

# Security

- Probably don't want to do reachability detection for completely unknown correspondents, but bar not too high
- So need room for auth data of some sort
- Must be light-weight!
- Avoid amplification and redirection attacks

# To Do

- Integrate new draft with Jari's
- Send text!