

Revisiting IP QoS: Why do we care, what have we learned?

ACM SIGCOMM 2003 RIPQOS Workshop Report

Grenville J. Armitage, Workshop Chair
Swinburne University of Technology
Melbourne, Australia
garmitage@swin.edu.au

Abstract- ACM SIGCOMM 2003 included a number of workshops, including the all-day workshop “Revisiting IP QoS: Why do we care, what have we learned? (RIPQOS).” The goal of RIPQOS was to critique the evolution and deployment of IP quality of service (QoS) mechanisms, from both the research and operational community perspectives. The workshop's name was a challenge to all interested communities to reflect on whether IP QoS has lived up to the hype or whether it is simply misunderstood. The workshop saw 6 papers, 2 short papers, a discussion panel, a range of opinions and lots of questions. This report attempts to capture the essence of our workshop's discussions, presentations and experiences.

Keywords- IP, Sigcomm, Workshop, RIPQOS

I. INTRODUCTION

There are few topics in the IP networking research and operational communities that elicit as much inconsistent opinion as “quality of service” (QoS). The very premise of QoS appears, on the face of it, to contradict the guiding principles of “best effort” service, a service model that has seemingly underpinned IP network engineering since the very beginning. Segments of the research community have taken the complexity and apparent contradiction as a challenge, and produced a substantial body of strong theoretical work showing how IP networking can evolve to support a variety of QoS schemes. Large segments of the operational community simply cannot see the point of adding QoS to networks that are humming along quite nicely as they are. A broad spectrum of people can't entirely agree on what QoS actually is. What's going on here?

The RIPQOS Call for Papers deliberately began with a provocative statement:

“For over a decade the Internet engineering and research community has debated, designed, and ignored IP Quality of Service tools and techniques. There's a sense that something might be needed, but little agreement on why and who will pay. At times the very notion of QoS has seemed to be a pointless waste of time, almost a solution waiting for a problem. This workshop is an opportunity for researchers and practitioners to discuss the history of IP QoS research and development, review what could have been done better, and perhaps develop a new focus going forward.”

We went on to give some specific questions that this workshop might consider:

“Papers are invited that provide well-argued opinion, speculation, or contrary positions. For example:

- *IP QoS schemes never quite seem complete. Is this just a great research game for academics?*
- *Where's the money? How do we make IP QoS pay when typical Internet applications don't care, and the user's don't know any better?*
- *Will online, multi-player games be the market segment that justifies end-user/access ISP investment in IP QoS tools and solutions?*
- *Isn't more bandwidth the answer?*

Of particular interest are papers that critique the evolution of IP QoS solutions to date and/or explain what sort of applications and user mindset will need to emerge before IP QoS solutions become cost-effective for ISPs to deploy.”

A bit pointed? Yes. Totally unfair to the research community? Well, no, not really. Our goal with RIPQOS was to start dialog on how IP QoS techniques and methodologies fare in the operational world outside of simulations and testbeds. The research community has taken IP QoS a long way. The question we wanted to consider at RIPQOS was whether we're entering a brave new world of QoS deployment or whether IP QoS should simply Rest In Peace.

In the end RIPQOS enjoyed 6 full papers, 2 partially developed papers and an invited panel discussion to wrap up the day's proceedings. The first four papers were grouped into two morning sessions under the heading “Challenges” - opinions on where QoS research should be going, a question of whether QoS has “failed to thrive”, observations on how QoS deployment must attend to operational and commercial realities, and a review of how DiffServ is solving real-world problems for real ISPs and customers today. After lunch we had two papers under the “Lateral Thinking” session, looking at two quite diverse topics of QoS and Denial of Service, and the potential for networked Games to emerge as an important QoS-sensitive application. The “Short Papers” session saw a discussion about QoS as a risk management tool and a brief proposal to add variable congestion control algorithms to existing transport mechanisms. Our discussion panelists wrapped up the day by considering the question of where the research community should go next in order to expand the role of IP QoS in operational IP networks.

The workshop proceedings are available through the ACM Digital Library [1]. What follows here is a summary of each session.

II. CHALLENGES

Two sessions covered the Challenges, four papers articulating some different perspectives on IP QoS past, present and future.

