L4S Architecture Low Latency, Low Loss, Scalable Throughput Internet Service

draft-briscoe-tsvwg-l4s-arch-00

B. Briscoe, K. De Schepper, M. Bagnulo

IETF97

Motivation

- Provide reduced latency for Internet traffic
- Achieve that by reducing the queuing latency suffered by TCP packets.
  - Current queues are inherent to the (large) saw tooth in TCP (either New Reno, Cubic or else)
  - Proved (in closed environments) that transport protocols with small saw teeth can significantly reduce the queuing delays (DCTCP)
- L4S aims to provide safe deployment of small teeth transport protocols in the Internet.

L4S Architecture overview

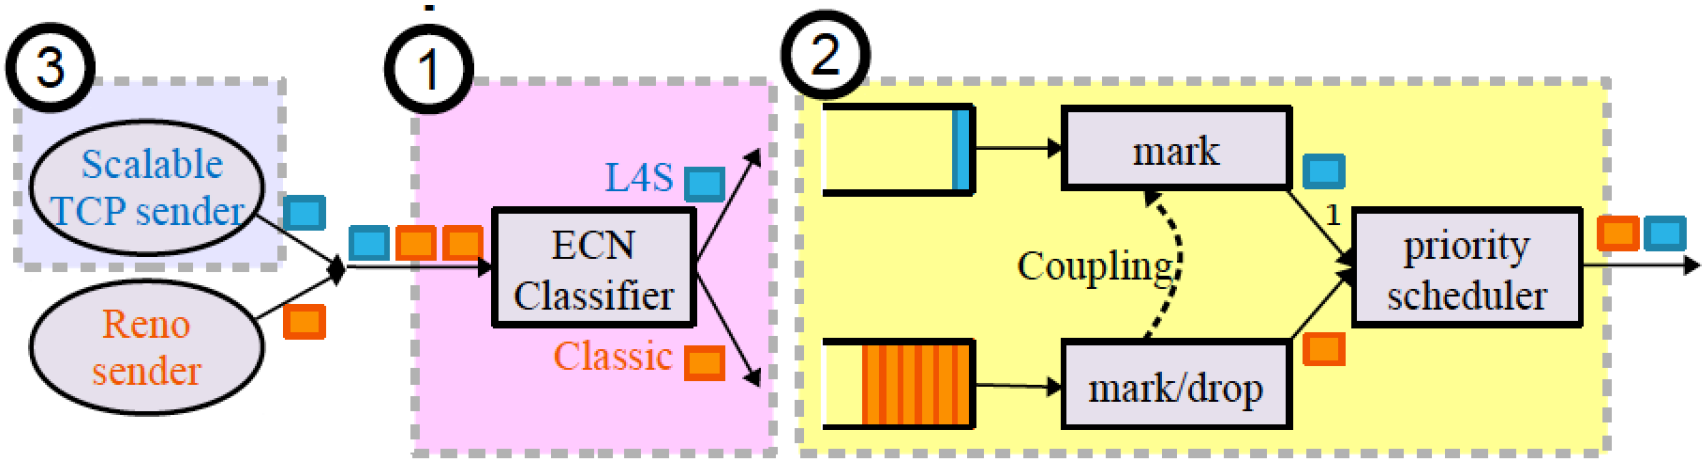

