QUALITY OF SERVICE IN THE INTERNET - THE FUTURE

Olev Kartau
IP and Access Solutions
Nokia Telecommunications
P.O. Box 370, NOKIA GROUP FIN-00045, Finland
Email: olev.kartau@nokia.com

April 30, 1999

Abstract

The Internet is overloaded and some users see the deployment of QoS as one of the ways to proceed. However, the QoS is not there yet. There are several research directions about technologies that may, and eventually will provide some kind of QoS in the future networks. *QoS Routing* tries to integrate QoS features into the network layer routing. *Multi Protocol Label Switching* is an integrated Layer2/Layer3 forwarding paradigm. *RSVP extensions* allow integration of RSVP and routing. The use of QoS mechanisms for congestion pricing are possible in the future, and potential coming of IPv6 QoS, also end-to-end QoS are topics for discussion.

1 Introduction

Internet business has different nature than more "traditional" ones, being studied for long time. For example, Internet Service Providers in the U.S. face churn rates that are five times greater than those of "mature" telecommunications services such as cellular phones, pagers, and long distance carriers. Ten percent of all subscribers cancel their services each month, usually frustrated by busy signals and slow connections, e.g. bad QoS. Nearly two-thirds that cancel are immediately signing up with different ISPs [1]. Provider change seems to be one of the customer QoS demand & feedback methods.
This paper describes some current problems with QoS in the Internet and discusses briefly what are the potential development directions in the near future.

2 Future visions

2.1 The near future of the Internet [2]

- the traffic amount will increase by 100% ...1000% annually
- services built on top of IP will be the dominant platform
- electronic commerce will grow significantly
- no great breakthroughs from initiatives like Internet2
- most end users continue to World Wide Wait, (slow modems)
- ATM switches will be the worldwide backbone standard, with speeds up to a terabit
- packet-switching replaces circuit switching - all traffic packetized by the year 2010
- telcos are going to lose their private data and voice business to the Internet
- some "traditional" concepts will be replaced (new: PC-based PBX, PC-based phone, etc.)

2.2 Next Generation Internet

- "Total Network Immersion", with people continuously on-line and equipped with smart sensors
- bodynets (hardwear and softwear)
- bio-electronics with computer interfaces
- "Interplanetary Internet" from NASAs Jet Propulsion Laboratory described in [3] will use an Interplanetary channel protocol based on CCSDS (Consultative Committee for Space Data Systems) recommendations. The first Interplanetary Internet building block launch is expected in the year 2003 [4].

3 Current problems

APC (The Association for Progressive Communications) Magazine asked several of the respected internetworking authorities

- Robert Redford, Cisco network services
- Gordon Cook, editor of the commercial newsletter the Cook Report on Internet
- Larry Roberts, the father of ARPANet
- Craig Partridge, developing multigigabit routers at BBN
- Vinton Cerf, TCP/IP's inventor, now at MCI
- Bruce Nelson, Cisco chief science officer
- Tom Lyon, founder of Ipsilon Networks
- Martin McNealis, Cisco IOS division

to diagnose current problems and prescribe "a diet for healthy growth". The full article is in [2].

3.1 Congestion

The congestion in the Internet is familiar to all - overloaded routers, saturated bandwidth, busy LANs and ISP access "busy" signals. The late movement towards flat-rate pricing has accelerated the traffic growth even more. Adding extra bandwidth typically offers only temporary relief. A relatively new trend - encrypted VPNs tunnelled across the public Internet instead of private leased lines - is likely to make the congestion problem even worse. Experts do not expect better times - there will always be congestion in the network. In the future, ISPs will start offering premium IP services to enterprise users as a way of funding bandwidth expansion.

3.2 Bandwidth bottlenecks

Capacity of fibre is not limitless, but WDM (Wave Division Multiplexing) promises enormous increases in throughput across existing lines. In the land the fibre links will not be a problem for long time, but the trans-oceanic links appear to be more difficult to upgrade to WDM. The fibre access regulations may become an issue. Lot of bandwith is wasted because of same data traversing same links many times. Improved Web caching products could reduce that tendency. The Web cache in the network will be an advantage if compared to a workstation off the side in a slow-speed way and needing frequent (re)configuration.

