

1



# QoS: What Went Wrong?

Grant Goodale

23 September 2004

2



# RIPQoS 2003

- Revisiting IP QoS? or *Requiescat in Pace?*
- Bruce Davie: *Deployment Experience with Differentiated Services*
- Crowcroft et. al.: *QoS's Downfall: at the bottom, or not at all!*
- Teitelbaum & Shalunov: *What QoS Research Hasn't Taught Us About Risk*



3



# What is QoS, anyway?

- Depends on who you ask.
  - Set of mechanisms for traffic classification?
  - End-user service?
  - Risk mitigation technique?
- Who are the stakeholders in QoS?
  - ISPs?
  - End users (Enterprises)?
  - Consumers?

4



# What QoS isn't

- It's not a way to create bandwidth
  - QoS is a zero-sum game
- It's not an end-user service
  - A platform for services (if you're an ISP)
  - An engineering tool (if you're an enterprise)

## A brief\* history of IP QoS

- ☘ Integrated Services (IntServ)
  - ☘ RSVP used to analyze and provision paths (signaled approach)
  - ☘ Bandwidth is reserved per flow
- ☘ Complex RSVP implementation
- ☘ Per-flow state information in routers (E2E again!)
  - ☘ Scalability issues
  - ☘ Policy enforcement issues

\* In this case, brief means "paraphrased from Jim Gogan's (gogan@unc.edu) presentation "IP Quality of Service: IntServ and DiffServ"

## A brief\* history of IP QoS

- ☘ Differentiated Services (DiffServ)
  - ☘ Traffic is classified at the edge (IPv4 TOS field / IPv6 class field)
  - ☘ Routers allocate resources per class of traffic
  - ☘ traffic handling can change on a per-hop basis
- ☘ No per-flow state
- ☘ Policy applied to aggregate flows
- ☘ More difficult policy interface between ISPs

\* Here, brief means "stolen from Cisco"

## Has QoS failed?

- ☘ Widely deployed support for IntServ and/or DiffServ
  - ☘ Cisco IOS
  - ☘ Juniper
  - ☘ Nortel
- ☘ Deployment is just the first step:
  - ☘ Configuration
  - ☘ management, packaging for resale, &c.

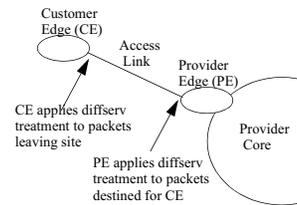
## Has QoS failed? (cont.)

- ☘ Significant adoption in certain markets:
  - ☘ L3 VPNs, enterprise networks
- ☘ Typically single-provider, single administrative domain
- ☘ High-priority traffic usually between nodes within the administrative domain
- ☘ VOIP (maybe)
  - ☘ Not clear if Telcos with IP backbones are actually using QoS

## Has QoS failed? (cont.)

### ISP edge deployment

- Deployed on access link between customer and ISP
- Difficult for ISP to categorize customer-bound traffic
- Works well as long as only customer's outgoing traffic needs classification



## Has QoS failed? (cont.)

### What about QoS at Layer 2?

- MPLS, 802.1p, AAL1-5, FR Traffic Shaping, &c.
- Often sophisticated, fine-grained
- Insufficient by itself - requires some Layer 3 support to be useful
  - Different transport layers = different mechanisms
  - Currently not popular outside of single-domain use

## Why don't ISPs offer QoS?

- Expensive
  - Possible infrastructure upgrades
  - Implementation costs: training, configuration, management
  - Consumer-end equipment
- Until recently, few clear services built on top of QoS
  - VOIP?
- No economic incentive for cooperation between ISPs
  - QoS seen as a differentiator

## Why don't ISPs offer QoS?

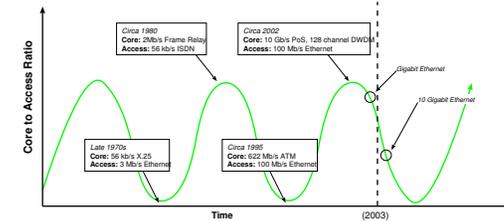
- What about security?
- Theft of Resources
  - Theft of bandwidth
  - Illegal promotion of packet PHB
- Denial of Service
  - Traffic class overloading (core)
  - source penalization (edge)



## The real reason: Fat Pipes

- ❁ No point implementing QoS in grossly over-provisioned networks
- ❁ Anecdotal evidence: 15% utilization = time to widen the pipe!
- ❁ A little-utilized best-effort fat pipe works well enough
- ❁ Core bandwidth far outstrips access bandwidth
  - ❁ DWDM at the core, 100Mb/s at the edge
  - ❁ Wasn't always this way, though...

## A Vicious Cycle



- ❁ Warning: “bogograph”
- ❁ Core bandwidth is ahead now, but it wasn't always that way
- ❁ Why didn't QoS succeed during one of the previous troughs?

## QoS R&D

- ❁ QoS interest was (is?) largely reactive
  - ❁ If there's no problem to solve, implementers focus elsewhere
  - ❁ Technical solutions have usually been too late
  - ❁ Providers historically have solved the problem by throwing bandwidth at it
- ❁ No shortage of research in the area: ~8000 papers in 10 years
- ❁ A few recent examples:

## QoS in overlay networks

- ❁ A lot of recent research attention
  - ❁ QRON
  - ❁ OverQoS
- ❁ Promising results with real-world applications on the WAN
- ❁ Crowcroft et. al. : No way!
  - ❁ Not TCP-friendly
  - ❁ Adds YACS (Yet another control system)

## One-bit QoS

- Proposed by Crowcroft et. al.
- Concept: Build very simple two-level QoS into the low-level internet
  - Routers signal congestion to end systems
  - End systems use higher-level algorithms to implement 'fair share' allocation
- Open question: which is more expensive, implementing the already extant DiffServ or deploying and implementing 'A Bit of QoS'?