Host Protocol Network

L4S Architecture overview
Host part - #3

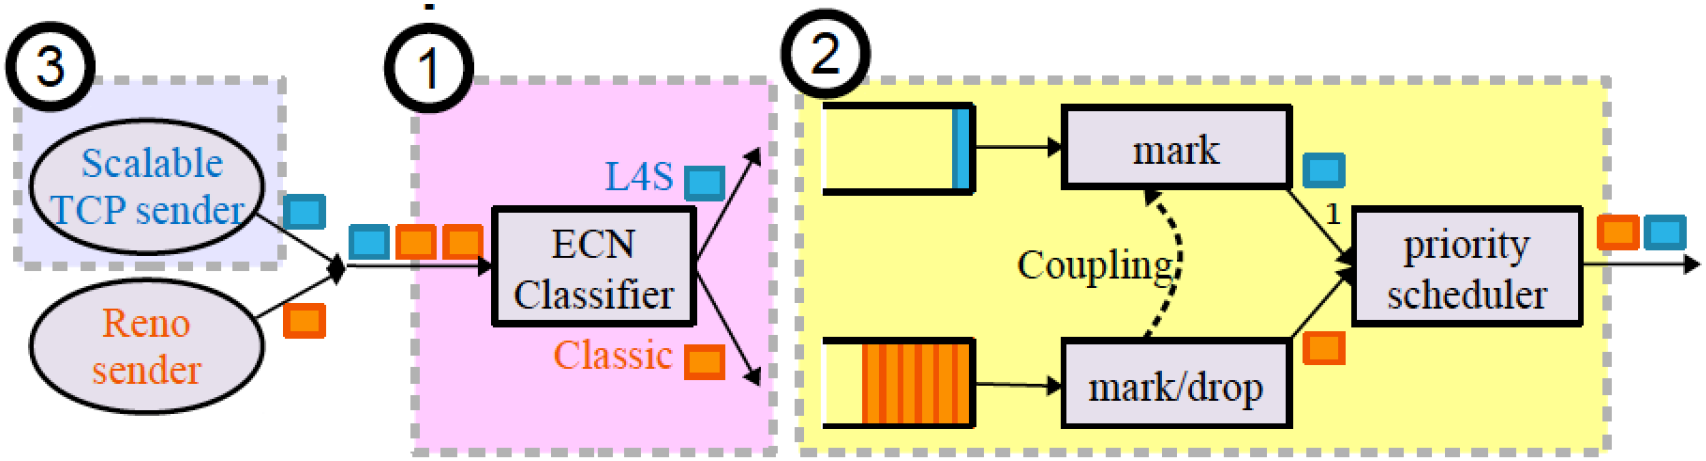

- Scalable congestion control already exists (DCTCP)
- Additional mechanisms are needed to safe deployment in the Internet
  - E.g. fall back to Reno if classic ECN bottleneck or in case of a loss, accurate ECN feedback
  - See others in draft
- Scalable congestion control for other transports

L4S Architecture overview
Network part #2

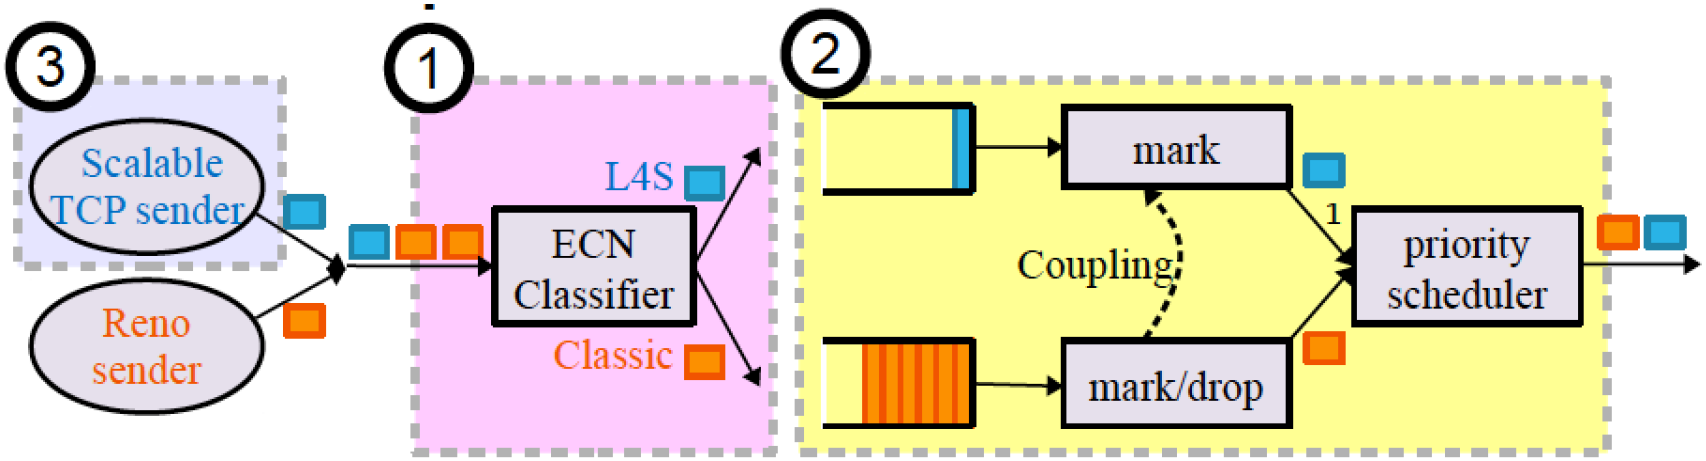

- AQM that is able to isolate the L4S traffic from regular traffic
  - If not L4s traffic would starve classic traffic and still not achieve low latency