A. *QoS's Downfall: At the bottom, or not at all!*

Jon Crowcroft (University of Cambridge) opened the day with a walk down memory lane to the Cambridge Ring, a “bogograph” with broad sweep armwaving, a reminder that QoS deployment means meeting the needs of stakeholders, and an argument that some minimal degree of QoS mechanism needs to be embedded in the very lowest levels of our networks [2].

Jon began by clarifying that, clearly, all networks deliver some sort of “quality”. But in the sense we're discussing QoS here, the issue is about a network's ability to offer different (and differentiable) levels of service quality over a shared infrastructure.

More importantly, there's a real problem for researchers in this field that isn't always obvious. Jon argued that we become trapped at particular points in the cyclical nature of the problem statement. The ratio of access network and core network capacities change over time, moving the congestion (and hence QoS) problem back and forth. QoS research tracks resource constraints in networks. Hence any given researcher's work tends to be a trailing indicator of where the access/core ratio stood at the time they embarked on their particular QoS scheme.

Many of the different viewpoints held firmly by people across our networking community may perhaps be understood in the context of the prevailing edge/core capacity ratio at the time each person embarked on their network engineering and research careers. Failing to realize that we're victims of the “wheel of time” will ensure we do not break out of the cycle. [There was also a bogograph (“bogus graph”) to back it all up, showing a sinusoid of “core to access capacity ratio” with a period of roughly twenty years.... and, as Jon readily admitted, scaled primarily to suit his talk.] Nevertheless, an intriguing position that our research is blinkered by historical trends.

Another issue identified in Jon's talk is the question of knowing who your stakeholders are when discussing, designing and proposing QoS schemes. Computing people, telecoms people, service operators, users/consumers.... they all have different timescales over which they evaluate the cost-benefit trade offs of new schemes and services. The drive to “converged” networking means a larger and more diverse range of applications, which has created demands for a multi-service (and hence QoS-enabled) IP layer. But each stakeholder has arrived at different points in the historical capacity ratio cycle, and thus perceive the technological problems differently.

Jon's fundamental points came towards the end:

- We need QoS at the lowest layer, below IP, and it only needs to be simple – a two-level (one bit) scheme will suffice. Overlay schemes cannot support “better” QoS control if the underlying links (networks) don't have at least two levels of service.
- The lowest level QoS mechanism needs to be exceedingly cheap to implement and deploy, encouraging innovative use with minimal inconvenience.
- There's no need for more “QoS architecture” work at higher levels, it's been done. We need to map QoS to marketable services, close the gap between mechanisms and revenue.
- We need QoS mechanism in the core. The current core/access ratio is trending towards the core being a congested resource, despite all current evidence to the contrary. Aim for the future, don't chase the near-term issue. Put two-level QoS into the optical core.

In other words, QoS needs to be deployed “bottom up”, yet the mechanisms at the lowest levels need not be complex at all. QoS needs to be “sold” correctly to the various stakeholders – if we hang around at “layer 8” (the political layer of the OSI stack) we'll continue to solve for congestion/capacity limits that have changed by the time our solutions appear for deployment.

B. *Failure to Thrive: QoS and the Culture of Operational Networking*

Gregory Bell (Lawrence Berkeley National Laboratory) took us on a different tack – he provided an operational engineer's contemplations on why IP QoS schemes have failed to thrive in enterprise networks, and the entrenched R&D methodologies that may yet ensure no complex IP QoS schemes ever manage to take off [3]. He observed that his insights are taken from supporting a research institution with about 80 subnets and 10,000 hosts – an enterprise rather than ISP environment.

“Failure to thrive” comes from the growth of children - thriving is the norm and anything else is cause for alarm. In networking it might be argued that withering is the normal course of events for most protocols and architectures. Yet Greg observed that IP QoS should be thriving by all accounts, considering the stature of the researchers who have done significant work in the area over the past decade and the abundance of literature generated in the area.

So why has IP QoS failed to thrive?

There appears to be a structural rift between the designers of protocols and QoS architectures, and those whose job it would be to operate such QoS-enabled networks. There's a disconnect in how researchers and operations people view the relationship between system complexities and system failures. Both communities recognize that complexity multiplies the potential for system failures. What the research community doesn't often internalize is the fact that failures in deployed equipment are often due to buggy implementation rather

than unintended operation of a perfectly implemented protocol. Operations engineers learn the hard way to assume failures will occur and act to avoid deployment of anything that increases the chance they'll be working weekends to fix their network.

