

June 2012
Geoff Huston

The QoS Emperor's Wardrobe

Back in 1997, with Paul Ferguson, I wrote a book on "Quality of Service" (QoS) in IP networks. A couple of years later I pushed out a revision of the book ("Internet Performance Survival Guide") that looked more generally at service performance in IP networks, and examined in voluminous detail the interactions within the network that impacted on the quality of the network service that was delivered to applications. As it turned out I was a little surprised that either book sold a single copy!

Why?

Because in writing the book it became painfully apparent that Quality of Service was in fact more myth than reality for the public Internet. When we take the public Internet and look at QoS there is an glaring credibility gap: we can't build it, and applications can't use it. If you really think that the network itself is the problem and QoS is the answer, then there is always another, very simple, response: get more bandwidth. That's as true now as it was almost twenty years ago. Nothing has changed.

So why revisit this topic now?

There is an industry group in Europe, the European Telecommunications Network Operators (ETNO) (<http://www.etno.eu>), which is an association of a number of network operators. In June of this year they released their proposal for changes in the International Telecommunications Regulations (ITRs). To quote from their web site:

This contribution calls for a new IP interconnection ecosystem that provides end-to-end Quality of Service delivery, in addition to best effort delivery, enabling:

- the provision of value-added network services, to both end-customers and over the top (OTT) players and content providers, and
- a reflection of the value of traffic delivery over network infrastructures.

Moreover, the contribution states that in order to ensure an adequate return on investment in high bandwidth infrastructures, operating agencies shall negotiate commercial agreements to achieve a sustainable system of fair compensation for telecommunications services.

By endorsing the concept of "quality based delivery", it will be possible to establish new interconnection policies based on the "value" of the traffic (not only on the "volume"), enabling new business models and implementing an ecosystem where operators' revenues will not be disconnected from the investment needs made necessary by the rapid growth of Internet traffic.

[<http://www.etno.be/Default.aspx?tabid=2500>]

As I understand their proposal, ETNO is advocating a new IP interconnection framework that is quite fundamentally based on this concept of end-to-end Quality of Service.

This immediately raises a question in my head: How "real" is end-to-end Quality of Service in the Internet?

Much has been written over the past couple of decades about the potential of Quality of Service (QoS) and the Internet, including some of my own personal contributions to this mountain of verbage. However, much of this material is strong on promise, but falls short in critical analysis. In an effort to balance the picture, lets take a brief recap on the various efforts to bring QoS to the Internet.

The QoS Service

The default service offering associated with the Internet is a best-effort service, where the network treats all traffic in exactly the same way. There is no consistent service outcome from the Internet best-effort service model. When the load level on the data path is low, the network delivers a high quality service. As the instantaneous load levels increase beyond the carriage capacity of the network, the network congestion levels increase, and service-quality levels decline along this data path. This decline in service is experienced by all traffic passing through a congestion point, and is not limited to the most recently admitted packets or recently established traffic flows.

For many applications, this best-effort response is perfectly acceptable. When network capacity is available, the application can make use of the resource, whereas when the level of contention for network bandwidth is high, each application will experience similar levels of congestion, and should adapt to the changing circumstance. A best-effort network service is a good match to opportunistic applications that can vary their data transfer rate in response to signalled network load.

The analogy to cars and a road system is pretty good here. Outside of rush hours a trip can be very efficient, but when the road system fills with cars, each car's journey time deteriorates. The amount of deterioration depends on the level of congestion on a particular route.

The objective of various Internet QoS efforts is to augment a best-effort service with a number of selectable service responses. We can juggle with the two basic actions of an active network router: queuing and discarding. Packets queued at the head of a queue should be processed faster than those queued at the end of the queue. And if a switch has to discard a packet, then some QoS-related rule set may be used to select which particular packet is the preferred candidate for discarding. In altering the queuing and discard behaviours selectively, service outcomes that differ from the best-effort service may be generated. These 'tailored services' may have lower average end-to-end delay, lower jitter, or greater bandwidth. These service responses are relative, where the service outcome is claimed to be no worse than best-effort at any time, and superior to best-effort under conditions of congestion load.

