

# DS-Lite Failure Detection and Failover

draft-tsou-softwire-bfd-ds-lite-04

Tina Tsou

Brandon Li

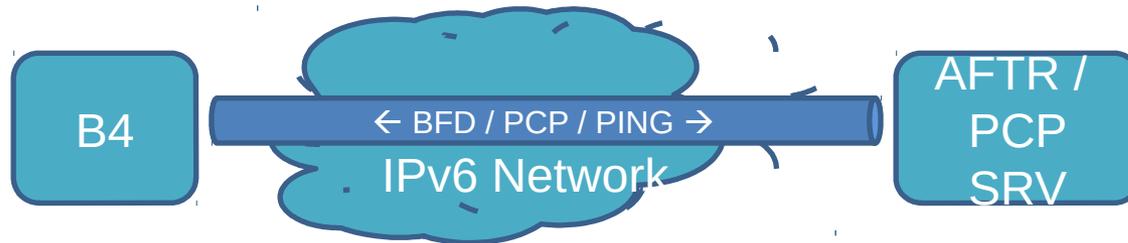
Cathy Zhou

Jürgen Schönwälder

Reinaldo Penno

Mohamed Boucadair

# Problem to Solve



- There is no status information of the DS-Lite tunnel, e.g., tunnel up or down, which brings difficulties for operations and maintenance.
- Protocols to resolve this problem: BFD, PCP, ICMP, ...

# BFD for DS-Lite Failure Detection

- BFD auto configuration
  - In DS-Lite, B4 has the AFTR address, which is sufficient to initiate a BFD session
  - Other parameters can be negotiated via signaling or static config, no manual configuration
- BFD packet rate
  - Long time period between BFD packets transmission, e.g., 10s or 30s

# PCP for DS-Lite Failure Detection

- If PCP is available in a DS-Lite deployment, PCP can be used for keep alive testing, and to trigger failover if a failure is detected.
- PCP is used to create a mapping with short lifetime, updates are sent periodically.
- If the PCP client detects a failure, e.g., a `NETWORK_FAILURE` error code is returned, or there is no response from the PCP server, the client will switch to another PCP server or AFTR

# ICMP for DS-Lite Failure Detection

- ICMP pings can be sent periodically, or triggered manually when necessary
- Since ICMP is an integral part of any IP stack, not extra implementation efforts are required on the B4 elements.

# Failover - Explicit

- The FQDN of the AFTR is sent to the B4 element via a DHCP option [RFC6334]
- The FQDN can be configured to resolve to multiple AFTR addresses
- Upon detecting a failure, the B4 element terminates the current DS-Lite tunnel and creates another tunnel to one of the other addresses.

# Failover – Anycast

- AFTR may use anycast address for receiving packets
- There might be issues with failures on the path if for example an ICMP error message fails to get delivered correctly (e.g., it is send to a different anycast server)

# Next steps

- Adopt?
- Need more reviews