Operational Guidance for Deployment of L4S in the Internet

draft-white-tsvwg-l4sops-02

Greg White, Editor

TSVWG @ IETF110

March 10, 2021

Contributors include: Bob Briscoe, Jake Holland, Koen De Schepper, Olivier Tilmans, Tom Henderson, Asad Ahmed

Scope

- Specific to the issue of L4S/Classic coexistence in RFC3168 bottleneck links (single queue bottlenecks & VPNs in fq)
  - Provide guidance to operators of L4S hosts
  - Provide guidance to operators of networks
  - Provide recommendations to researchers

- Not just in-band detection and fallback
  - Experiments to identify RFC3168 instances
  - Pre-deployment measurements
  - In-band monitoring + administrative fallback
  - Mitigations for network operators

- Note: General requirements, definition of the L4S experiments, and other operational guidance for L4S are in the L4S drafts
  - ECN-L4S-ID draft requires hosts to be capable of disabling L4S functionality by application and/or administrative control

Status

- Individual Draft
  - Draft-00 (July 2020) discussed @IETF108
  - Draft-01 (Nov 2020) discussed @IETF109
  - Draft-02 (Feb 2021)
    - Addressed comments:
      - Discuss VPNs + fq_codel
      - Include discussion on fairness
    - Improvements in Introduction
    - Section dedicated to detection mechanisms (currently a pointer)
    - Discussion of general-purpose servers vs. specialized servers

Summary of Draft-02

- Discusses prevalence & severity
  - L4S coexists well with classic traffic aside from shared RFC3168 queues
    - Single queue RFC3168 bottlenecks – believed to be rare
    - VPN traffic in fq_* bottlenecks – less rare
  - Mainly an issue for long-running flows at high data rates with long RTTs
  - Can result in L4S flows outcompeting classic flows
  - Summarizes historical issues around per-flow fairness
- Discusses detection of RFC3168 bottlenecks via experiments
  - Detecting FIFO vs. FQ is of interest
  - Points to methods for detecting RFC3168 in “Fallback” report[1]
  - Does not recommend L4S network nodes uniformly disable RFC3168 ECN (Alex Burr’s idea)
    - See Bob Briscoe’s slides: https://bobbriscoe.net/presents/2103ietf/l4s-exclusive-ecn-marking.pdf
- Provides guidance for operators of L4S hosts
  - Prior to deployment, conduct experiments on presence of RFC3168 bottlenecks
  - Take action if warranted, guidance depends on context:
    - General purpose servers (e.g. web) vs. Specialized servers (e.g. cloud gaming)
    - Edge servers vs. other hosts
- Provides guidance for operators of RFC3168 bottlenecks
  - Several options outlined to eliminate any coexistence issues

1. Briscoe, B. and A.S. Ahmed, "TCP Prague Fall-back on Detection of a Classic ECN AQM", ArXiv , February 2021, https://arxiv.org/abs/1911.00710

Mailing List Comments (since draft-02)

- Include more references to data about the deployment of RFC3168
  - And interpretations, e.g. fq vs fifo
- More info on ways to cache/maintain a list of detected 3168 paths
- Discussion of risk of incorrectly classifying a path
- “Disable RFC3168 ECN Marking” section needs to be written more clearly
- Cite RFC7567?
- Mention that RFC3168 FIFOs aren’t prohibited, so could be deployed in the future
  - Non-ECN FIFO AQMs *do* exist, some of these could turn on ECN support
- Discuss Risks
  - Risk = Severity * Likelihood
  - Who suffers? (active participant vs. innocent bystander)

WG Adoption?

- Adoption call announced March 9
- Concludes March 24