Alternatively, QoS service responses may be distinguished by providing a consistent, and therefore predictable, service response that is unaffected by network congestion levels. These are quantitative service responses, where the characteristics of the service can be measured against a constant outcome. A quantitative service may be one that constrains jitter to a maximum level, or one that makes a certain bandwidth available, within parameters of bounded jitter. Such constant-rate services may be better than best-effort services when the network is under load, but they may also be worse than best effort when the network is unloaded.

What motivates the provision of differentiated service profiles within the Internet? One could attribute much of the motivation to a desire to provide a network client with a range of service-quality levels at a range of prices. Obviously this is a broad agenda, where there are requirements to extend specific network services to applications, requirements to adapt network services to particular transmission characteristics, and requirements to manage network resources to achieve particular response characteristics for an aggregated collection of traffic.

Approaches to QoS

The relevant efforts within the Internet Engineering Task Force (IETF) attempted to address the obvious need for interoperable technical standards for QoS mechanisms within the network that would support some form of true application level end-to-end distinguished service on the Internet.

The initial approach to QoS was that of the "Integrated Services" architecture (Intserv). This approach focuses on the application as the trigger for QoS. Here, the application first signals its service requirements to the network in the form of a reservation request, and the network responds to this request. The application proceeds only if the network has indicated that it is able to carry the additional load at the requested service level by committing to the reservation. The reservation remains in force until the application explicitly requests termination of the reservation, or the network signals to the application that it is unable to continue the reservation. The essential feature of this model is the "all-or-nothing" nature of the service model. Either the network commits to the reservation, in which case the application does not have to monitor the level of network response to the service, or the network indicates that it cannot meet the reservation. This approach imposes per-application state within the network, and for large-scale networks, such as the global Internet itself, this approach alone is simply not viable.

Taking the cars and roads analogy, Intserv is a bit like trying to construct a new lane in the road system every time a premium service passenger wants to take a journey. It's an expensive undertaking that simply cannot scale!

So the IETF tried again, and this time it looked at the core of the network, and examined those mechanisms that could provide differentiated service outcomes with appropriate scaling properties. This approach, the "Differentiated Services" architecture (Diffserv), included dropping the concept of a per-application path state across the network and using instead the concept of aggregated service mechanisms. Within the aggregated service model, the network provides a smaller number of different service classes and aggregates similar service demands from a set of applications into a single service class. Aggregated services are typically seen as an entry filter, where on entry to the network each packet is classified into a particular service profile. This classification is carried within the IP packet header, using 6 bits from the deprecated IP Type of Service (TOS) header to carry the service coding. The network then uses this service code in the packet header to treat this packet identically to all other packets that carry the same service code. While this approach does possess some possibility to scale across the entire Internet, there are numerous unresolved issues relating to the quality signalling between individual applications and the network. There is no uniform definition of the aggregated services, nor any particular level of assurance that an individual instance of a flow within a particular service class would receive any particular service response, as DiffServe deals in aggregate outcomes, not outcomes on a per individual packet flow basis. This aggregated service model does not allow an individual application to sense if it is receiving the necessary service response from the network.

Again, taking the cars and roads analogy, Diffserv is a bit like adding an express lane to the freeway. When the freeway is lightly loaded the express lane is no faster than any other lane., And when the freeway is completely congested the express lane suffers precisely the same fate.

QoS Deployment - or Not!

Neither approach alone is adequate to meet the QoS requirements. The Integrated Services approach imposes an excessive load in the core of large networks through the imposition of a per-application path state. The Differentiated Services approach provides superior scaling properties through the use of aggregated service elements, but includes no concept of control signalling to inform the traffic conditioning elements of the current state of the network, or the current per-application requirements.