So the real question faced by IP QoS community is "is this deployable?" Greg gave us some history of a related service with a checkered past – IP multicast. At LBNL the enterprise networking team had a range of bad experiences with routers running shipped, yet buggy, IP multicast code. Greg's main point is that the IP QoS research community needs to recognize that operations engineers are as afraid of buggy QoS implementations as they are of the architectural issues surrounding QoS-enabled service deployment.

Deployment thus depends on a variety of factors - QA by vendors, critical mass of users, debugging tools, enterprise knowledge, trust between neighbouring domains, and business case. For the operational staff the existence of debugging tools is crucial – if I deploy a new technology to offer a new service, where are tools I can use to cover myself when things go wrong (network failure) or users complain they aren't getting the XYZ service level they thought they would? And a related broad concern, can this new technology be deployed incrementally in such a way that it doesn't disrupt existing best-effort IP service?

Faced with these concerns it is hardly surprising that many network administrators will seek out the safer route of just adding more bandwidth. Greg sounds a cautionary note to the QoS community not to dismiss "throwing bandwidth at the problem", because in reality this is often a reasonable and pragmatic engineering response to network congestion. It certainly seems more appealing to network engineers than "throwing protocols" at the problem! The reality in many enterprise networks is that adding links is a well understood and low-risk action that usually solves impending congestion issues. (Greg acknowledged that the economic and practical constraints facing ISPs may differ greatly from those in enterprise networks, he discussed some examples where long haul telco links simply couldn't be upgraded easily.)

Perhaps the essence of Greg's cautionary tale is contained in this quote from his paper:

"Attempting to architect QoS without taking into account its economic and institutional context is roughly analogous to designing a city with reference to local culture, climate or geography."

Greg's talk concluded with some animated discussion of his conclusion that unless the research and development communities start paying attention to the deployment issues, QoS will continue to suffer from a failure to thrive.

C. Beyond Technology: The Missing Pieces for QoS Success

After the morning break Carlos Macian (University of Stuttgart) gave us a view from the other side [4] – the issues that affect ISPs more so than the enterprise network focus on Greg's talk.

Carlos began by observing that the ITU and IETF have different perspectives on QoS - the former seeing it as the collective service performance affecting the user's satisfaction, the latter seeing it as a set of end to end performance metrics to be met along a network path. Yet there's a common issue – QoS is a "...measurable performance that satisfies some(one's) demand".

Although reliability, availability and security are also metrics by which QoS can be evaluated, Carlos specifically focused on the timescales of packet loss and jitter. He also noted that QoS is not simply a network issue – the end user applications have a significant impact on the end user's sense of service quality. Thus QoS is truly an end to end issue.

And the main motivation for ISPs to deploy QoS is money. How to provide priority service to someone who would be willing to pay for better than best effort?

Overprovisioning is often not a viable option for ISPs, ultimately boiling down to economic realities. Links of suitable capacity may be available and yet priced too high to justify, or there may simply be no links available between different sites on a provider's IP network. Alternative mechanisms are then deployed to share the existing links in a controlled and differentiated manner (using the more general notion of "differentiated", not the IETF's Differentiated Services model necessarily).

To be successful QoS needs a mixture of technical mechanisms (buffer allocation, capacity, protocols) and economic consideration (price differentiation, market segmentation, service bundling, perhaps auctions).

Although the technological pieces may be largely in place, the business environment still has a long way to go. The solution needs to be completely end to end or not at all, which implies inter-domain relationships to support QoS across provider boundaries. This creates a hugely complex problem for relationships between ISPs – indeed, the traffic exchange relationships need to be more explicit and formalized than they are today. Network operators may need to expose internal performance data about their networks, which is understandably sensitive information. Will we end up with market consolidation? Enforced interconnect rules through regulation? Stagnation by simply letting overcapacity solve QoS where it can and ignore the rest of the network?

