

L4S Issues 26, 27, 22

TSVWG Virtual Interim

Feb 20, 2020

Greg White

g.white@cablelabs.com

Issue 26: Admission Control

- Issue summary:
 - In Dual-Q, high rate unresponsive traffic in the L4S queue could cause problems
 - NQB draft recommends Queue Protection to prevent problems with “misbehaving” NQB-marked flows
 - Shouldn’t DualQ-coupled-AQM also recommend QP to prevent problems with misbehaving ECT(1)-marked flows?
- Non responding traffic of any type can hurt any type of other traffic due to the coupling
- Overload protection makes sure that this is also the case in extreme non-responding cases that saturate the coupling
- L4S does not create a DoS vector
- Proposal:
 - Many L4S nodes may also support the NQB PHB, where QP is recommended already
 - For those that don’t, additional policers can be added to protect the L4S queuing latency from misuse on an as-needed basis.

Issue 27: The term “Classic”



- Lots of opinions expressed - no consensus
- Proposed Compromise
 - Existing TCP
 - Don't refer to existing standard TCP as “Classic”
 - Instead use “non-L4S TCP” or “TCP using classic congestion control”
 - Congestion Control
 - Continue to use the term “Classic Congestion Control” – defined in the document to mean Reno-friendly or $1/\sqrt{p}$
 - RFC-3168 ECN
 - Continue to use the term “Classic ECN” – defined in the document to mean RFC-3168 ECN
 - Dual-Queue Description
 - Continue to use the terms “Classic Service”, “Classic Queue”, “Classic Packets”, “Classic sender”, “Classic AQM”, “Classic Traffic” to differentiate from “L4S *”

Issue 22: Deployment Feasibility

- Issue Summary:
 - DualQ classifies CE-marked packets to the L4S queue
 - A non-RACK-capable, RFC-3168 ECN-Capable TCP Connection, operating on a path where an RFC-3168 ECN AQM bottleneck precedes a DualQ bottleneck, could experience packet reordering
 - Likely one or more CE packets arriving ahead of some lower-sequenced ECT(0) packets
- Question:
 - In the case of a single CE marked packet, wouldn't the sender see one dupack with ECE flag, then the normal ack sequence?
 - In the case of multiple (n) consecutive CE marked packets, wouldn't the sender see n dupacks with ECE flag, then the normal ack sequence?
 - In both cases, wouldn't the sender simply cut its cwnd one time? Since that is the desired behavior anyway, is there a problem to solve?
- Proposal:
 - This seems like an unlikely situation to happen in practice
 - This topic is already discussed in [Appendix B.1](#) of the ecn-l4s-id draft
 - This seems like a relatively low priority item, but it could be tested if need be.