Is either of these services, or a combination of these two approaches, enough to motivate widespread QoS deployment on the Internet?

For the past decade the response has been a clear “No.”

Perhaps this strong negative response should be further qualified. The existing tools are insufficient to support widespread use of QoS-based services on the multi-provider public Internet. The qualification is that within the enterprise network environment there are much stronger drivers for QoS mechanisms and much greater levels of administrative control over the overall network architecture, while within the multi-provider public Internet, these drivers are not apparent. The enterprise approach may also have some parallels within a single IP carrier’s network, or even across some forms of bilateral agreements between carriers. However, such approaches are not a sufficiently widespread feature of the public Internet service environment.

Let’s look more closely at the public Internet and QoS to see why there is a mismatch between the two. The major stumbling blocks in attempting to address how QoS could be deployed in the public Internet are both engineering and economic in nature.

From an engineering perspective, we need to remember that in order to actually deliver any reasonable assurance of a quality-differentiated service, the service-quality mechanism chosen must be deployed across all networks along the end-to-end paths of the quality-service traffic. In a heterogeneous multi-provider environment such as the public Internet, this outcome is very unlikely. Within the tens of thousands of component service providers that make up the global Internet, such uniformity of action is highly improbable. The IPv6 transition structure correctly identifies the first step as isolated “islands” of IPv6 functionality, interconnected by some form of IPv6 “bridges.” While the potential scenario of initial QoS deployment may be similar, in terms of isolated islands of deployment of QoS services, there is a much stricter requirement for the “bridges” across the non-QoS-aware parts of the network; namely, that they do not distort the service outcomes. In effect, this scenario requires a QoS response from a non-QoS system. This is obviously a major impediment for QoS deployment.

The engineering issues are deeper than simply the considerations of transition within a potential deployment scenario. The issues include:

- The need for QoS-enabled applications that can predict their service requirements in advance, and be able to signal these requirements into the network.
- In the case of the differentiated service approach of admission controls, there is a requirement for the interior of the network to be able to signal current load conditions to the network admission systems.
- This architecture also requires that the admission control points be able to use admission-decision support systems in order to include consideration of the service load, the current network load, and the policy parameters of the network that may allow some level of pre-emption of various admission decisions in order to meet high-priority service requirements.
- The signalling and negotiation aspect of QoS extends into the inter-domain space, where two or more service providers need to negotiate mutually acceptable service profiles, and associated service access. This extends beyond the addition of bilateral agreements and encompasses the requirement to add QoS attributes to inter-domain routing protocols. The tools and operating techniques required to support this functionality remain poorly defined.
- Measurement of service performance remains an area in which existing measurement tools are lacking. While it is possible to instrument every active device within a network into a network

management system, such an element-by-element view does not readily translate to the end-to-end view of application service performance.

From an economic perspective, we must remember that no Internet retail tariff structure includes a concept of end-to-end tariffed transactions. All tariffs are access based, because application transactions are not readily visible to the Internet network. It should come as no surprise therefore to observe that no financially stable structure of inter-provider interconnection transactional-based financial settlements has existed on the Internet.

However, end-to-end QoS transactions demand a different economic model to that used in today's Internet. The initiator of the end-to-end QoS transaction is in effect electing to generate an end-to-end service profile that pre-emptly network resources that would otherwise be distributed to other users of the network. If such a profile is requested, the initiator should pay the initiating provider a retail tariff to cover the entire end-to-end cost of maintaining the network state that would support this service profile, and the initiating provider must then indicate a willingness to financially settle with other transit networks that lie on the end-to-end path, in order for these transit networks to also devote network resources to service the traffic associated with this transaction. The arbitrary nature of the Internet transits, the dynamic nature of inter-domain routing, and the lack of transaction setups in any scalable form to support QoS mechanisms make this entire scenario highly improbable within our current understanding of inter-provider connections and inter-provider policy-management mechanisms in the Internet.

