# unofficial technology marvel of the millennium

BY GLENN RICART

As the year rolls over to an even 2000, the Internet appears to have captured the public eye as the technological innovation with the most promise for changing our lives. If only someone were in charge of the Internet, they would surely brand it as the "Official Technology Marvel of the Millennium." But no one is in charge. The distributed leadership of the Internet is by itself a big tip-off that this is an innovation crafted largely by higher education. Higher education first developed the Internet technology under the sponsorship of the government's Advanced Research Projects Agency (ARPA) and later was a leader in commercializing that technology.

It's a technology with a deep impact because of its disruptive influence on how we think about distance and time and because of the Internet's ability to dethrone technology and commerce leaders and crown new ones. Messages sent on the Internet span great distances in hundredths of a second. When used effectively, the Internet has helped compress the amount of time it can take to get something done, and now we speak of activities accomplished in "Internet time."

*Glenn Ricart's career has spanned federal government (National Institutes of Health), academia (University of Maryland), and major corporate enterprise (Novell). About every fifteen years he co-founds a successful startup; this is one of those times, and the startup is CenterBeam.*

Illustration by Brian Rea

In the next couple of pages, we'll take an Internet-time look at how this cyclone has grown over the past thirty years, with plenty of drive from higher education, and at where it appears to be headed. I'll conclude with some suggestions for additional leadership roles for higher education.

## A Cyclone of Information

We have the Internet to thank, at least in part, for our recent economic prosperity. The longest peacetime economic expansion in the history of the United States began in March 1991, assisted by the Internet. A study commissioned by the Global Internet Project and written by Takuma Amano and Robert Blohm showed that fully half of all new jobs in 1996 were due to the Internet (see <http://www.gip.org/gip9e3.htm>).

The Internet is a cyclone of information that derives energy from better technology and cheaper bandwidth. Its high winds are capable of destroying even established information sources and institutions, and its future path is largely unpredictable.

## A Brief History in Internet Time

There's a bumper sticker that states, "If you can read this, thank a teacher." If you can click on and interact with Web sites globally, you might also want to thank an academic. Despite the ubiquitous association of the Internet with dot-com, e-this, and i-that, the Internet was largely created and originally deployed by academics.

The story begins thirty years ago, in the late 1960s, when ARPA decided to link computers at a set of top research universities receiving ARPA funding. The big mainframes then typically in use were designed to queue large numbers of batch jobs. This was efficient for the computer, but researchers usually had to wait overnight to get their results. ARPA decided to change that model by providing online access to computers, typically PDP-10s, that the agency bought for university research. However, with no queue of batch jobs to keep the computers busy, how could the research community make maximum use of these expensive new machines? The answer was to invent a way to load-balance the new machines with work from multiple universities. Larry Roberts at ARPA found just the way to do that by taking advantage of a new idea called *packet switching* then being developed by Len Kleinrock and his team at UCLA.

ARPA contracted with Bolt, Beranek, and Newman (later, BB&N) to create the Interface Message Processor, or IMP—the first production packet switch. Developed to military specifications, the IMP was as big and heavy as the typical commercial refrigerator. A baseball bat could only slightly dent it. But with an IMP at each university and 56Kbps data lines between, one PDP-10 could talk to another. Using BB&N software called RSEXEC (Resource Sharing Executive), academic researchers at one university could use the computer at another. The protocols then in use were comparatively simple. Addressing was by a single byte. Could an experiment like this ever involve more than 256 computers?

By 1973, Bob Kahn at ARPA had worked with Vint Cerf at Stanford University and others to devise a new set of protocols with a bigger address range and a way of doing computer-independent communications called TCP (Transmission Control Protocol). They shared their invention with the networking community at a meeting at Sussex University. Yet this was not to become the standard transportation mechanism of the Internet until 1983.

Coordination between universities was needed, prompting Ray Tomlinson to invent network e-mail in 1972. Files needed to be moved between computers, leading to the creation of what would become FTP (File Transfer Protocol). With e-mail and FTP, the rate at which collaborative work could be conducted between researchers at participating computer science departments was greatly increased. Although we didn't quite recognize it then, Internet time had begun.

In fact, the collaborative advantages of Internet time were so great that computer science departments not on the ARPAnet felt disadvantaged and started their own, equivalent network: Computer Science Network, or CSNET. So not only was the original ARPAnet a network of universities, the second Internet network, CSNET, was also a network of universities. By the early 1980s, ARPAnet and CSNET had grown to nearly one hundred universities and military installations.