So back to a critical question – how can money be earned? The QoS area seriously lacks a billing and accounting model that can work inter-provider and inter-domain. The telephone industry is not a good guide because it essentially offers only one level of service – either the phone call is there, or it is not. Billing and accounting are thus simplified by aggregating call durations. There's no simple analogy to IP networking. And since telephony providers are already making far more profit than ISPs, why would telephony migrate to IP? Some IP architects have proposed "brokers" to handle accounting and billing between domains, but this is still very much work-in-progress and doesn't address the need for compelling business reasons for domains to play nice with each other.

In the end Carlos summarised that we have no integrated QoS architecture, but even more importantly the community needs to address business models and trust models between providers.

D. Deployment Experience with Differentiated Services

Rounding off the morning was Bruce Davie (Cisco Systems) with an entertaining and far more positive position on IP QoS – despite the doom and gloom from certain quarters there is a thriving deployment of basic Differentiated Services (DiffServ) in real-world networks solving real-world problems today [5].

Bruce started off with a fair observation that the very premise of RIPQoS was somewhat confrontational to the QoS community. He felt obliged to step forward with a DiffServ success story and to ask whether there were lessons here for the broader deployment of IP QoS mechanisms.

There are at least two valid definitions of QoS – an application-based “what can I get from the network?” or mechanism-based “what technologies can I utilize?” Bruce’s talk would focus on the use of certain mechanisms on the presumption that overprovisioning was not the solution (otherwise we should simply stop all QoS R&D work). Traffic engineering was not defined as a QoS mechanism for this talk. (Bruce also observed that if the network offers good-enough service without additional service levels this could be considered “QoS”.)

The first critical clarification Bruce made was to observe that most DiffServ deployment today was in private networks or managed VPN contexts. It is commonly deployed by ISPs supporting customers who have high expectations of mixing Voice over IP (VoIP) and regular IP traffic over shared links running at very high utilizations. Usually the customer is driven to run their external links at high load because such links are exceedingly expensive. There’s little to no DiffServ deployed in the core of the public internet.

VPN scenarios are tractable because they avoid the issues of maintaining QoS across inter-domain or inter-provider boundaries, or having to mix IP traffic from uncontrolled and uncontrollable sources. The VoIP sources and sinks are usually under the control of the customer directly asking for QoS, or the provider offering the QoS-enhanced service.

Typically the VPN customer has a number of remote sites connected to the VPN provider’s core network through low bandwidth links. VoIP traffic is marked with the DiffServ “Expedited Forwarding” (EF) bit, and routers on the customer premises edge and provider edge provide priority treatment to packets with the EF bit set.

DiffServ is an exceedingly useful tool when edge bandwidth is an expensive commodity. It is especially useful when migrating customers from older private networks – such customers expect their voice trunks to operate much as they did when using e.g. Frame Relay, even though they’re now sharing their external links with other intra-VPN IP traffic.

At least one national carrier (Telecom Italia) has deployed an entirely self-contained VoIP network for public telephone calls, with DiffServ turned on. Since the VoIP gateways are under the carrier’s control, provisioning and policing are tractable problems.

Bruce then considered the situations where DiffServ is not attractive, or very difficult to use. For example, many ISP customers do not have the same expectation of service quality as someone who migrates from traditional circuits like Frame Relay. Regular IP access is a difficult environment to offer QoS assurances because we run into inter-provider issues, and it is hard (or impossible) to establish a-priori the set of communicating sites (and thus almost impossible for the ISP to internally provision their network correctly).

It is also hard to sell QoS mechanisms into environments where Best Effort service appears to be working just fine (e.g. where bandwidth is not a problem) – what is the benefit to a customer of a premium service if BE is fine? Who would want to be the first ISP to downgrade their BE service below their competitor’s just to offer a “premium” service that’s about as good as their competitor’s BE service?

DiffServ is also a hard sell when the customer only expresses their needs in terms of end to end performance, rather than per-hop behaviors (PHBs). And finally, inter-provider QoS agreements are really hard!

Bruce wrapped up by observing that future QoS research needs to look at deployment issues rather than develop new mechanisms. The catalysts for future deployment of QoS will include a subsiding of the bandwidth glut, development of customer demand, regional ISPs co-ordinating together, and the development of standard service definitions (to allow comparison shopping by customers).

III. LATERAL THINKING

Coming back from lunch we jumped into the Lateral Thinking session with two quite diverse topics – Denial of Service and the potential for networked Games to emerge as an important QoS-sensitive application.

