Technical Update

Bill Wyman
University of Colorado at Boulder

"Internet 2 is about the future of the Internet and of digital global telecommunications."

Internet 2 is about the future of the Internet and of digital global telecommunications. It's hard to write about the future without becoming nostalgic. So, I won't try. My first article for the IALL Journal, which had been cajoled out of me by my late friend and colleague, Marie Sheppard, appeared in the winter of 1995 subtitled "A Report from Behind the Front Lines," the piece mapped out the panoply of the pre-web Internet, a medium then consisting more of potential and possibility than application and actuality. The World Wide web was still in gestation. While it is true that the Uniform (or Universal) Resource Locator (URL) and Hypertext Markup Language (HTML) date back as far as 1989, the ascendancy of the web did not begin until a cadre of computer scientists at the University of Illinois' National Center for Supercomputing (NCSA) developed the first successful program to display HTML in a graphical window, dubbed Mosaic. (Talk about nostalgia.) Within a matter of months, Netscape, the web browser, began showing up on campus desktops everywhere. Supplanting any number of would-be competitors, Netscape soon became synonymous with the Internet, the web, and a new vocation cum avocation—surfing. Many of us watched the familiar .edu domain grow like Topsy, but few of us picked up on the what the commercial sector had in store for the Net. By the time the National Science Foundation bowed out of its Internet backbone management role and turned the care and feeding of Internet to an amalgam of private sector and academic interests in January 1995, the benevolent neoplasm of the Internet surrounded us on campus, at the office, and in the schools. The marketing capabilities of the new global multimedium were proving to be irresistible, and the .com domain moved toward center stage to giddy applause.

Your Postal Service. The result: unprecedented growth and traffic. Also slowness, delays, and congestion due to competition for bandwidth. A new word entered our vocabulary—

latency—a technical term for what happens when you click on a web link somewhere, and nothing happens because too much is happening everywhere. What we in higher education had come to regard as our network now seemed to belong to everyone and anyone. Tweek your request, and presently a new package of code makes its way back to you via the Internet as appropriate.

"More bandwidth, yes, but the requirements created by synchronizing the audio and video components of multimedia transmissions, by broadcasting and multicasting, and by the demand for delay-free teleconferencing can only be met by change, innovation, and a gargantuan mound of cash."

Transfer interrupted. If the analogy between the web and the U. S. Postal Service made you shudder just a bit, your experience probably provides good reasons. Some of the same nasty tricks of fate that befall the mail can and do affect hapless web packets all the time. Ever tried to access a page on the web and had the progress bar on your browser stop dead in its tracks? Chances are good your surfing session was interrupted by packet loss. Those imaginary major cities along the Internet are actually data-switching points, and there are only 28 or so of them worldwide. When they get overcrowded, packets get lost or dropped. Think of it as the web equivalent of a "Blue Christmas." Here the Internet improves immeasurably on the post office. Because those packets contain only virtual merchandise, their code contents will be automatically and instantly replicated and re-sent so the goods may get delivered anyway. You can do much the same thing manually by pressing "Stop," followed by "Reload" in your browser.

Gridlock. You've just had a brush with latency. That wasn't so bad, was it? As real world experiences may suggest, however, arithmetic growth in the number of web sites and users begat geometric increases in packet traffic. Automatic re-sending of lost packets worsens the overcrowded conditions that caused the packet to go astray in the first place. Human ingenuity, a factor universally unmanageable, coupled with the profit motive, now gone global, ensured that new web site after new web site, servicing every variation on every imaginable need, would appear by the day and by the hour. We ran out of time to browse all the web had to offer, and we ran out of time to keep track of it. New software robots were ushered in that scoured the Internet ceaselessly and assembled keyword-indexed resource lists millions and millions of items long. The creation of these lists and our accessing them via the web consumed even more bandwidth. Public awareness of network congestion hit a new high when the America On-Line shot itself in both feet at once by offering their subscribers a lower price and unlimited access to the Internet without the means to deliver either successfully. Overnight the problem advanced from being merely inconvenient and unpopular congestion, to risky and expensive gridlock.

Quality of Service. Those celebrated AOL unfortunates and their attorneys are not the only users dissatisfied with the performance of the commercial Internet. New multimedia applications and inchoate digital telecommunication schemes, such as video-over-IP now under development by NYSERNET and Information Technology Services at University of Colorado, show existing and future needs in ever sharper relief. More bandwidth, yes, but the requirements created by synchronizing the audio and video components of multimedia transmissions, by broadcasting and multicasting, and by the demand for delay-free teleconferencing can only be met by change, innovation, and a gargantuan mound of cash. Ironically, two democratic hallmarks of Internet architecture—the stateless or connectionless interface, discussed earlier, and best effort delivery for transmission of packets—have re-emerged as impediments to new applications and actualities. Again, like the Postal Service, anyone can use the Internet who has the price of admission. And it's pretty low. All users get an equal crack at a share of the bandwidth that they can call theirs for one transmission's worth of packets. In return, the system promises to make its best effort to deliver those packets to their destination. They may be delayed and re-sent. They may be scrambled or corrupted. They may be lost or dropped. Recourse for the user is limited to a casual, or heated, or repeated "oh well . . ." A current solution proposes a reservation system in which users would queue up and take turns using a larger share of the bandwidth, and pay for it accordingly. This is part of the billowy politics of Internet 2. The entire issue of who will get what buzzes around the catchword "QoS" or "Quality of Service." Watch for it in your neighborhood soon.

"The entire issue of who will get what buzzes around the catchword 'QoS' or 'Quality of Service.'"

