Tactical Traffic Engineering

draft-li-rtgwg-tte-00

T. Li & C. Barth, Juniper Networks

A. Smith & B. Wen, Comcast

Strategic vs tactical resource optimization/management

- Conventional traffic engineering approaches for resource management used by RSVP-TE and SR-TE often leverage estimates of the ingress traffic demands, during path placement
  - Path placement strategy is to avoid potential congestion

- However, unforeseen and/or dynamic events, can skew these estimates by significant enough margins to result in unexpected network congestion
  - Recomputed paths that address the new demands may take a considerable amount of time, leaving the network in a sub-optimal state

Real time TTE (1/5)

- Set of mechanisms that would allow the network to react in real-time to avert congestion and optimize traffic flow
    - Recognizing congestion
    - FIB entries & backup paths
    - Activation / deactivation
    - Mitigating downstream congestion
    - Flow distribution & selection

Recognizing congestion

- When is link is nearing congestion and when has congestion abated
  - Each link that is protected by TTE is sampled periodically for its current utilization
  - The boundaries of acceptable utilization are defined by high and low utilization thresholds
  - To avoid oscillation, the link must be outside of acceptable utilization for some consecutive number of periodic samples before any action is performed

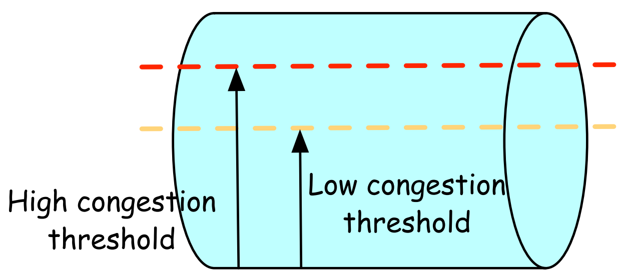

FIB Entry & backup paths

- Flow manipulation
  - TTE manipulates traffic flows by changing the IPv4 / v6 prefixes found in the Forwarding Information Base (FIB), or by changing label entries found the Label Forwarding Information Base (LFIB)
- Several mechanisms exist that potentially create backup paths for a single flow (LFA, FRR, TI-LFA, …)
  - A key property of a backup path is that its loop free and avoids the same link that the primary path is using
- TTE makes use of backup paths by turning them into active paths in parallel with the primary path.
  - This creates an Equal Cost Multi-Path (ECMP)

Activation / deactivation

- Activation - TTE selects a flow and makes appropriate data plane changes so that traffic is balanced between the primary path(s) and the backup path(s)

- Deactivation - TTE shifts traffic away from its backup path(s) back to the primary path(s)

Mitigation further downstream congestion

- Any change to the traffic flow may have an impact in multiple places on the network
  - When TTE is activated, it may shift traffic to an entirely different path, not just around a single link, and the change may result in congestion along the new path
- Networks that are engineered to support protection against link failures should already take this into account

Prefix selection

- When a link is outside of its bandwidth thresholds, TTE must select certain paths to activate or deactivate

- Which paths and flows to select is a critical decision that affects how quickly TTE converges to a solution where the link bandwidth is within its thresholds
  - Random
  - No Elephants
  - Maximum fit
  - Best fit
  - Maximum fit with elephants

Thank you