A. Quality of Service and Denial of Service

Ben Teitelbaum (Internet 2) introduced an interesting argument that QoS solutions will never see meaningful deployment if they are not designed for worst-case conditions – i.e. when network components are suffering from denial of service (DoS) attacks. In other words, most QoS researchers should be designing QoS schemes that protect against (and work in the presence of) adversarial conditions. However, if QoS is deployed to protect against DoS then how do customers ever verify they’re getting DoS protection?

Ben’s fundamental definition of QoS is the regulation of the impact of congestion. The reality of today’s IP networks is that “best effort” is generally pretty good, and most of the time operators find “adding bandwidth” to be the operationally and financially pragmatic solution to growth in traffic loads. QoS schemes are attractive only insofar as they offer protection against

service degradation during worst-case network traffic load conditions. QoS is not attractive if deployed primarily to optimize an operator's use of their fibre capacity.

Ben identified two broad classes of QoS - "elevated priority service" and "non-elevated priority service". The former delivers treatment equal or better than the default best-effort service (e.g. based on the DiffServ Expedited Forwarding or Assured Forwarding models), while the latter provides something equal or lesser than default best-effort service (e.g. scavenger services that make use of spare capacity without disrupting existing best effort traffic). Ben observed that only elevated priority services can protect against adversarial traffic overload conditions. Although easy to deploy, non-elevated and default best effort can both collapse to zero service under suitably severe congestion conditions.

Defining a DoS attack is also problematic. Is it a security or resource management issue? How does an operator discern the difference between legitimate and adversarial increases in traffic load? If I'm a researcher pushing the limits of a gigabit link am I "attacking" people along the path during my test? The mere fact that a network resource begins to congest or become overloaded cannot be used as a criteria, so it becomes a question of intent - a difficult concept to infer with completely automated tools. Ben suggests that operators instead look to protecting certain traffic known to be legitimate rather than attempting to automate the detection of illegitimate traffic. This is QoS by any other name.

Elevated priority schemes may or may not provide protection against DoS (e.g. against a third party attempting to disrupt the service guarantees offered between any two hosts by pushing the network conditions outside normal parameters). Unfortunately, any elevated priority scheme that does not protect against DoS will be essentially undeployable - it adds cost and complexity while offering almost no benefits over best-effort during good periods and no guarantees when the network is overloaded.

The costs and complexities of elevated priority schemes that can protect against DoS are similar to ones that do not, yet the benefits are tangible during times of network overload. Ben observed that this makes such schemes potentially deployable if the costs are low enough. In other words, a QoS scheme will only be deployable and attractive in the field if it inherently protects legitimate traffic from all other sources of traffic competing for network resources - legitimate or otherwise.

In essence DoS management is a QoS problem. Ben identified a number of ways that QoS researchers need to re-orient their thinking on this point. First, start thinking like an adversary. Consider the various ways an adversary can disrupt traffic flows along your statistically provisioned network paths (e.g. compromised hosts elsewhere on your network) and ask how existing QoS technologies could mitigate the impact of a DoS attack. Develop a minimally complex

QoS architecture that protects against DoS, then we'll have a better chance to encourage deployment.

Of course, Ben then pointed out the obvious problem - customer verification of the offered service. How can a customer confirm they're getting protection from DoS? Generating a DoS attack on their own ISP is unlikely to be met with much delight in other quarters.

Ben's talk left us with more questions than answers. He definitely challenged us to consider DoS-protection as the most viable selling point for QoS technologies, because under "normal" network conditions best effort service and engineering practices are usually the simpler solution.

B. Networked games --- a QoS-sensitive application for QoS-insensitive users?

An entirely different talk was presented by Tristan Henderson (University College London) on the issue of whether networked games really demand QoS as many people have asserted. Tristan presented the results of some research he'd done into people's satisfaction with playing Half-Life/Counterstrike online.

Tristan first observed that online games could be broken into three broad categories:

- First Person Shooters (FPS, such as Quake3, Half-Life, or Doom)
- Massively Multiplayer Online Role Playing Games (MMORPG, such as Everquest or Star Wars: Galaxies)
- Real Time Strategy (RTS, such as Civilization or Age of Empires)

Typically these games would use UDP-based client-server communication. Network latency is broadly a major concern for each category, more concerning for fast-paced interactive games. A range of studies unrelated to online gaming have shown 100ms to 300ms as the acceptable range for round trip delay in human interactions. Some game related studies have claimed limits down around 100-250ms.

