Operational Guidance for Deployment of L4S in the Internet

draft-white-tsvwg-l4sops-01

Greg White, Editor

TSVWG @ IETF109

November 18, 2020

Scope

Jake/Pete:
Don’t completely ignore FQ

- Specific to the issue of single-queue RFC3168 bottleneck links
  - Provide guidance to operators of L4S hosts
  - Provide guidance to operators of networks
  - Provide recommendations to researchers

- Note: General requirements, definition of the L4S experiments, and other operational guidance for L4S are in the L4S drafts

Coexistence of L4S & Classic Traffic

Network Bottleneck		Prevalence	Fairness Result	
AQM	Sched.			
L4S ECN	DualQ	*Future*	DualQ-Coupled-AQM provides reasonable fairness across a range of conditions	✓
L4S ECN	FQ	*Future*	FQ enforces fairness	✓
Classic ECN	FQ	Common	FQ enforces fairness*	✓
Classic ECN	FIFO	?	**Problem: Classic flows get less throughput than L4S flows**	**X**
Non-ECN	any	Common	L4S flows become Reno-friendly in response to loss	✓
Taildrop	any	Very common	L4S flows become Reno-friendly in response to loss	✓

* Except mixed traffic in tunnels, or in case of hash collision

Classic flow/ECN = RFC3168

Background (issue 16)

https://trac.ietf.org/trac/tsvwg/ticket/16

- RFC3168 defined Classic ECN behavior:
  - In an AQM, apply CE marks to ECT0 & ECT1 packets instead of drop
- L4S (ECT1) vs. Classic (ECT0) response to CE
  - L4S flows aim for a higher CE marking rate than Classic flows
- When sharing a single-queue RFC3168 bottleneck, L4S flows will out-perform Classic flows
  - Long-running capacity-seeking flows
  - Worse in moderate BDP connections (e.g. BDP≈100-150pkts)
    - Less impact at low BDPs
    - Less impact at high BDPs (Cubic operating in Cubic mode?)
  - Unfairness > 10:1 can be shown in some cases
    - e.g. Figure 1 in https://bobbriscoe.net/projects/latency/ecn-fallback_tr.pdf

3. Operator of an L4S host

- 3.1 CDN/ISP Servers (constrained set of networks/clients)
  - Prior to deployment
    - Consult with network operators on deployment of single queue RFC3168 bottlenecks
    - Consult with network operators on deployment of L4S
    - Perform testing to assess presence/absence of RFC3168 bottlenecks
  - In-band detection/monitoring via mandatory features in [l4s-id]
    - Real-time response (fallback)
    - Non-real-time response (disable for future connections)
- 3.2 Other hosts
  - In-band methods as above
  - Per-dst path testing

TODO:
Define specific tests

TODO:
Define specific tests

4. Operator of a single queue RFC3168 bottleneck

- 4.1 Configure AQM to treat ECT1 as NotECT
- 4.2 Configure Non-Coupled Dual Queue
- 4.3 WRED with ECT1 Differentiation
- 4.4 ECT1 tunnel bypass
- 4.5 Disable RFC3168 ECN marking
- 4.6 Re-mark ECT1 to NotECT prior to AQM (e.g. ingress)

Sebastian:
Are there routers that can do this?

Need input from network gear vendors & operators

5. Role for Researchers

- 5.1 Measurement campaigns to detect RFC3168 bottlenecks
  - Active testing
  - in-band/out-of-band
  - distinguish FIFO from FQ
- 5.2 Measurement of L4S vs Classic performance

Rüdiger:
Can we define specific test requirements?

- Also, how often do multiple long-running flows coincide within one bottleneck queue?

Need input from researchers

Additional recent mailing list comments

- Sebastian Möller
  - Reword “more precise flow balance”
- Jonathan Morton
  - Scrap L4S and choose SCE+DSCP instead
- Jake Holland
  - Need references to tests showing the unfairness issue
  - Should this be “Informational” or “Experimental”?

Next Steps

- WG adoption?
- Commitment to provide text?