

Some Congestion Experienced

Some Congestion Experienced (SCE)

A RFC3168 backward compatible approach to high fidelity ECN signaling and congestion control

Pete Heist
Network Analytics
Rodney W Grimes
FreeBSD

Jonathan Morton
Principal Architect: sch_cake and Cobalt AQM systems
TekLibre, LLC
IETF tsvwg, March 25, 2019



The problem

- Transports spend most of their time either:
 - Building a bottleneck queue → high delay
 - Under-utilising the path capacity → low goodput
- The network doesn't give them enough information!
 - Binary signal per RTT → oscillation
 - Bang-bang control theory

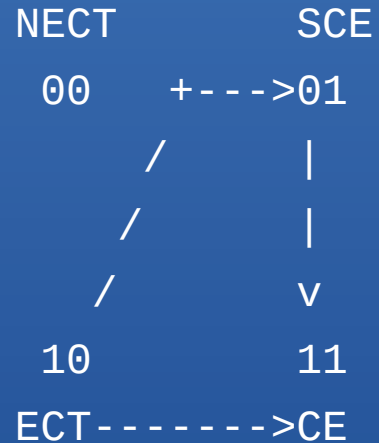
The solution

- Congestion should be an earlier signal than loss
- ECT(1) is unused, treated identically to ECT(0)
 - RFC 3168, actual measurements cited in RFC80XX
- Rename ECT(1) as SCE, ECT(0) as ECT
 - Some Congestion Experienced
- Other ECN codepoints retain current meaning & semantics

Two Congestion Signals

- SCE (some congestion experienced)
 - High fidelity congestion information
 - 100% marking means CE threshold reached
 - Less marking means less congestion
- CE (congestion experienced)
 - Coarse congestion information
 - Continue using as defined now
 - May be encountered without SCE (existing middleboxes)

State Diagram



NECT: Not ECN Capable Transport

ECT: ECN Capable Transport

SCE: Some Congestion Experienced

CE: Congestion Experienced

Experimental Plan

- Experiment #1 – First, do no harm – success!
- Experiment #2 – Test marking behavior – success!
- Experiment #3 – middlebox (OpenWrt) router tests - ongoing
- Experiment #4 – integrate with TCP
 - Requires a feedback path receiver → sender
- More details on our web site at:
 - <http://www.bufferbloat.net/wiki/projects/ecn-sane>

Proportion

192,168,48,197;44974-192,168,48,86:5201