So long Ethernet? Ethernet, the most abundant and cost-effective networking technology available, is likewise democratic. It does not distinguish among packets carrying a distance education course, a teleconference, or a researcher's dataset, and a no-priority printing job, the arrival of the next villain in an networked DOOM game, or that junk email about some pyramid sales scam. The packets sent by whomever happened to be first will monopolize all 10 megabits of available bandwidth on Ethernet, and every other user must wait for an opening. Ethernet packets also vary in length, and this, too, causes delays. Some technologists suggest that QoS will require us to scrap Ethernet, and switch to the competition, such as ATM (Asynchronous Transfer Mode). The huge installed base of Ethernet and the daunting costs of replacing it have fueled innovation, and recently its champions have proposed a

QoS solution consisting of a new Resource Reservation Protocol (RSVP), Real-time Transport Protocol (RTP), and IP Multicast. RSVP can be used to allocate a portion of bandwidth to intensive real-time audio and video applications, such as teleconferencing, while leaving some bandwidth available for regular local area network traffic. RSVP can also call ahead to routers along the Internet and reserve the resources necessary for audio to reach you without dropouts and for video to come through without herky-jerky motion or jitters. RTP works at the application level, instructing software how to deal with latency gracefully. IP multicasting directs programs only to those network users who wish to receive them, rather beaming them to all users whether they are interested or not, thereby conserving network resources. What may become RSVP-RTP, and other QoS schemes, and what role they will play in Internet 2, is open to speculation.

"Some technologists suggest that QoS will require us to scrap Ethernet, and switch to the competition, such as ATM (Asynchronous Transfer Mode)."

Internet 2, the One. Around the time that the Internet went commercial (January 1995), many of its original modelers were still working on college campuses. They were both pleased with the medium's success and also concerned about clamorings from their constituents—many of whom were colleagues working on campus—that for them quality of service had done little but decline in recent years. Academics thrive on information and instrumentation, and in the context of networking, that means rapid access to digital libraries and remote control and monitoring of experimental apparatus whether it is located across town or in another hemisphere. It was also clear that large-scale distance education would require some new type of internet. On October 1, 1996 representatives from 30 public and private universities and colleges convened in Chicago and anted up $25,000 each to form an Internet 2 higher education consortium. Its purpose was to foster high-speed global networking, including new hardware, middleware, and software applications.

Internet 2 (I2), the Other. October 1996 was re-election season, and many smiled when Candidate Bill Clinton announced on the 15th of the month that, once reelected, his administration was prepared to sink $100 million into the "Next Generation Internet" (NGI). The two events were not unconnected, except by verbiage, and the modelers of I2 saw reason to rejoice. As NGI was fleshed out in the environment of Washington at election time, however, its parallels to the Chicago I2 concept seemed to wane. NGI's goals turned out to be in two parts. The first sets out to create a new network about 10 times faster than the current Internet which will connect 100 institutions together. This makes it look a great deal like "the other Internet 2." However, its purpose is to promote experimentation with new applications, rather than with networking technologies.

Internet 3? The second part of the NGI proposal will create yet another network that will be 100 to 1,000 times faster and have only 10 connections. This will be dedicated to networking research rather than for use by researchers who need networking. Real world experience suggests some thorny questions: how will our Federal bureaucracy, which tends to lumber, leap to the task of developing and deploying some of the most ephemeral of technologies, and quickly? How can the United States government, which is forbidden from making industrial policy, persuade the developers and manufacturers of computing and networking technology to do it Our Way? For the moment, forward motion on either part of NGI is not evident.

Internet, the Next Steps. About 100 private and public colleges and universities comprise the present Internet 2 consortium in partnership with private industry and the Federal government. The cornerstone of its program, which is described in detail on a web site at http://www.internet2.edu has been summarized as "to ensure advanced network services are available on interoperable backbone networks that are competitively provided by many vendors," as well as to revitalize the cooperation between higher education, industry and government, including the Next Generation Internet (NGI), that brought us the original Internet. The National Science Foundation (NSF) has again emerged as an important player in several regards. NSF manages the new national vBNS (very high performance Network Backbone Service) and indications are that vBNS will be the initial backbone for Internet2. NSF has also awarded grants to 50 of the I2 member campuses on a competitive basis as seed money for developing new technologies that will be required. I2 will use the familiar TCP/IP (Transmission Control Protocol/Internet Protocol) now standard on our desktops. Users on certain member campuses can expect connectivity before the end of 1997 (some people are connected to and using vBNS already), and plans project its availability to all 100 member campuses by the end of 1998. A QoS scheme should be in place by about the same time. Universal availability is possible if adequate funding can be found here and abroad. I2 will be interconnected to the commercial (a.k.a. commodity) Internet. Campuses unable to join the consortium due to financial constraints will be required to wait until the connectivity and other technologies inherent in I2 are stable, so that cost drop

"The huge installed base of Ethernet and the daunting costs of replacing it have fueled innovation, [such as]...a new Resource Reservation Protocol (RSVP)...."

and they and can be transferred inexpensively. This is much the same process that took place with the NSF-sponsored Internet beginning 15 years ago.

The nature of the working relationship between I2 and NGI is still to be determined. The parallels in their missions are unmistakeable, and each would benefit from as much cooperation as can be mustered. This is the time for some judicious lobbying at the member campus level, coupled with expressing our needs and wants to our elected Federal officials in an informed manner. These efforts can contribute significantly to letting the public-private-government partners of I2 know how much we regard the Internet as our preferred medium for communication and education, and how important access, speed, interoperability, and reliability are to the future of our profession. ■

Bill Wyman is a member of the Information Technology Services staff at the University of Colorado at Boulder. He retired from 25 years of teaching German part-time in 1994 to begin another full-time career in computing and networking. Bill enjoys getting feedback on his columns, positive and otherwise. His email address is Bill.Wyman@colorado.edu.