Tristan's question was basically whether players actually really cared about QoS, and in what terms did they actually articulate their thoughts about QoS. He made the observation that although we intuit a need for QoS to make games popular, online games appear to be remarkably successful anyway in today's IP networks. So is there really a need for "good" QoS for games?

Tristan set up two Half-Life/Counterstrike servers in London, equivalently connected to the Internet except that additional controlled latency could be added to one or the other server. The servers both ran for a few months to build up their popularity, and then Tristan began adding nominal latency of 50ms to one and then the other server. Usage patterns clearly showed that simply adding 50ms would discourage people from even joining their server (client software has a method for potential players to rank servers according to network latency before playing).

The server logs also showed that players tended to stay on a server once they'd joined even when network latency got worse for short periods of time. Interestingly, the relative increase in delay didn't seem to affect the likelihood of a player leaving. Players who'd already been playing a long time seemed to be more tolerant of brief bursts of additional delay (perhaps tolerance increases with immersion in the game). But note that players who play regularly (as opposed to players who've been logged into a single playing session for a long time) are no less likely to leave when delay gets worse than a player is not a regular on the server.

Finally, there was clear evidence that a player's "success" (in terms of kills, or "frags" per minute) was adversely affected by increased latency, as was the likelihood of the player dying frequently.

Tristan used these insights to ask an unsettling question – if players are relatively insensitive to network delay once they've decided to join a game server, how will an ISP attract these typically price-sensitive customers to a premium-cost IP service with improved QoS? Or put another way, will a player pay for QoS improvements when they appear insensitive to QoS degradation? (Tristan also commented that in some earlier work he'd surveyed game players who indicated a strong unwillingness to pay for network QoS on principle.)

Tristan noted that his study only considered absolute delay, and did not test player's sensitivity to jitter or packet loss. There's more research to be done to better understand if those QoS parameters have more influence on player satisfaction than raw latency. There's also some interesting questions about how an ISP might vary the QoS delivered to any particular customer based on how far "into" a game they are (e.g. good at the beginning (to attract players) and then gradually declining to a tolerable level as the game progresses).

IV. SHORT PAPERS

Two short presentations came after the afternoon break, representing ideas under development at the time of review.

A. *What QoS Research Hasn't Understood About Risk*

Ben Teitelbaum (Internet 2) came back to the stage with a short discussion that he acknowledged wasn't necessarily consistent with his first presentation. (The chair also noted that it was a quirk of double-blind reviewing that led us to have two papers by Ben!)

Ben's basic thesis is that neither customers nor ISPs need or want hard performance guarantees. Rather, each wants tools and understand and manage risk. He observed that the design goals for much QoS research can be captured by words attributed to S.Keshav, "*The Holy Grail of computer networking is to design a network that has the flexibility and low cost of the Internet, yet offers the end-to-end quality-of-service guarantees of the telephone network*". Ben's contention is that this design goal completely misunderstands the market's real use for IP

QoS, and that engineering for true end-to-end QoS destroys the apparent cost advantage of the Internet.

QoS is essentially a cost/benefit proposition. Many technical solutions for IP QoS focus on congestion-management or congestion-avoidance. However, avoiding congestion entirely (or close to entirely) can be a rather expensive proposition. What other tools can be brought to bear on the process of managing the risks associated with using a congested network?

Customers are essentially rational, although they have complex and divergent utility functions. They typically want the ability to trade-off between a number of service criteria, including good performance, low costs, simple pricing, low transaction overheads, and means to manage exposure to worst-case performance. Technical QoS schemes can only remove a component of risk exposure, and at significant cost.