## Is QoS the right solution?

- Many axes of 'quality': latency, bandwidth, jitter, &c.
- QoS schemes that only address one axis may be unsuitable for some applications
- What is the customer's desired result?
  - Hint: it's not 0 packet loss
  - An end to expensive over-provisioning
  - Protection from / mitigation of risk

## Is QoS the right solution?

- Over-provisioning: "Throw bandwidth at the problem"
  - Expensive
  - Inefficient
- Doesn't necessarily work!
  - Doesn't prevent high-priority traffic from being dropped due to bursts of lower-priority traffic
  - The 'backhoe' problem (Of course, QoS doesn't help here either)
  - How do you deal with failures?

## Warranted QoS

- a.k.a. Layer 8 (Financial)
- Recognizes that customers don't want 100% uptime
- Requires careful statistical analysis of traffic (ongoing effort)
- In the end, unlikely to succeed
  - Sets unrealistic customer expectations ("110% uptime!")
  - DoS attacks can be engineered to exceed any assumptions based on statistical modeling

## Risk management

- ☘ “Bit insurance” - also Layer 8
  - ☘ Risk assumed by a neutral third party
  - ☘ Financial compensation for network failures is usually acceptable
- ☘ Requires buy-in from ISPs to be effective
- ☘ If ISP QoS revenue is significant, could be hard to get buy-in
- ☘ Unlike QoS, insurance can be end to end without cooperation between individual ISPs

## Application-layer solutions

- ☘ FEC, application-level buffering, TCP
- ☘ Impose some additional overhead
  - ☘ May or may not be reasonable, depending on application
- ☘ Same potential problem as QoS overlays - all users incented to use apps that are loss tolerant, congestion results
- ☘ Loss-tolerance techniques aren't infallible; probably unsuitable for mission-critical apps

## The future of QoS?

- ☘ Demand for QoS may resume in the future
  - ☘ Access bandwidth may close the gap with core bandwidth (No more bandwidth glut)
  - ☘ Possibility of increased interest in failure-case behavior
  - ☘ Emergence of the QoS “killer app” (VOIP? Online gaming?)
- ☘ ISPs may grow more comfortable with QoS through experience in other environments
- ☘ Standardization of QoS services between providers may speed adoption

## ...or not.

- ☘ Core bandwidth has some headroom, still ramping at impressive rates
  - ☘ Plenty of dark fiber at this point
  - ☘ Lots of research
- ☘ Deployment issues remain thorny, unsolved
  - ☘ Need a standard way of arranging QoS between ISPs
  - ☘ Some research, but gets into Layers 8 and 9
- ☘ If throwing bandwidth at the problem will work, QoS will be left out.

## Other RIPQoS papers

- ♣ Gregory Bell, LLNL: *Failure to Thrive: QoS and the Culture of Operational Networking*
- ♣ IP QoS has failed to thrive in enterprise networks due to disconnect between researchers and operations people
- ♣ Ops: complexity isn't the only potential source of failure
  - ♣ Buggy implementations hinder adoption (the deployment problem again)
  - ♣ Comparisons to IP multicast
- ♣ Lack of solid debugging tools

## Other RIPQoS papers

- ♣ Carlos Macian: *Beyond Technology: The Missing Pieces For QoS Success*
- ♣ Difference of opinion between ITU and IETF on what QoS is
- ♣ QoS isn't just a network issue - apps have an impact too.
- ♣ Need a number of business environment pieces for QoS to succeed
  - ♣ Billing & accounting model
  - ♣ Trust model

## Other RIPQoS papers

- ♣ Ben Teitelbaum: *Quality of Service and Denial of Service*
- ♣ QoS solutions have to be designed for worst case conditions
- ♣ Not attractive to customers if it's primarily helping the ISP stretch their bandwidth budget
- ♣ Elevated priority schemes may not provide DoS protection
  - ♣ If they don't, they're essentially undeployable
- ♣ Researchers need to think adversarially
- ♣ Open question: how do customers verify DoS protection?

## Other RIPQoS papers

- ♣ Tristan Henderson: *Networked Games --- a QoS-sensitive application for QoS-insensitive users?*
- ♣ Research into the satisfaction of online gamers
  - ♣ Most games use UDP-based protocols
  - ♣ To gamers, latency is all that matters
- ♣ Turns out, gamers don't care about QoS once connected
  - ♣ Higher latency will keep them from joining a server, though
- ♣ How do you get price-sensitive customers to pay more for a service they're insensitive to once connected?

## Other RIPQoS papers

- Panos Gevros: *Internet Service Differentiation Using Transport Options: the case for policy-aware congestion control*
- A means to control TCP congestion control dynamically in an efficient manner
- Existing QoS schemes won't survive the combinatorial explosion of application to the whole internet
- Instead, modify endpoints' congestion control parameters in response to changing network conditions
- ISPs would offer this service to end users



## Additional Sources

- *Security Issues in a Differentiated Services Internet* - A. Striegel, Department of CS and E, Notre Dame
- [http://www.cisco.com/en/US/tech/tk543/tk766/technologies\\_white\\_paper09186a00800a3e2f.shtml](http://www.cisco.com/en/US/tech/tk543/tk766/technologies_white_paper09186a00800a3e2f.shtml)
- <http://www.unc.edu/~gogan/qos/>

