

Advice on network buffering

draft-fairhurst-tsvwg-buffers-00

Gorry Fairhurst, Uni of Aberdeen
Bob Briscoe, BT

IETF-86 Mar 2013



Existing buffering advice: § 13 RFC3819 Advice for Internet Subnetwork Designers (BCP)

- Started 1999, published Jul 2004.
- Tentatively recommended RED & ECN
- Recommended large buffers
 - $\text{link_bandwidth} * \text{link_delay product} * N$
- At that time
 - L2 equipment had very short buffers
 - research on sizing buffers was immature (just)
[McKeown Sizing Router Buffers, SIGCOMM' 04]
- We want to fix the advice

Proposed flow of logic

- A long-running TCP will fill a tail-drop buffer if it is the bottleneck
 - hard to test, because intermittent
 - conditional on coincidence of 4 pathologies
- Therefore *implementers* should use AQM & ECN
 - in every buffer: subnet, router, middlebox or host
 - later section lists candidate AQMs
- If line rate adjusts, buffer should adjust accordingly
- If no AQM in existing buffers
 - advice for *operator* on buffer sizing
- If no auto-adjustment in existing buffers
 - advice for *operator* on static buffer sizing

enlisting help of ICCRG

- draft is currently a fairly empty vessel
 - individual -00 version
 - intended for IETF tsvwg
 - intended status: best current practice (BCP)
- specific sections on buffer sizing for
 - host, router/switch/middlebox (edge & core)
 - flow isolation
- need consensus on content for these sections