Delay and Loss Traffic Engineering Framework for MPLS

draft-fuxh-mpls-delay-loss-te-framework-06

November 8, 2012

IETF 85, Atlanta

Co-Authors

Xihua Fu	**ZTE (Editor)**
Dave McDysan	**Verizon (Editor, Presenter)**
Vishwas Manral	**HP**
Andrew Malis	**Verizon**
Spencer Giacalone	**Thomson Reuters**
Malcolm Betts	**ZTE**
Qilei Wang	**ZTE**
John Drake	**Juniper Networks**

Changes in -06 from Version 05

- Significant rewrite based upon review comments
- Moved problems statements, requirements to separate draft as suggested by MPLS Review Team
- Mapped potential approaches to requirements stated in problem statement draft
  - Cited other Internet drafts where appropriate
  - Described potential protocol extensions needed for some approaches
- Strived to remove specific solution descriptions and replace with more general descriptions of architecture and functionality needed
- Added section to describe how scalability and stability concerns can be addressed
- Changed Intended Status from Standards Track to Informational

Approaches to Communicate Performance Information

- Extend IGP with performance (delay, loss, delay variation) information
  - Existing drafts in OSPF and IS-IS WGs
- Augment RSVP-TE signaling with performance information
  - Request and concatenated value feedback
  - Existing MPLS wg draft
- Augment PCE(s) with performance information
  - Requestor can consult, or PCEs can communicate amongst themselves to make better path selection decisions
  - Method to inject performance information and associate with links, nodes, area/level, domain TBD
- (G)MPLS methods communicate significant changes between layers
  - Could be IGP, RSVP-TE or something else

Performance Information Estimation, Hysteresis and Automatic Responses

- Ideally, measure but a good estimate may be all that is needed since measure is statisical and over a relatively long time interval
- Use thresholds and timers to create hysteresis in IGP, PCE information base to dampen changes
- Provide capabilities (interfaces) for other processes to:
  - Automatically attempt reroute if end-end performance measurement is unacceptable
  - Automatically attempt reroute based upon performance estimate threshold crossing

Addressing Scaling Challenges

- Performance estimate changes limited to order of minutes by definition
- Augmented IGP flooding performance parameter change frequency within one area/level controlled by configuration parameters
- Augmented PCE information base performance parameter change frequency within one area/level controlled by configuration parameters
- Re-computation and re-signaling of LSPs whose composition of performance parameter values changes to unacceptable controlled by configuration parameters
- Declaration of links, nodes, FA-LSPs as unacceptable/acceptable controlled by configuration parameters
- Frequency of a lower layer network indicating a significant performance change controlled by configuration parameters
- Re-computation and re-signaling of LSPs whose measured end-end performance is unacceptable controlled by configuration parameters

Next Steps

- Pre-requisite should be wg adoption of companion problem statement/ requirements draft
- Does this version remove concerns about being too solution specific?
- What is the best way to decide between approaches that can solve the problem in different ways?
  - Augmented IGP/RSVP-TE versus augmented PCE?