The first university, to my knowledge, to decide to embrace the Internet protocols campus-wide was the University of Maryland at College Park. Walt Gilbert, head of systems at the Computer Science Center, recommended in 1981 that the campus eschew the proprietary networks from IBM, DEC, and others and, using the fact that all vendors selling to the military were required to have TCP/IP software, convert the campus to this protocol.

Being able to provide campus-wide computer services using a network instead of a mainframe was highly attractive to me when I joined the University of Maryland in 1982, but there were a few problems to over-

**transmission control protocol**

**By 1973, Bob Kahn at ARPA had worked with Vint Cerf at Stanford University and others to devise a new set of protocols with a bigger address range and a way of doing computer-independent communications called TCP (Transmission Control Protocol).**

# Looking Forward

### Bandwidth Begets Reality

I must sheepishly admit that when I used my first 1200-baud modem, I concluded that modems need never get faster because 1200 baud was as fast as I could read. Of course, that was as fast as I could read the character-based interface then in use. An increase in Internet bandwidth to the home (and business) will continue to fuel graphics, detailed pictures, and greater realism in general in Internet interactions. In 1999, there was an explosion of MP3 audio. When sent over ISDN or cable modems or DSL lines or on-campus or business lines, superb-quality binaural audio could be sent in real time over the Internet. For the first time, we could saturate, with full fidelity, one of the five human senses.

I believe that we will continue to send more-realistic sensory information over the Internet. Pictures will acquire more pixels, movies will become smoother and more detailed, and an additional sense of reality will flow from greater control over the point of view. In distance education, each student will be able to have multiple simultaneous views (instructor, whiteboard, other students, relevant lectureware) and to control those views (move virtual cameras), as well as sending back her or his own multiple views (student, notebook, whiteboard, and/or Web information) to be shared with the instructor and the class.

### Connectivity Nearly Everywhere

Nearly every Internet-connected device will talk wirelessly. Low-power radio cells will connect and relay information. Bluetooth (see <http://www.bluetooth.com>), or a similar low-power high-bandwidth radio standard, will connect Internet devices in proximity to each other, enabling instant classroom and conference networks. The same technology will foster body-area (personal) networks that connect your wristwatch with your cellphone/pager (on your belt or in your purse) with the personal databank (multiple gigabytes) stored in the heels of your shoes with the voice and video I/O (including heads-up display) from your eyeglasses (cyberglasses). The wallet will disappear, since the waterproof wristwatch can produce any identification or authentication needed as well as store the anonymous certificates of value spent online and increasingly accepted as cash in what used to be called the "real world." Even if you kick off your shoes or set down your personal phone, your wristwatch remains with you.

At universities and businesses, wireless access points will fill the work and study spaces with overlapping microcells of connectivity (as we already do at CenterBeam!). I predict that before 2010, more than half of American homes will also have at least one low-power radio cell connected to 100Kbps or greater Internet bandwidth. Good hosts will naturally provide connectivity for their visiting guests. Universities will once again "re-wire" to provide multiple overlapping connectivity cells.

### Internet-Enabled Everything

Even though the personal computer may continue to evolve as the main Internet-access device, lots of other Internet-smart appliances will come to the market. For instance, the wristwatch, after authenticating its wearer to a personal computer, will learn the schedule for the upcoming week and will pass it on to the alarm clock when its wearer goes to bed. Televisions will, unprompted, record your favorite shows (oops, this is already happening; buy a Replay or Tivo). With Internet transmission and control, every telephone will easily place conference calls (see AT&T Click2Dial), and voice mail will be retrieved from your e-mail program (see <http://www.ureach.com>). When people wearing identifying wristwatches come to your front door, a voice will announce the names and/or roles of your visitors.

The personal automobile will also be a mobile personal-information store. In addition to repeating and amplifying the signals of the personal network devices of its occupants so that they have internal connectivity, the automobile will easily carry a very large cache of your digital "stuff." It will be easy to retrieve favorite music, talk, interactive games, pictures, and movies. Each of your visitors who brings a car will be able to conjure up home videos and simulations from the car parked nearby. While traveling, passengers will have the choice of looking out the window at the real world or looking in the window of their in-car display.

### An Ecology of Devices and Information

Internet-enabled "everythings" will form interdependent information ecologies. External dollars (funded by advertising or subscriptions) encoded into packets will pay for processors and storage devices throughout the network to cache and transport their information. In some cases, the information alone will be valuable enough to power the ecosystem, such as providing personal information (useful for targeting ads) in return for stock quotes or real-world weather forecasts.

