

# MPLS Fast ReRoute as Redundancy Indicator

Pete Templin

NANOG46

# Problem: “Are we redundant?”

- PHB: “Are we redundant?”
- NANO: “Well, boss, I think so, but there are so many paths through the network that I don’t know.”
- Idea: Let’s *virtually* consume some capacity as an indicator of capacity margin

# Idea

- TE tunnels that may or may not carry your traffic, but consume TE bandwidth across TE-eligible links in a manner representative of reality.
- FRR tunnels that may or may not protect your TE tunnels, but consume TE bandwidth across TE-eligible links in a manner representative of the TE bandwidth needed.

# Strategy

- Assumptions:
  - MPLS TE is in place.
  - Some method of tunnel bandwidth adjustment.
- Method:
  - One or more low-priority FRR tunnels in parallel with your existing TE method (tactical/strategic).
  - FRR tunnels may or may not need to exclude their protected interface, based on your needs.

# Tactics

- FRR tunnels should “consume”  $X\%$  of the total available RSVP bandwidth on the link they’re protecting.
  - $X=50$  for  $2N$  protection,  $X=??$  For  $N+1$ , etc.
- Perhaps 3-4 tunnels - 25%, 12%, 6%, 6%?
- Have a “worse” priority than TE tunnels, so the FRR tunnels preempt out of the way first.

# Outcome

- Low-priority FRR tunnels will seek a path through the network with enough bandwidth.
- If unable to find room, tunnels will go down, signifying insufficient **redundant capacity**, leaving room for primary (traffic-bearing) paths.

# Gotchas

- Overlay this tactic with existing FRR using Extreme Care™.
- You can deploy FRR without telling your TE tunnels to request protection, isolating the FRR redundancy indicator from your TE mesh.
- Your mileage may vary.