The coordination structure of the public Internet would have to change from the state we have today if we want to use QoS based services. The necessary changes include:

- A common selection of a set of QoS mechanisms to deploy,
- A different packet forwarding mechanism that relies on inter-AS path pinning,
- An altered inter-AS routing environment that was QoS enabled in some manner,
- Ubiquitous deployment of these mechanisms across both service provider and client networks,
- The adoption of a uniform set of retail tariffs for QoS services,
- The definition and common acceptance of multi-party QoS-related financial settlements that support fair and equitable cost distribution among multiple providers, and
- The definition of commonly accepted service performance metrics and related measurement methodologies to allow end-to-end and network-by-network service outcomes to be objectively assessed.

This is a significant agenda for the industry at large to undertake, and more so in an environment that features diversity and vigorous competition between various public Internet service providers.

An additional factor is also working against QoS deployment in the public Internet space. The sustained delivery of increasing network capacity and an associated dropping of unit costs continues in bringing network carriage capacity to the level of an abundant commodity across large parts of the Internet world. Over the past 10 years we've moved from transmission capacities of hundreds of megabits to gigabits and we are now heading into multi-terabit capacities. As the unit costs of network capacity decline in the face of increasing levels of capacity of transmission systems, the market niche that QoS could occupy in managing a scarce resource shrinks, rather than grows.

The driver for QoS deployment is not that the best-effort service is not good enough. The problem that QoS is attempting to address is one of allocation of network capacity at those points in time when the network is under sustained load, or, in other words, taking on the task of rationing capacity when there is not enough network capacity to meet the cumulative total of every demand that is being expressed at the time. When a network is under load, the QoS response is to place additional control functionality in both applications and in the network to manage this rationing function so that the degradation of the delivered service is selective. Obviously such an activity imposes additional costs on the network operators and the network clients. Such additional costs have not created any additional network capacity. The total sum of demand remains in excess of capacity after the deployment of QoS mechanisms.

The alternative approach is to incur additional cost by augmenting the capacity of the network. This approach also minimizes the impact of load on the network, but by increasing capacity it also increases the capability of the network infrastructure to support more customers and generate higher revenues. This approach also imposes additional costs onto the network, but in an environment of abundant transmission capacity, our experience so far, in more than twenty years of scaling the Internet, is that volume economics works. Bigger has continued to be cheaper, as larger networks drive the transmission network's unit cost base down, and this in turn raises the pressure on competitors to also scale up their networks to access similarly lowered price points in the market. The alternative form of management of excess demand, that of rationing, even when described is such a beguiling term as "Quality of Service" is not a commercially viable alternative in a highly competitive environment of competing carriage services. Rationing attempts to create higher price points for a service level when the competition is achieving at lower price points. In terms of game theory, spending money on QoS is a losing response when your competitor is spending their money on augmenting capacity.

Where does this leave QoS and the public Internet? In asking for QoS to be deployed within the existing incarnation of the public multi-provider Internet, we are simply asking for the wrong thing. When we contemplate what the Internet will need for tomorrow's uses then everybody's interests appear to be served through a path of adding more capacity, rather than contemplating various forms of rationing and enforced scarcity in the underlying fabric of the network.

The QoS Emperor's Wardrobe

If all the above is true, then why do we continually see QoS pedalled in today's Internet? If QoS truly is a dead end, then it's a very persistent dead end that that won't simply go away and die! Why is QoS such a persistent meme in the Internet?

To summarize much of the preceding material, QoS has major barriers in the world of the public Internet, including:

- no convergence on a single QoS technology
- no understanding of network-to-network signalling technologies
- no clear concept of feedback control
- no clear concept of application signalling
- no uniformity in network architectures
- high degree of complexity and operational cost
- the continuing existence of cheaper alternatives

If QoS is just dysfunctional networking snake oil then why have vendors added QoS capabilities into their function set? Why do we see respected standards bodies push concepts such as "Next Generation Networking," which attempt to integrate QoS into their basic architectural models? Are these folk all deluded? Or are they following a strange set of market signals that lead them in another direction?