3.3 Delay growing

Delay caused by the flow control mechanism is what some experts think as the biggest unsolved issue. With current TCP flow control, delay is growing faster than the growth of the network. TCP starts transmitting slowly, then waits for feedback from the network before it can speed up. Reaching "flow" state takes around five round-trips. As the number of router hops has increased - now typically 15 to 20 hops - TCP is slow to reach faster transmission rate. And during congestion, TCP backlogs network nodes with data until it can slow down. One solution is for the network to cooperate in the process. The switches could examine their current capacity, and demand and notify users before they transmit, reducing the cycle of interdependence between round-trip feedback and buffer overload. Cerf has worked with Roberts to explore proactive notification schemes. He foresees two uncertain aspects: notifying the TCP protocol of the network status, and coping with heterogeneity if proactive notification is not implemented throughout the entire Internet. Weighted Random Early Detection (WRED) is a way to deliver classes of service without changing the TCP protocol. WRED begins dropping lower class packets first. Cerf: Weighted RED is a part of a solution, but by no means all of it, because it involves discarding traffic what is less attractive than controlling the rate at which traffic arrives in the network in the first place.

3.4 Routers as bottlenecks

Routers are overloaded with allocation work trying to share resources without a functional QoS protocol and without good flow control. The relationship between development of routers and bandwidth is cyclical. The experts predict that the routing market will bifurcate. At one end will be small, inexpensive LAN routers where cost is the major factor, while backbone operators will use larger and more expensive routers. Crossing of layers 2 and 3 may be hard to continue in the future because it can not tolerate changes in protocols. Switching will become more important, but there will be tensions between packet-switching and virtual circuit approaches (opinions differed already back in 1973 when Roberts started X.25 and Cerf promoted TCP/IP). ATM needs improvement: VCs have to be set up much faster than they are.

3.5 Quality of Service - how ?

Most experts agree that something QoS -related will be implemented to some degree in the near future, but don't agree how. QoS will happen because enterprise customers will demand it. QoS may become an important business driver for ISPs, who have opportunity to develop better services at higher margins. Initially there will be two rates: cheap, often flat-rate consumer access; and premium business service. It is important that enterprises pay for better service in the public internet, instead of building private Inter-

nets, because then also basic service users can benefit from the investments made by big players.
All difficulties are not technical. Developing service-level agreements (SLAs) among ISPs and telcos requires complex negotiation, and will take time. Cross-provider, cross-nation coordination is going to require much effort and time. The telephone industry, with a ITU coordination body took a long time to develop universal interconnect structures, and there's no similar international agency for the Internet. The "SLA game" may exclude smaller ISPs because they don't have the function and performance to play that game.
There are several models that technically help to implement QoS features, from which Cisco's Tag-switching, precedence bits in IP headers, RSVP and ATM are the most widely known for now. See [5] for a broader outline.
Experts foresee that QoS will develop gradually, but once started, may spread quickly. It will probably start with some core services deployed by the "big guys", with QoS at the edge and classes of service in the core. One possible implementation would be with priority queues and SNMP-managed channels, with multilevel queuing within the routers. If there's a significant competitive advantage for one of the big players delivering QoS, then others will follow (a "domino effect").
Some experts say that RSVP is "a dead horse" and there is no replacement. Both RSVP and ATM signalling are much too slow for either Web access or voice, and RSVP imposes too much overhead on switch processors, which won't have enough processing time to set up enough connections to use the bandwidth they have to offer. RSVP is not scalable. RSVP may be usable on private IP networks, for instance when delivering video.

4 QoS issues

4.1 The best environment for QoS

The best environment for QoS, the most common denominator, is the IP protocol suite. There is no ubiquitous single-transport technology for telecommunications industry, and trying to build network with a pervasive link-layer technology (like ATM) is successfull only in smaller scale.