L4S Architecture overview
Protocol part #1

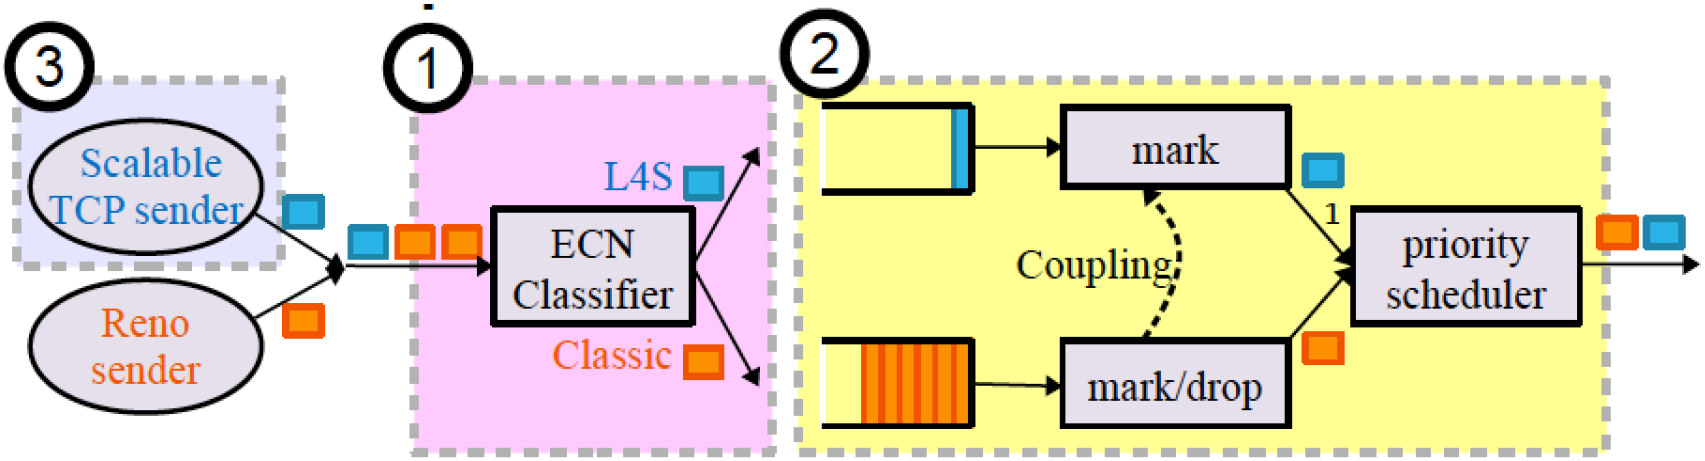

- A mean to distinguish L4S packets from regular packets, so the AQM can process them accordingly

Status update

- Source Code
  - Dual Queue Coupled AQM, With Curvy RED for Linux (access available shortly), With PI2 for Linux [UPDATED]
  - Data Centre TCP (DCTCP) fo Linux (in the mainline kernel), FreeBSD patch, ns2 patch.
  - Accurate ECN TCP Feedback for Linux [COMPLETED, but not fully tested]
- IETF specs
  - Low Latency, Low Loss, Scalable Throughput (L4S) Internet Service: Architecture <draft-briscoe-tsvwg-l4s-arch> [NEW]
  - A proposed new identifier for Low Latency, Low Loss, Scalable throughput (L4S) packets <draft-briscoe-tsvwg-ecn-l4s-id> [UPDATED to reflect the proposed process]
    - enabled by <draft-black-tsvwg-ecn-experimentation> [NEW]
  - Dual-queue AQM: : <draft-briscoe-tsvwg-aqm-dualq-coupled>
  - scalable TCP algorithms, e.g. Data Centre TCP (DCTCP) <draft-ietf-tcpm-dctcp>
  - Accurate ECN: <draft-ietf-tcpm-accurate-ecn> [UPDATED]
  - Adding ECN to TCP control packets: <draft-bagnulo-tcpm-generalized-ecn> [UPDATED]
- Papers
  - Article in the IETF Journal describing the Demo in Bits-N-Bites at the IETF in Prague, July 2015. “Ultra-Low Delay for All” IETF Journal, Nov 2015.
  - “Ultra-Low Delay for All: Live Experience, Live Analysis“, Proc. ACM Multimedia Systems; Demo Session (May 2016).
  - “PI2: A Linearized AQM for both Classic and Scalable TCP,” Proc. ACM CONEXT 2016 (To appear Dec 2016). [NEW]

Next steps

- Please review and comment.
- From the authors perspective, the document is in good enough shape for have a call for adoption.