Identifying Modified ECN Semantics for Ultra-Low Queuing Delay

draft-briscoe-tsvwg-ecn-l4s-id-01

Bob Briscoe

[simula . research laboratory]

- Koen De Schepper, Inton Tsang

NOKIA Bell Labs

IETF-95 Apr 2016

The authors were part-funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700). The views expressed here are solely those of the authors.

Low Latency Low Loss Scalable throughput (L4S) – background

Latency so low it sticks to your finger

- Recall: demo at Prague IETF (aqm wg & bits-n-bites)
  - see https://riteproject.eu/dctth/ for videos, papers, etc
- L4S could incrementally replace “best efforts”
  - ultra-low queuing delay
  - zero congestion loss
  - scalable throughput (beyond Reno, Compound, Cubic)
- Eventually for *all* Internet traffic
- Aim: to be worth the deployment hassle – so much better than today

very high level

- problem: TCP is the elephant in the room
- solution: build another room without the elephant

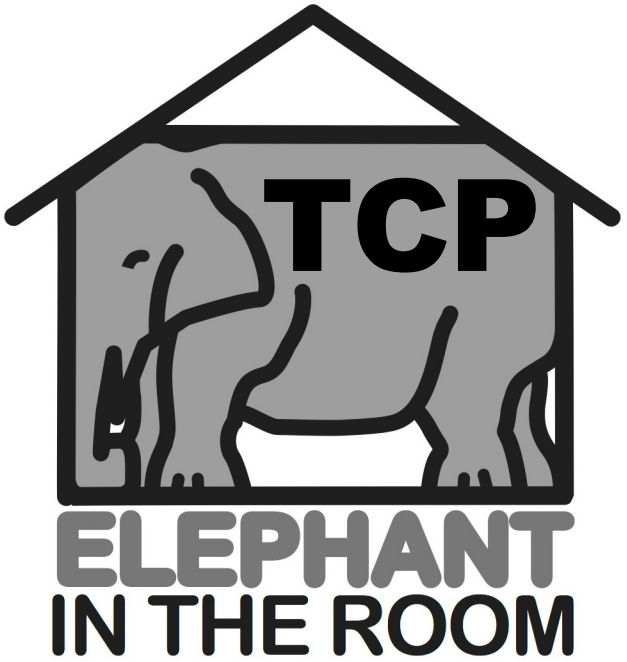

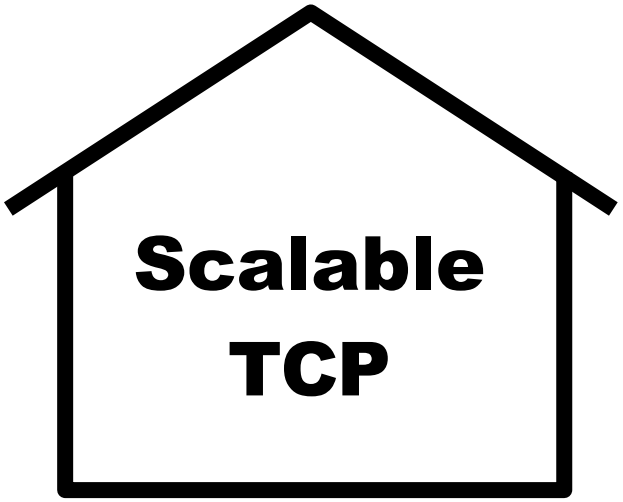

3 parts to standardise

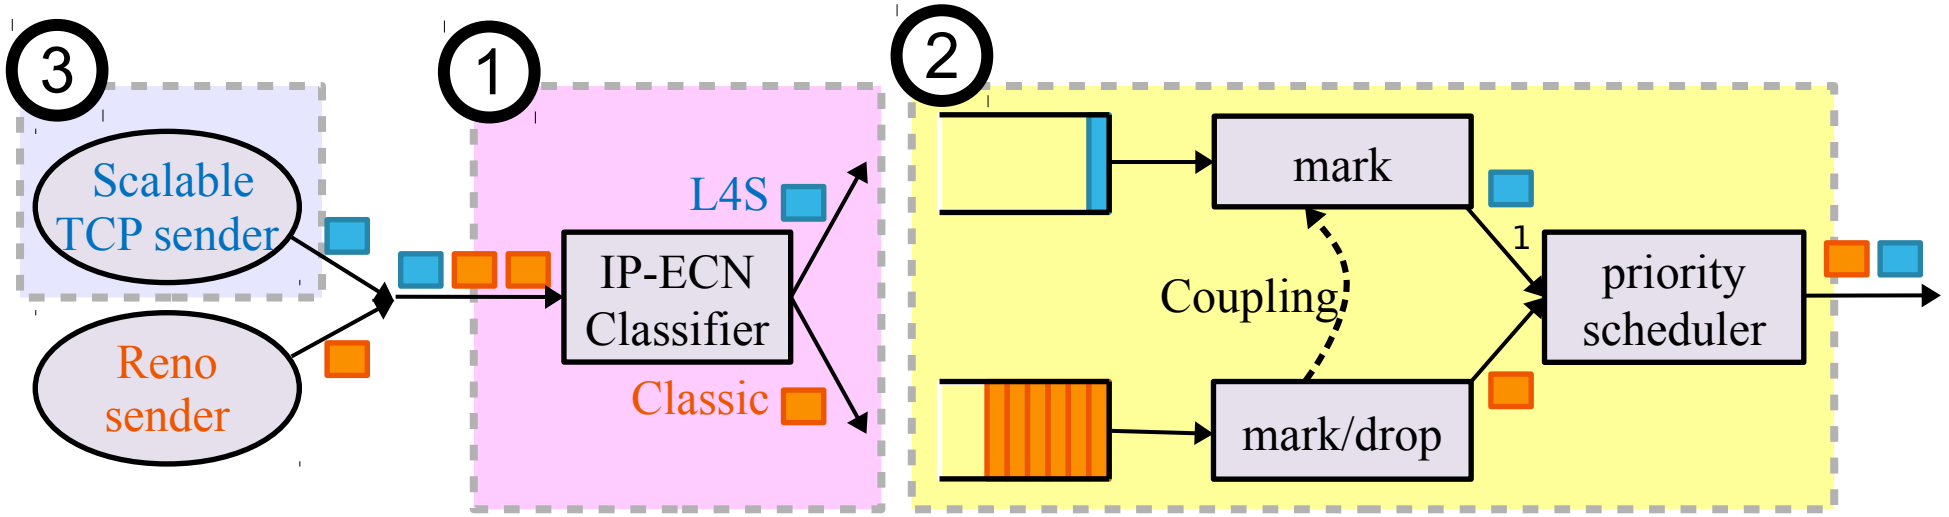

1)	The identifier	draft-briscoe-tsvwg-ecn-l4s-id	tsvwg
2)	The DualQ AQM	draft-briscoe-aqm-dualq-coupled	aqm
3a) 3b)	Scalable transport	draft-ietf-tcpm-accurate-ecn draft-ietf-tcpm-dctcp (bis)	tcpm

this talk

- #1, #2 are as general as possible
- #3a) fixes TCP feedback, other transports are already OK
- #3b) one concrete example transport behaviour: DCTCP

choice of identifier

- Three possibilities; all involve compromises
  - two other possibilities quickly discounted
  - ECT(1) + CE chosen
  - reasoning recorded in Appendix A of draft
  - table highlights solely the distinguishing issues

Issue	DSCPx + ECN		ECN*	ECT(1) + CE	
	initial	eventual	initial	initial	eventual
end-to-end?		?			
tunnels			?	?	
lower layers		?			?
codepoints		?			?
reordering					?
control packets					?

Legend	
	Poor
	Ordinary
	Good
?	Optimistic

CAVEAT: The table is not meant to be understandable without referring to the text.

---