I suspect that its again an outcome of market forces and the desire to differentiate. In a competitive deregulated market individual producers thrive when they deliver what their customers want to pay for. So if a customer says "I want QoS" then the best answer that a vendor could give is "You will find all the QoS knobs and levels you would ever want in our latest product." The vendor could've said "Well, for your environment QoS is all nonsense - save your money and buy a cheaper router. The QoS knobs and levers are all just cosmetic in your case, and they probably would do more damage than good." But if a vendor responded in that way the customer would simply head to a competitor and buy it anyway. There is a certain amount of wish fulfilment in buying a router as much as there is a certain amount of wish fulfilment in buying many consumer products, as any retailer would tell you! So vendors perpetuate the myth because customers want to believe. But this is circular. Customers see vendors selling QoS and believe the vendor knows something that they don't. Collectively this circle of desire and fulfilment creates an aura of plausibility that sucks in more customers, which further reinforces the myth, which creates more demand, and so on.

At this point we see others enter in to the field and their starting assumption becomes that QoS is real - that it exists and that it works. Aside from the somewhat breathtaking levels of credulity this exposes, this now gets difficult. Debunking this myth now involves taking on an entire industry. Vendors really have no interest in saying out loud: "Well you are right, we knew it was inappropriate to your needs. But you wanted to pay for it so we sold you exactly what you said you wanted to pay for!" All those parts of the industry that also believed the myth have absolutely no interest in saying out loud: "Well, actually we were wrong. We just believed these other folk because they looked like they were the experts!"

Now the problem is of similar dimensions to that of the Emperor's new clothes. A significant part of this industry now has invested a certain amount of their reputation and expertise into what is, for the public Internet, nothing more than a myth. And if this myth is commonly accepted to be just that, then these same folk have to manage a credibility problem. How could a group of network operators, a group who supposedly use this equipment every day, who rely on it as the basis of their business activity - how could they be so trusting? Are they really at the point they simply accept everything their vendors tell them as the literal truth? And even if the answer to that question is "yes", I'd guess that noone really wants to hear that answer!

So debunking this kind of embedded myth is hard. Some carriage operators, particularly those with a strong heritage from the former telephone industry still fondly cherish the hope that QoS is a way of restoring their lost fortunes and lost powers; they crave for a way of fishing out those golden voice packets from an ocean of worthless bit-torrent, streaming video, youtube and web surfing packet dross, and charging both ends of this voice conversation by the minute for treating these golden packets with loving care and tender attention - first class treatment, if you will. Debunking the QoS myth on the public Internet necessarily entails debunking some widely held and fondly nurtured myths within the carriage sector of this industry, and that's a message that has some uncomfortable undertones. We would be trampling all over these precious hopes for a brighter future role for some carriage providers in a QoS-enabled Internet. And that wouldn't be appreciated!

For the public Internet, the QoS Emperor's wardrobe is indeed completely void.

But saying so out loud is not an easy task.

Disclaimer

The views expressed are the author's and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.

About the Author

Geoff Huston B.Sc., M.Sc., has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and has been active in the Internet Engineering Task Force for many years.

www.potaroo.net

The QoS Emperor's Wardrobe In the process of revising the International Telecommunications Regulations, it has been suggested to base a new IP interconnection framework on the concept of end-to-end Quality of Service. This raises a question in my head: How "real" is end-to-end Quality of Service in the Internet? <https://labs.ripe.net/Members/gih/the-qos-emperors-wardrobe>
<https://labs.ripe.net/logo.png>. The QoS Emperor's Wardrobe. Geoff Huston â€” 24 Jun 2012. Tweet.