4.2 QoS Routing

QoS Routing (QoSR) is a routing mechanism where flow paths are determined based on some knowledge of network resources availability and QoS requirements of flows [6]. QoS Routing is considered one of the missing pieces in the puzzle of delivering QoS in data networks. QoSR presents several problems :

- determining whether the QoS requirements for a flow can be accommodated on a particular link or a path
- QoSR dependency on the stability of underlying routing infrastructure
- RSVP and routing cooperation efficiency concerns, especially when end-to-end paths change
- QoSR scalability concerns in big networks - path computation is difficult

A modified link-state routing protocol would appear to be an ideal candidate for *intradomain* QoSR, but the additional route information propagation overhead may make it very inefficient. It is hard to provide scalable QoSR, and reducing the number of peers and aggregation as solutions to scalability problems would bring also new problems. Proposals [7],[8] suggest QoS extensions to the OSPF routing protocol.
For *interdomain* QoS Routing, the link-state protocol is not the best choice, and it is unclear what would provide dynamic interdomain QoS Routing [9].
Knowledge of the stability and scalability of large-scale routing is inadequate even in the existing network, so speaking of QoS Routing for large scale is highly speculative. As with dancing circus bear, the fact that it can dance is simply amazing. Getting it to sing is possibly asking too much.

4.3 QoS Routing and RSVP

Currently, there is a noticeable disconnect between RSVP and underlying routing. At least one IETF draft proposes a method for RSVP to request information and services from the local routing process - RSRR interface (Routing Support for Resource Reservation) [10]. RSRR works via exchange of asynchronous query and reply messages. Another IETF draft [11] describes extensions to RSRR, proposing support for explicit paths and QoS routing. Still, given the importance of QoS path management, relying entirely on the dynamics of an underlying routing protocol can be harmful, for example because of risk of routing and RSVP interfering each other. It is unclear whether mechanisms for RSVP and routing cooperation will be adopted or not, and if yes, then how broadly.

4.4 QoS and MPLS

Multi Protocol Label Switching, sometimes referred as "layer 2.5", tries to blend the best of similar concepts into a standardized framework and protocol suite, thus improving the performance, scalability and flexibility of routing services. Future possibilities regarding MPLS and QoS are:

- direct mapping of IP precedence field (3 bits) to CoS field in the label. (this is Cisco's contribution to label switching)
- capability to build explicit label-switched paths, using MPLS traffic engineering tools and download of paths to network devices.
- possibility to integrate MPLS traffic engineering with RSVP

4.5 QoS and IP Multicast

Currently, the capability to impose any form of QoS structure on multicast traffic is not well understood, and the IP multicast QoS subject is only speculative [9]. Experts say that a much better multicast mechanism than either ATM or IP have is needed, providing complete flow control and complete QoS across many-to-many multicast, with the ability to add and drop flexibly.

4.6 QoS as a Congestion Pricing Algorithm

It has been suggested that pricing mechanisms are a viable method of relieving congestion. One model called *Network-Activated QoS* works like a performance insurance - traffic matching QoS criteria is marked as being of an elevated precedence at ingress device, and will get premium service associated with pricing. It is interesting that in the absence of congestion, also any other traffic will see similar performance, so it is unclear how this model could attract users that already are paying for normal best-effort service.
Another model is *Customer-Activated Precedence Traffic* where customer marks the traffic instead. In case of performance degradation, the customer can increase the precedence level of transmitted packets, which is associated with an incremental cost.
In transition to a QoS-defined service, there will probably be a number of increasingly sophisticated implementation phases for QoS congestion pricing.

4.7 Interprovider structures and end-to-end QoS

QoS is going to be deployed in two different environments - the private networks with particular objectives, known applications, and concrete business needs, and the public Internet with ISPs trying to sell competitive service. The public-side deployment will be challenging, because of the need to interpret QoS related traffic uniformly. It is realistic to predict that QoS will not be deployed uniformly on the public Internet, resulting in parts of network that ignore QoS signalling, generating an issue for stateless QoS environment.
Current vision is that local QoS structures will appear within the local Internet with certain number of participating ISPs. End-to-end QoS is then achievable when both ends are attached to same provider (or group of providers). We will see changes in both interprovider and end user agreements. Does this "bottom-up" growth model result in truly uniform end-to-end QoS deployment, is a subject of speculation, the big issues being not only technical, but also commercial, business practices, policies etc.

4.8 Bidirectional QoS

QoS mechanisms discussed this far are unidirectional. The question is whether this is adequate, or is bidirectionality necessary in QoS picture. This bidirectionality issue is relevant only to end-to-end controlled traffic flows, and most readily is applicable to TCP traffic. One example is an issue of ACK packets taking different path than the actual data flow.

