Internet2 and Quality of Service: Research, Experience, and Conclusions

Steven C. Corbató
Managing Director for Technology Direction & Development
Internet2

Ben Teitelbaum[1]
Program Manager
Internet2

May 2006

The Internet2 community has been deploying high-performance Internet networks for nearly ten years. Today the networks operated by the Internet2 community—including campus, regional, and nationwide networks—connect over 5 million individuals at more than 240 research and education institutions. These same networks interconnect with research networks in dozens of other countries and to broader education networks in 35 states. In addition to enabling the research and education community to conduct their work in ways not possible on the commercial Internet, a notable feature of this large-scale network infrastructure is that it provides an ideal test bed for new Internet technologies, services, and applications.

As many of these networks began to be deployed in 1997, a large portion of the Internet2 technical community assumed that implementing Quality of Service (QoS) or "premium services" would be a key capability for a high-performance network environment. QoS technologies would prioritize certain applications, such as streaming video or videoconferencing, to assure that their bits traversed the network with the required reliability and fine-tuned performance. Many engineers, network researchers, and application developers believed these applications required QoS because overall demand for limited bandwidth would result in network congestion; absent QoS, the degraded performance resulting from congestion would not allow these applications to work well or at all[2].

Therefore, between May 1998 and October 2001 technical leaders[3] from the Internet2 community worked to develop and deploy an advanced Internet

---

[1] Steven C. Corbató previously managed the nationwide Internet2 Abilene backbone network. Ben Teitelbaum is the former chair of the Internet2 Quality of Service Working Group and led Internet2's QoS testbed known as the QBone.
[2] Real-time video and audio applications, such as videoconferencing do not tolerate packet loss or jitter (variation in delay) well, and while DVD-quality videoconferencing can be achieved with under 10 megabits per second (Mbps) of bandwidth in each direction, today's highest-bandwidth videoconferencing technologies require up to 1,500 Mbps.
[3] The Internet2 QoS Working Group was begun in fall 1997 and recommended in May 1998 that, "Internet2 deploy services that leverage the then newly-emergent differentiated services architecture for IP QoS."

Protocol (IP) service based on QoS technology[4]. Network scientists and engineers from academia and industry together explored various architectures and technical foundations for implementing a "premium" service. These efforts included multiple workshops and design meetings supported by both industry and government research agencies such as the National Science Foundation and the U.S. Department of Energy[5].

The result of this work was a technology specification and a working implementation of that specification that delivered the promised quality of service function across multiple networks. With the participation of several national and regional networks, including the Internet2 Abilene backbone network, and with the support of the U.S. Department of Energy, a nationwide multi-year nationwide QoS test bed was created. Both the sustained test bed, and successful high-profile demonstrations of it, proved that the technology the Internet2-organized effort had developed and tested was both workable and deployable.

However, practical experience during the same timeframe showed a far simpler and far more cost effective means for ensuring high-performance networking: simply provide an overabundance of bandwidth to end users to ensure that the odds of network congestion are minimized. This approach avoided practical deployment obstacles to implementing any effective QoS across a multiple network environment such as the Internet. Specific obstacles include: coordinating upgrades to QoS technology across every network; changing dramatically network operations, peering arrangements, and business models; and developing suitable means to verify QoS service delivery by users, providers, or both.

The "overprovisioning" approach to ensuring network performance has been made possible by new technology that provided geometric increases in networking capacity at rates that matched or exceeded Moore's Law.[6] These capacity increases have been accompanied with much lower price increases for the products providing the same function. Thus, the effective cost per bit for any network segment's bandwidth has dropped dramatically. With these economics in effect, network operators on all scales—national, regional and local—have made independent decisions in the course of normal operations

---

[4] While Internet2's experience is specifically with QBone Premium Service (QPS) or "Premium," an interdomain virtual leased-line IP service built on diff-serv [IETF RFC2475], the conclusions reached apply to any IP quality of service (QoS) architecture offering a service guarantee.

[5] For example, papers presented at the "First Joint Applications/Engineering Quality of Service Workshop" held in Santa Clara, California on 21-22 May 1998 with the supported both by Bay Networks and the National Science Foundation, can be found online at: http://qos.internet2.edu/wg/documents-informational/199805-Proceedings.pdf. Additional documentation of the Internet2 community's work on QoS can be found at: http://qos.internet2.edu/wg/documents.shtml

[6] Moore's Law refers to the observation in 1965 by Gordon E. Moore, co-founder of Intel, that the complexity of integrated circuits doubles every 24 months with improvements in manufacturing methods.

that simply increased link capacity as demand from users for bandwidth has grown.

With increases in capacity at all levels of networking, applications originally assumed to require QoS have been increasingly deployed and used without explicit QoS technology. Even during the period when the QoS development and testing was underway the use of applications such as bandwidth-intensive video over Internet2 networks greatly increased. Users and application developers simply responded to the increasing availability of bandwidth by increasingly developing and using innovative applications, despite having no QoS.

Today, the Internet2 community uses off-the-shelf leading edge technology in its networks and operates on the very same fiber footprint that the commercial Internet uses. The most significant difference between the Internet2 environment and the commercial Internet, including home broadband, is simply connection speed[7]. The slowest network connection between any two desktops within the Internet2 community is typically the 100 megabits per second (Mbps) link between the computer and the local area network and many of these links are moving to a gigabit, or 1000 Mbps[8]. Regional and nationwide networks are increasingly using 10 gigabit per second (10,000 Mbps) technology, with 40 and 100 gigabit per second technologies expected to be available within the next few years.

Internet2 networks today routinely enable applications that require 10 to 100 Mbps of bandwidth or even higher. Scientists—from astronomers to cancer researchers—routinely exchange extremely large datasets in real-time, students participate in remote classes using DVD-quality video-conferencing, and doctors share critical life-saving medical data instantly. Critical to supporting the continued development of networking technologies, Internet2's infrastructure also supports network researchers investigate new Internet technologies, architectures, and protocols.

The Internet2 community's success with the approach of enabling advanced applications by providing ample bandwidth to end users, and its experience with QoS, suggests that rather than introduce additional complexity and additional costs to implement prioritizing techniques, commercial Internet providers should focus on supplying an abundance of bandwidth to end users. A gigabit per second connection to many homes currently served by broadband is technically feasible, and would allow consumers to connect computers and other devices with off-the-shelf equipment for less than $15

---

[7] Importantly, there continues to be considerable effort, such as the Internet2 End-to-End Performance Initiative, to ensure the various network components work well together so that users can make the most of nominal network speeds. Other capabilities that distinguish the Internet2 infrastructure from the commercial Internet include IP multicast and IPv6 distinguish the Internet2 infrastructure from
[8] One example of this is the University of Wisconsin-Madison, which recently deployed gigabit-per-second class network across its campus. Additional details can be found at http://www.doit.wisc.edu/network/upgrade/

per endpoint. This is the approach being successfully implemented in countries such as Korea, Japan and elsewhere.

Instead of implementing QoS, simply increasing network speeds leverages the decreasing cost-per-bit trend of new networking technologies and avoids the pitfalls QoS implementation. The elegant simplicity of the best-effort service model provided by IP is one of the essential reasons for the success of the Internet. Together with the inherent strengths of connectionless networking and the IP's end-to-end design principle, the best-effort service model has enabled a fast, dumb, cheap, and wildly scalable Internet which has, in turn, provided a foundation for manifold innovative uses, unconstrained by a centralized view of how the network can or should be used.