

L4S Status Update

Bob Briscoe, Independent

<ietf@bobbriscoe.net>



Koen De Schepper, **NOKIA** Bell Labs <koen.de_schepper@nokia.com>



Olivier Tilmans, **NOKIA** Bell Labs <olivier.tilmans@nokia-bell-labs.com>

Asad Sajjad Ahmed, Independent <me@asadsa.com>

Joakim Misund, Uni Oslo <joakim.misund@gmail.com>



tsvwg, IETF-110, Mar 2021

Focus

- Prague CC implementation
 - Faster getting up to speed (GUTS)
- Parallel activities
 - Prague requirements survey + text updates
 - Ops guidance on Classic ECN AQM issue

CC: Congestion Control
ECN: Explicit Congestion Notification
AQM: Active Queue Management

New Prague CC Draft

- New full write-up
 - draft-briscoe-icrg-prague-congestion-control-00
 - only published yesterday
 - answers FAQs
- Transport-protocol- and OS-agnostic
 - plus copious transport- & implementation-specific notes
 - majority: TCP Prague over Linux
- There is not 'One True Prague CC'
 - L4S Prague Requirements [draft-ietf-tsvwg-ecn-l4s-id] are for that
 - This draft describes a reference design with implementation notes

Prague CC: Reality + Promise

Linux code:	code	none (disabled)	research private	research opened	L4Steam git	Linux RFC	Linux mainline
Requirements				base TCP		TCP Prague	
L4S-ECN Packet Identification: ECT(1)							mandatory
Accurate ECN TCP feedback				sysctl option			mandatory
Reno-friendly on loss							inherent
Reno-friendly if classic ECN bottleneck							default off -- on later
Reduce RTT dependence (low RTT dominance)							default on
Scale down to fractional window				research git			research git
Detecting loss in units of time				default RACK			default RACK
Performance improvements							
ECN-capable TCP control packets				module option off			default off -- on later
Faster flow start				research git			research git
Faster than additive increase							in progress
Continuous additive increase							default on
Reduce RTT dependence (high RTT weakness)							todo
Burst avoidance for TSO sizing & pacing (<1ms)							default
Performance-bug fixes							
integer scaling & fractional carry (alpha, cwnd, etc)							fixed
PRR undershoot spike							fixed

- Maturity Status (paraphrasing from the abstract):
 - Implementation does not satisfy all the Prague req's (yet)
 - IETF might relax certain req's as an outcome of the process of trying
 - In two cases, research code replaced by placeholders until fully evaluated
 - 1) faster flow start
 - 2) faster additive increase
- §2: Proven code, defaults
- §3: Optional code (under development / evaluation) + ideas

Prague CC – recent work

- Prague CC is not only about adding code to DCTCP for each Prague L4S req't
 - draft corrects this impression
 - subtle but critical changes to:
 - feedback averaging, AI, MD
 - burst limiting / pacing
 - segmentation offload

Linux code:	code	none (disabled)	research private	research opened	L4S team git	Linux RFC	Linux mainline
Requirements					base TCP	TCP Prague	
L4S-ECN Packet Identification: ECT(1)							mandatory
Accurate ECN TCP feedback					sysctl option		mandatory
Reno-friendly on loss							inherent
Reno-friendly if classic ECN bottleneck							default off -- on later
Reduce RTT dependence (low RTT dominance)							default on
Scale down to fractional window					research git		research git
Detecting loss in units of time					default RACK		default RACK
Performance improvements							
ECN-capable TCP control packets					module option off		default off -- on later
Faster flow start					research git		research git
Faster than additive increase							in progress
Continuous additive increase							default on
Reduce RTT dependence (high RTT weakness)							todo
Burst avoidance for TSO sizing & pacing (<1ms)							default
Performance-bug fixes							
integer scaling & fractional carry (alpha, cwnd, etc)							fixed
PRR undershoot spike							fixed

L4S Drafts – Summary

- Milestone for main drafts: Oct'21; aiming for earlier
- draft-ietf-tsvwg-l4s-arch-08
 - stable, complete, unchanged
- draft-ietf-tsvwg-ecn-l4s-id-14
 - Tom Henderson's review
 - updates to L4S Transport (Prague) Reqs (§4) & supporting appendix (thx to survey respondents, Asad). Ongoing...
- draft-ietf-tsvwg-dualq-coupled-aqm-14: stable, except...
 - imported L4S Network Reqs by reference [ecn-l4s-id; §5]
 - Ingemar Johansson's review
- draft-white-tsvwg-l4s-ops-02
 - updated – in adoption call (thx to Bob, Jake, Koen, Olivier, Tom, Asad)
- draft-briscoe-iccrp-prague-congestion-control-00
 - new

L4S Status Update

Q&A