4.9 QoS and IPv6

It is a common misperception that the IPv6 protocol somehow includes a magic knob providing QoS support. Two components of the IPv6 may provide a method to deliver Classes of Service - the 4-bit *Priority field* and the 24-bit *Flow Label*. Of those, the priority field does not offer any substantial improvement over the IP precedence field of the IPv4. It offers eight more levels of distinction, but that is not of benefit - where no one has implemented even a simplest two-level distinction of service classes, increasing the possible number of service classes is not a reason to prefer IPv6 over IPv4. A rough consensus of polled service providers indicates the need for **no more than three** service levels (for matters like configuration simplicity). The *Flow Label* may be more useful, but the only immediate use for it is in conjunction with a RSVP, to associate a flow with a reservation. Currently it has however no additional value, because recently devised methods are available to build flow state in routers without the need for a label in each packet. As for conclusion, IPv6 does not offer any substantial QoS benefits above and beyond what already is achievable with IPv4.

5 Application examples

Here are some examples of trends that are happening right now.

5.1 Policy-based networks

On Apr 26 1999 Lucent rolled out the first of a series of software applications called *RealNet Rules* that will allow administrators to configure policies on user access, application priority, and performance. The system uses LDAP to distribute policy information from a central directory to the devices across a network, and has price as $10,000 per server. Many other networking vendors (including Extreme Networks, Cisco, Nortel Networks, and IBM) plan to introduce or enhance policy-based networking software during next months.
"Lucent, Extreme set to unveil policy networks" Article By Stephen Lawson, Apr 23, 1999, InfoWorld Electric (http://www.infoworld.com).

5.2 QoS and the applications

Cisco and Oracle provide templates and guidelines for administrators to set up *CiscoAssure Policy Networking software for Oracle Applications.* Using the templates, administrators can give networks the intelligence to recognize Oracle applications, and use priorities when dealing with such traffic. This is one example how policy-based networking systems will allow to favor some users and applications over others on their networks. Similar deal already existed between Cisco and PeopleSoft.
"Cisco Systems and Oracle this week will team up to help IT managers prioritize Oracle applications on Cisco networks" Article By Stephen Lawson, Apr 15, 1999, InfoWorld Electric (http://www.infoworld.com).

6 Research activities

6.1 The NASA Research & Education Network (NREN)

NREN (http://www.nren.nasa.gov/eng.html) is NASA's project in developing the Next Generation Internet (NGI). NREN has plans to demonstrate the effectiveness of QoS over backbone and "end-to-end QoS for NASA applications", exploring both router-based and switch-based capabilities. They also have goals to demonstrate the effectiveness of RSVP over the backbone and for NASA applications, and to test classes of service in its network. The NREN/NGIX-West has been accepted as a backbone network for the Internet2 Quality of Service testbed (QBone, www.internet2.edu/qos/qbone/), an interdomain testbed for DiffServ.

6.2 The Energy Sciences Network (ESN)

The ESnet (http://www.es.net) QoS project has two main objectives:

- to supply useful, deployable, flexible QoS to the Energy Research community as an ESnet service
- to demonstrate how administratively heterogeneous QoS might be implemented and deployed

They claim that these objectives will be reached in two steps:

- using class-based queuing (CBQ) machinery in ESnet Cisco routers and selected site routers.
- developing a "bandwidth broker application" to control access to the CBQ machinery

6.3 Stanford Research Institute (SRI)

SRI has a project called "QoS Middleware For Group Applications" (http://www.csl.sri.com/~shacham/Intro.html), where they try to define QoS for a multi-party session with heterogeneous participants. They have developed the concept of heterogeneous multicast to overcome the problem of a common denominator. In this concept the group elements participate in the session based on their individual capabilities, without having to suffer the effect of other participants' limitations. For example, in a video teleconferencing session the video streams are encoded hierarchically so that they can be viewed at different fidelity levels. SRI current implementation of heterogeneous multicast builds the mechanisms for QoS specifications and handling into the applications. As the next step, they are developing middleware and protocols for performing these tasks on behalf of many different applications. SRI and Sun Microsystems will be implementing the software layer called the Middleware Daemon (MD) in Java.

6.4 Internet2

Internet2 QoS (http://www.internet2.edu/qos) researchers have set some not-so-humble goals, looking for end-to-end, wide-area QoS functionality for inter-institutional links running across separately managed network clouds. They discuss future QoS requirements in [12].