### Jackets

In an even more interesting twist, packets within the network may directly carry logic or executable code. If the packets carry the popular Java programming language, we might call them Java packets or, more succinctly, Jackets. Jackets, like smart cards, can self-authenticate their creators. They can carry routing, filtering, and translating functions from one place in the network to another to improve efficiency or reduce bottlenecks.

Jackets will also be very useful for dealing with latency in the growing Internet. Internet traffic has grown by about 6 percent per month for the past decade, and it shows no sign of significantly departing from come. First, the newfangled IBM PCs did not have a TCP/IP implementation; we wrote one under contract to IBM. Second, there were no routers, and IMPs were much too expensive (reportedly about $250,000 each). But we did have David Mills, who, with intrepid students Mike Petry and Louis Mamakos, implemented IMP-like routing capabilities in comparatively inexpensive PDP-11 minicomputers. Not only did these PDP-11s network the campus, but the National Science Foundation (NSF) eventually gave us a contract to implement and support them in the original NSFnet backbone connecting supercomputer centers.

In 1982 and 1983, the Southeastern Universities Research Association (SURA) held meetings chaired by Jesse Poore, then of Georgia Tech, on where in the Southeast supercomputers might be placed to garner political support from throughout the area. The meeting ended inconclusively on the main issue; multiple institutions wanted a supercomputer. I was tasked to ask the NSF for money to connect one institution in each state of the region to a new southeastern network. Regardless of where the supercomputers ended up, the rest of the universities would be able to access them.

It was an immediately winning argument everywhere except at the NSF, where members of the supercomputer panel were supportive but wanted to spend their money on supercomputers and not on networks. After much discussion between SURA and NSF, the Division of Networking and Computing Research Infrastructure was created. Its first grant was to SURA for the creation of the proposed network.

Jack Hahn got SURA's network operational in January 1985 by connecting the University of Maryland to George Washington University, narrowly edging Richard Mandelbaum's NYSERnet (New York State Education and Research network) for "first operating network" honors. NYSERnet did become the first to employ T1-speed circuits, and NEARnet (Northeastern Education and Research network) became the first to use Ethernet-speed circuits. The San Francisco Bay Area Regional Research network (BARRNET) became the first to use the even faster T3 circuits.

SURAnet, however, became the first ISP, or Internet Service Provider, agreeing to connect all comers without the typical restrictions of an acceptable-use policy. This was a tough decision. Along with SURAnet co-founders Henry Schaffer and Morty Taragin, I reasoned that if we did not accept commercial traffic, industry would eventually form its own network, which would overshadow the academic network. If we accepted commercial traffic from the beginning, the academic influence would at least set precedents on what would be a single Internet.

Making sure the Internet did not splinter took additional planning. The federal networks (MILNET, ARPAnet, NASA Science Internet, and later ESnet [Energy Sciences Network]) were enticed to interconnect first at a FIX (Federal Internet eXchange) we created and located in College Park and later, for redundancy, at NASA Ames in California as well. Next, the University of Michigan under Doug Van Houweling set up a successor supercomputer backbone (also known as NSFnet) that interconnected the so-called regional networks as well. Third, as fully commercial Internet providers such as UUNET and PSINET arrived on the scene, the FIX was enlarged to interconnect them as well, creating MAE-East (Metropolitan Area Ethernet–East Coast). Until about 1994, most Internet networks were interconnected at College Park, Maryland, or via MAE-East, making College Park and the greater Washington, D.C., area arguably the center of the Internet.

**mosaic browser**

**The University of Illinois at Champaign–Urbana took the Web idea and made it graphical using the MOSAIC browser—the root of all Web browsers now in use.**

The software that drives the Web also has a story rich in academia. Nearly all Internet protocol stacks now in existence are derivatives of the protocol implementation work that was done at the University of California at Berkeley. The BSD (Berkeley Standard Distribution) implementations of the Internet protocols became the de facto definitions for any area in which the RFC standards documents were vague. The Stanford University Network (SUN) became the birthing place for the eponymous company that produced the hardware of choice on which to run these protocols.

For a significant time, the main applications were TELNET (to access a remote computer), SMTP (to send a message), and FTP (to move a file). The next breakthrough came from the University of Minnesota, which defined a new protocol called GOPHER (which also happens to be the mascot for this mascot-challenged university) in 1991. GOPHER brought up on your screen a text page that contained a number of choices for the succeeding screen. By typing in a number, you could navigate among the GOPHER screens, search, and download files.

