

Network Performance Isolation using Congestion Policing

[draft-briscoe-conex-policing-00.txt](#)

Bob Briscoe, BT

IETF-87 ConEx Jul 2013

draft status

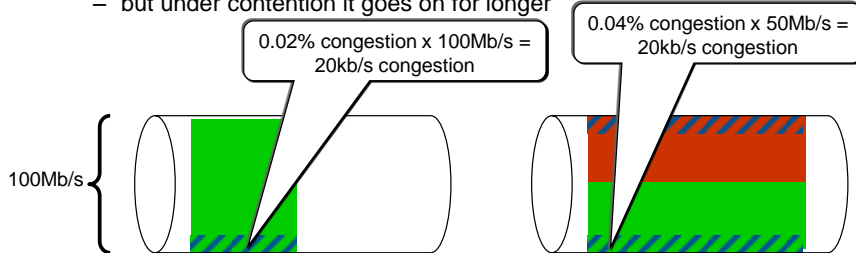
- Network Performance Isolation using Congestion Policing
[draft-briscoe-conex-policing-00.txt](#)
- new individual draft Feb '13
 - mostly cut & pasted from conex-data-centre, but new info too
 - not specific to data centres
- largely complete (33pp)
- immediate intent: encourage review & consider for WG adoption
- audience: ConEx & non-ConEx – assumed sceptical
 - how it works is simple
 - why it works is outside people's comfort zones
 - how can you isolate users with no per-user config on switches?

what's in it, and what's new

- Frontpieces (Abstract, Intro)
- 2. Example Bulk Congestion Policer **[new]**
- 3. Performance Isolation: Intuition **[corrected]**

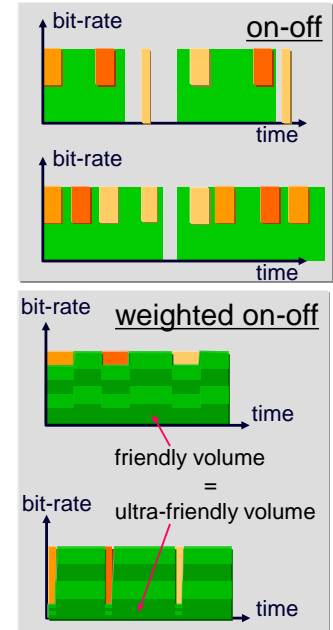
performance isolation intuition

- congestion policer enforces bit-rate, $x = w/p$, where
 - w is a constant for that tenant (the policy)
 - p is % congestion
- similar to a so-called 'scalable congestion control'
 - but for aggregate (hose) from source, made up of flows
 - TCP is evolving towards this (Compound, Cubic, DCTCP etc)
- property easy-to-say but hard to grasp:
 - same rate of congestion per tenant, however many other tenants capacity is shared with
 - congestion-bit-rate in one flow is the same
 - but under contention it goes on for longer



intuition built-up as follows

- 'scalable congestion control' as boundary case (previous slide)
- single link
 - long running flows, single link
 - similar to (weighted) round-robin
 - on-off flows
 - congestion-volume accounts for how often a tenant is *not* 'on'
 - weighted on-off flows
 - longer flows shift away
- network of links
 - congestion-volume allows for how many links tenant is 'on' in
- transients



4. Parameter Setting **[new]**

- Tailpieces (Security, Conclusions, Acks)

Network Performance Isolation using Congestion Policing

[draft-briscoe-conex-policing-00.txt](#)

Q&A