Perhaps an alternate goal for QoS research is "*How can network services offer customers and providers flexible management of exposure to poor network performance?*"

A complete risk management solution needs to include economic tools. There are few (if any) markets where customers actually demand infallible service. Typical systems use a mix of technical and economic/regulatory mechanisms (e.g. warranties, insurance, and certification). Warranted performance appeals to customer's desire for simplicity in costs, but demands new tools for providers to correctly estimate their ability to warrantee any particular service agreement. The likelihood of customers demanding compensation for performance degrading outside warranted levels leads to opportunities for third-party performance insurance. Insurance is a monetary risk management tool, there's no reason it cannot be applied to the network service provision industry. Yes, this would introduce a whole new insurance industry and require government involvement to establish rules for liability, etc. But a necessary evil. And yes there will be costs to administer, but in most places insurance acts to lower risks for goods and services. And finally, certification processes and procedures would complete the tool-kit, allowing customers to compare ISPs and requiring performance monitoring and reporting by ISPs along the lines of the phone companies today.

QoS needs a multi-disciplinary approach from this point forward if we are to see any truly deployed solutions at all.

B. *Internet Service Differentiation using Transport Options: the case for policy-aware congestion control*

Panos Gevros (University of Cambridge) presented a rather different discussion, on technical means to optimise TCP's congestion control behavior in a dynamic fashion.

Panos first observed that QoS allows us to differentiate between traffic belonging to different users and to guarantee quantifiable performance levels. Most mechanisms discussed so far have been network-centric, e.g. end-to-end guarantees, router mechanisms such as

DiffServ, etc. Such mechanisms may be OK for small scope environments, but there is a combinatorial explosion when we try and use QoS for the whole Internet.

Panos' basic thesis is that we should engage the endpoints in the control loop – modifying each endpoint's congestion control behaviour to influence performance. Many networks are well-provisioned. So if there is no congestion, users might be able to transmit more aggressively (compared to regular/default TCP as currently deployed). We could modify TCP's slow start, congestion avoidance or retransmission timeouts.

The idea behind “transport options” is that an ISP could offer a service that changed the congestion control behaviour of endpoints dynamically, on behalf of the user, according to certain classes of service offered by the ISP to the user. Classes of service are offered in relative terms (no guarantees) - so you know that a class is worse than another, but that is all.

This work on transport options is still in progress, and has a number of open issues. It offers to put ISP congestion policy control further out to the network edges and into the operating systems of end users.

V. DISCUSSION PANEL

We wrapped up the day with a short, informal panel discussion involving Jon Crowcroft (University of Cambridge), Bruce Davie (Cisco Systems), Jennifer Rexford (AT&T Research) and Ran Atkinson (Extreme networks). As chair I basically let each person comment on what their main thoughts were after participating in the day's event.

Jon: Control planes are complex - a signalling network is as complex as the network itself. We need a distributed partition to partition BE from EF/AF. But how to bootstrap? Network is ever-evolving - a moving target. Need a new model of the evolution of the network so we can predict where this moving target will be?

Bruce: Range of opinions today, from DiffServ is a success to DiffServ is a fundamental failure to QoS is a solution in search of a problem. Is the glass half-full or half-empty? How can we get Internet-wide QoS?

Jennifer: How can you have QoS when:

- A typo by a network operator can bring down service? Half of network outages come from misconfiguration
- Routing anomalies might throw away your traffic?
- Users don't know how to predict their bill?
- You don't know who to blame for a QoS violation?

Ran: Universities are more interested in degrading traffic – e.g. degrading game or p2p traffic. Degrading is better than ACLs since applications will port-hop if blocked. 5-10% of universities employ some sort of ACL to downgrade undesirable traffic. Some universities use quota-based systems, e.g. a rate-limiting up to a quota per IP address.

Extreme's products are QoS feature-rich, but most customers don't switch those features on. Ran thinks there is less QoS out there than Bruce does.

In the late 80s the DDN (Defence Data Network) in the US, based on T1s used John Moy's OSPF which supported the TOS bit. So e.g. file transfers would have lower precedence and be routed via a satellite (which had higher capacity). TOS bit was used as in RFC791.

Sprint will (for a fee) use CBQ between customer routers and the aggregating router. DSCP bits are cleared in the core because in Sprint's core there is a funny DiffServ setup for prioritising BGP traffic (as opposed to customer traffic). Also Sprint core is overprovisioned.

The IAB Network Management workshop had two points: operators desperately need better config/management tools (most ISPs use a set of Perl/MySQL/Tcl/Expect scripts to manage routers), and reducing operations costs.

Bandwidth is increasing faster than our ability to encrypt or authenticate data. This is a problem for interdomain QoS since ISPs will be unwilling to forward marked packets unless they can verify that it is not a DoS.