6.5 SIMA

Simple Integrated Media Access (SIMA) is a way to provide QoS using Differentiated Services. See http://www-nrc.nokia.com/sima/index.html.

6.6 DTM

Dynamic Transfer Mode (DTM) is a network protocol developed by Swedish company NetInsight AB (http://www.netinsight.se).
DTM is for high speed networking for dynamic transport of integrated traffic. It is a transport network architecture based on circuit-switching augmented with dynamic reallocation of bandwidth.
DTM includes switching and a signalling protocol and can thus, in contrast to e.g., SDH/SONET, set up multi-rate channels (circuits) on demand, and the capacity of a channel can be changed according to traffic characteristics during operation. Additionally, resources can be reallocated between nodes according to the current demands. In this way, free bandwidth is allocated to nodes with highest demands, providing an autonomous and very efficient dynamic infrastructure.

7 QoS observations

7.1 There is no such thing as a Free Lunch

Ongoing efforts to provide "perfect" solutions illustrate that attempts to solve all possible problems result in technologies that are too complex, poorly scalable, or do not integrate into the diversity of the Internet.

7.2 Marketing matters

The current marketing model in the Internet is an undifferentiated one, and the current best-effort model is remarkably simple, and used uniformly. The marketing questions are:

- is there another market for a different model within the same environment ?
- is QoS finacially viable ?
- does user see the difference ? (does not in an uncongested network)
- subscription service or a user-spcified per-packet option ?
- effects on others (non-QoS) users ?

7.3 Economics matters

Having QoS means making a set of compromises, some of which are economic, others are technical. There must be a balance between network engineering, network design, and scales economy. Big companies getting into ISP business face a paradigm shift from "planned capacity matches planned demand" model to "all you can eat" model. In such environment, regardless of the levels of installed bandwith, the capacity is not readily available ! The question is:
Is the QoS needed because of fundamental economic imbalances in the evolution of telecommunications industry ?

8 Conclusion

QoS is possible in the Internet, but it does come at a price of compromise - there are no perfect solutions. Some expectations will probably not be fulfilled, since guarantees are simply not possible in the Internet, at least not for the foreseeable future. What is possible, however, is delivering differentiated levels of best effort traffic in a manner which is predictable and fairly consistent.

References

[1] Jim Hu. *ISPs face high churn rates.* CNET News.com February 5, 1998, http://www.news.com/News/Item/0,4,18879,00.html.

[2] Dan Tebbutt. *Internet 2000: Highway to the horizon.* Oct. 97 apcmag.com article, http://apcmag.com.

[3] Vinton Cerf. *What is the Internet?* Slides from a talk in UNC, Aug. 31, 1998, http://ils.unc.edu/ils/releases/cerf_talk/.

[4] Vinton Cerf. *A talk at the Telecom Israel 98.* Review by Claus Descamps and Adrian Rawlings, http://www.de.infowin.org/ACTS/IENM/NEWSCLIPS/arch1998/981198no.htm.

[5] Cisco. *Advanced QoS Services for the Intelligent Internet.* http://www.cisco.com/warp/public/732/net_enabled/qos_wp.htm.

[6] E. Crawley et al. *A Framework for QoS-based Routing in the Internet.* IETF Internet Draft, draft-ietf-qosr-framework-01.txt, July 1997.

[7] Z. Zhang et al. *QoS Extension to OSPF.* IETF Internet Draft, draft-zhang-qos-ospf-00.txt, June 1996.

[8] R. Guerin et al. *QoS Routing Mechanisms and OSPF Extensions.* IETF Internet Draft, draft-guerin-routing-ospf-01.txt, June 1996.

[9] G. Ferguson, P. Huston. *Quality of Service*. Wiley Computer Publishing, New York, 1998.

[10] D. Zappala. *RSRR: A Routing Interface for RSVP*. IETF Internet Draft, draft-ietf-rsvp-routing-01.txt, November 1996.

[11] R. Guerin et al. *Extended RSVP-Routing Interface*. IETF Internet Draft, draft-guerin-ext-rsvp-routing-untf-00.txt, July 1997.

[12] Ben Teitelbaum and Ted Hanss. *QoS Requirements for Internet2*. Draft, http://www.internet2.edu/qos/may98Workshop/html/requirements.html.