* only feasible to use ECN alone if Classic ECN becomes obsolete
"ECN" shares the eventual scores of "ECT(1) + CE"

meaning of this new identifier?

- Original goals of ECN included lower delay with modified TCP
  - but too many combinations to standardise a winner
  - so [RFC3168] defined 'Classic' ECN behaviour as *equivalent* to drop

    ...could allow new criteria to be developed for setting the CE codepoint, and new congestion control mechanisms for end-node reaction to CE packets. However, this is a research issue, and as such is not addressed in this document.
  - so ECN inherited the dilemma of drop-based TCP...

- Proposed meaning of L4S identifier:

> The likelihood that an AQM drops a Not-ECT Classic packet MUST be roughly proportional to the square of the likelihood that it would have marked it if it had been an L4S packet.
>
> The constant of proportionality does not have to be standardised for interoperability, but a value of 1 is RECOMMENDED.

If we choose ECT(1): what do we preclude?

- would obsolete ECN nonce [RFC3540]
  - other ways to do feedback integrity without a codepoint
- various “less severe than CE” schemes from the research community, incl. for flow start-up
- see Appendix B

- consider carefully before consuming the last ECN codepoint

Intended Status

- If Proposed Standard, it would update:
  - ECN in IP [RFC3168]
  - ECN in TCP [RFC3168]
  - ECN in SCTP [RFC4960]
  - ECN in RTP [RFC6679]
  - ECN in DCCP [RFC4340]
- Draft written as Experimental
  - can obsolete ECN nonce (also experimental)
  - if experimental, cannot update all the above proposed standards
- Discuss

Next Steps

- Please review and comment
  - brief draft (8pp without boilerplate & appendices)
- Plenty of discussion (on aqm) when issue first raised
  - quiet since
- Industry L4S assessment activities in progress
  - TSVWG will need visibility of this before it can adopt

Q&A

large saw teeth can ruin the quality of your experience

Q) why is queuing delay of DCTCP so low?

A1) Finer saw-teeth of a Scalable TCP