

QoS Appliances Cause Cancer And Contribute to Global Warming

Stanislav Shalunov <shalunov@internet2.edu>

Spring Internet2 Member Meeting, Washington, 2002-05-07

Middleboxes Designed for T1 speeds

- QoS appliances are designed to handle a really scarce resource
- The target market used to be organizations with a congested T1 (and it still shows)
- High-end routers today do 10 Gb/s, high-end QoS appliances do 622 Mb/s
- Deployed routers today do 2.5 Gb/s, deployed QoS appliances do 100 Mb/s
- There's more than an order of magnitude difference in speeds
- You want to use equipment that has extra capacity and leeway

This stuff changes your packets!

- Often wants to rewrite TCP window size
- Will it rewrite TCP options?
- These packet changes will lead to very-hard-to-track-down end-to-end performance problems.

Complexity that Hurts Performance and Transparency

- Configurations are complex; is some behavior the intended result of QoS appliance configuration?
- Configurations can change with time of day (“OK to share files at night”); this baffles someone who tries to debug a performance problem
- Configurations are frequently updated; the behavior of the network changes every week (or month, or day, ...)
- Making the network more opaque now leads to trouble later

Unintended Drops

- TCP throughput $\approx 0.7 \frac{MSS}{RTT \sqrt{\text{loss}}}$ (with adequate window)
- Let $MSS = 1460$ bytes, $RTT = 70$ ms.
Then $\text{loss} = 1/100,000 \Rightarrow$ TCP throughput ≈ 37 Mb/s.
- Along with managing Internet1 capacity usage, you can kill your Internet2 performance.
- “I am connected at 100 Mb/s. Why do I only get 37 Mb/s?”
- How do you measure 0.001% loss (so much for *five nines*)? ping won't do it.
- Routers are tested for this sort of things. QoS appliances aren't designed to do that!

“But it produces such nice utilization stats”

- Yes, it can gather very interesting statistics (and you should know what your network is used for)
- Should this statistics be gathered in a device that also does packet forwarding or in a passive monitor? (Or perhaps even NetFlow analysis would be sufficient for your needs?)
- If a passive monitor is short on CPU cycles the only bad thing that can happen is a little inaccuracy in the stats.
- If a packet forwarding device is overwhelmed, packet forwarding suffers along with statistics gathering.

Usage-based billing: the right way of doing same

- The real alternative to QoS appliances: usage-based billing
- Bill top 1% of your Internet1 users per bit and the problem goes away
- Or, make top 1% clean the toilets (it's not the money, it's the natural and required negative feedback loop)
- Don't play cat-and-mouse games: on the net, the mouse always wins (by using, e.g., http, https, proxies, random port numbers, ssh, etc.)
- If you play a cat-and-mouse game, after the mouse wins you not only lose your ability to control what's going on. You can't even tell what's going on in your own network!

What we should be doing instead

- Usage-based billing (to take care of habitual overuse by people who just don't care because they're not the ones paying).
- Easy to deploy passive monitors (to have understanding of network dynamics, usage, and performance).
- NetFlow analysis (do less than passive monitors, but without new boxes that cost money).
- Scavenger service (to let people who care police themselves).