Maybe ISPs could just negotiate levels of EF between each other. e.g. renegotiate every day, if a customer sends too much traffic, that is their problem. Need to be able to assign blame.

As to each panel member's favourite research problem:

Jon: a model of the complexity of network architectures (not just the complexity of components)

Bruce: interprovider QoS

Jennifer: configuration management - models of protocol configuration state

Ran: configuration management - how do I configure a network (not just a single box)

kc: how to get funding? The NSF won't consider a lot of these problems as research. Maybe Dave Clark's knowledge plane could be used as a platform for research.

What is left for "traditional" QoS research? Bruce: no more queuing algorithms!

VI. CONCLUSIONS

It would be inappropriate to conclude that RIPQOS has answered everyone's questions about QoS. But the day did see two broad themes emerge:

- There's a lot of respect for the complex theoretical work that has been done on device-level congestion management schemes in bursty, best-effort IP networks
- If IP QoS is to be truly deployable the research and development communities need to shift gears and begin answering the market's actual questions – use

the technical device-level mechanisms to develop systems-wide toolkits for monitoring and managing QoS schemes, and recognize that QoS is only part of what customers demand from their ISPs.

The key message from RIPQOS is that QoS is not dead, but as an IP QoS R&D community we need to reach out and include business, systems control, and marketing expertise in our efforts to get IP QoS meaningfully deployed and used.

ACKNOWLEDGMENTS

This report relied heavily on our RIPQOS scribe, Tristan Henderson, and informal post-workshop discussions with kc claffy, to supplement my flakey memory. All errors of transcribing are, naturally enough, mine.

The workshop itself could not have materialised if it were not for the RIPQOS program committee - Mark Allman, kc claffy, Tristan Henderson, Geoff Huston, Derek Mcauley, Kathie Nichols, John Wroclawski, and Sebastian Zander. The ACM SIGCOMM2003 Workshop organizers (and in particular Craig Partridge) also deserve thanks for first believing in this workshop and then providing logistical support for workshop advertising, website and registrations.

And finally, the workshop would have been nothing without our presenters (listed below) and our discussion

panelists at the end of the day – Jon Crowcroft, Bruce Davie, Jennifer Rexford and Ran Atkinson.

REFERENCES

- [1] "Proceedings of the ACM SIGCOMM workshop on Revisiting IP QoS: What have we learned, why do we care?," August 2003, <http://portal.acm.org/citation.cfm?id=944592&coll=ACM&dl=ACM&CFID=13079983&CFTOKEN=66295760> (URL as of October 2003)
- [2] Jon Crowcroft, Steven Hand, Richard Mortier, Timothy Roscoe, Andrew Warfield, "QoS's Downfall: At the bottom, or not at all!" Proc. ACM SIGCOMM 2003 Workshops, p. 109, August 2003
- [3] Gregory Bell, "Failure to Thrive: QoS and the Culture of Operational Networking," Proc. ACM SIGCOMM 2003 Workshops, p. 115, August 2003
- [4] Carlos Macian, Lars Burgstahler, Wolfgang Payer, Sascha Junghans, Christian Hauser, Juergen Jaehnert, "Beyond Technology: The Missing Pieces for QoS Success," Proc. ACM SIGCOMM 2003 Workshops, p. 121, August 2003
- [5] Bruce Davie, "Deployment Experience with Differentiated Services," Proc. ACM SIGCOMM 2003 Workshops, p. 131, August 2003
- [6] Stanislav Shalunov, Benjamin Teitelbaum, "Quality of Service and Denial of Service," Proc. ACM SIGCOMM 2003 Workshops, p. 137, August 2003
- [7] Tristan Henderson, Saleem Bhatti, "Networked games --- a QoS-sensitive application for QoS-insensitive users?," Proc. ACM SIGCOMM 2003 Workshops, p. 141, August 2003
- [8] Ben Teitelbaum, Stanislav Shalunov, "What QoS Research Hasn't Understood About Risk," Proc. ACM SIGCOMM 2003 Workshops, p. 148, August 2003
- [9] Panos Gevros, "Internet Service Differentiation using Transport Options: the case for policy-aware congestion control," Proc. ACM SIGCOMM 2003 Workshops, p. 151, August 2003