Of course, we all know that Tim Berners-Lee at CERN (a university and government consortium) created what we now recognize as the Web and its http: (HyperText Transfer Protocol), which has become the single-most-used applications protocol of the Internet. Pointers could easily go to other sites, creating a single Web out of what might have been many individual GOPHERs. The University of Illinois at Champaign–Urbana took the Web idea and made it graphical using the MOSAIC browser—the root of all Web browsers now in use.

## Leadership Plan for Higher Education

It's always a tough call, but I foresee several trends for the coming decade of the Internet (see the "Looking Forward" box). If all this is going to happen, what role should universities and other educational institutions play in guiding this evolution? An excellent start is provided by the Internet2 project (see <http://www.internet2.edu>). Internet2 is exploring what new leadership roles higher education can and should play in creating the next generation of the Internet. There are important initiatives in exploring the effects of very high bandwidth collaboration and simulation, cross-network middleware, and new applications.

I'd like to highlight three areas that need higher education's leadership but that may be underappreciated.

### *Identity*

A significant inhibitor to many serious uses of the Internet is the lack of strong authentication of personal identity. A close link to real-world personal identity is needed before the Internet can be used to sign contracts, vote, unlock a car door, or trigger an emergency response. Such an authentication should be based, at a minimum, on something you know plus something you own. I've nominated a futuristic wristwatch, perhaps one that can verify its wearer by monitoring distinctive characteristics of an electrocardiogram, in the same way that exercise machines can take a pulse. The wristwatch would provide electronic authentication of its wearer by infrared or low-power radio.

Or perhaps the device should provide authentication of any of a number of roles. For example, it seems reasonable for me to have separate identities for my place in

## Looking **Forward**

that growth rate. About half of the growth appears to come from new people and devices on the Internet, increasing routing costs and hop counts. Hop counts in the Internet can now often reach forty or more to access a desired server. Even with continually faster routers and backbones, we are building in about an extra ten milliseconds of delay every year to our end-to-end Internet connections.

Jackets can help reduce this latency. By executing within the network, they can bring what would have been server functions toward the requesting device. Conversely, they can bundle a related series of requests (e.g., of the people at the meeting downtown tomorrow, find out whom I've recently met and/or who has sent me mail). That bundle of requests must contend with the full end-to-end latency only once, effectively reducing network delays.

The ecology of information will also extend to the execution of Jackets. Routers may execute code-rich Jackets in return for the Jacket telling the router about delays and congestion it encountered upstream in the network. Or entrepreneurs may add Jacket processors to the network in order to attract Jackets that offer nanopayments in return for their execution. (In a striking simile in today's real world, Akamai is offering free Web caches to ISPs who will install them.)

### Substitutes for Bandwidth

Jackets help us to see that there are substitutes for using bandwidth. While raw bandwidth is becoming much less expensive, disk-drive storage costs are declining at an even faster rate. Therefore, as time progresses, it will become increasingly more cost-effective to store information on a local disk than to send it again over a network. One of the same motivations that we saw helping to spawn Jackets—latency—is also at work in favor of local storage. These local-storage devices often go under the name of *Web caches,* but much more than Web information can be usefully cached: security certificates, files, copies of objects, and yes, Jackets.

Local computing can also substitute for bandwidth, but the economic argument is less clear-cut than the latency argument at the present time. Jackets can self-optimize by moving themselves through the network to the location that offers to execute them least expensively while still meeting the user's response-time criteria.

### Services

Jackets will often be used to implement online services. By contrast, the fraction of people who will buy and install a software package will decline. Software stores are doomed. Services delivered online will increase dramatically in popularity because: (a) they are available when needed, (b) you pay for only what you use, (c) prices can be low because there is no software piracy, (d) packaging and distribution costs are nearly zero, and (e) you are always using the latest and presumably most wonderful version. While the costs of commodity goods sold on the Internet are heading toward zero profit margins (and intentionally negative profit margins in the case of www.buy.com), the differential quality becomes satisfying service that builds loyalty.

### Player Rotation

It's very hard for the leader in a field to develop and deploy the next innovation in that field. Clayton Christensen recently memorialized that idea effectively in *The Innovator's Dilemma.* Longtime innovators and players in data transport, like AT&T, are less likely to be leaders in the future of Internet transport than are newer players, like Qwest, which don't have a large, established customer base to protect. Technology innovation induces player rotation. Join a startup!

my family, my responsibilities at my workplace, and my membership in American Coaster Enthusiasts. Each role has a set of personal data that is shared within that role but that is separated from data in my other roles. Each of these roles is a separate online identity. As the owner of all the identities, I can make a change that affects them all; perhaps I've moved and have a new address and telephone number. Calendaring busy/free information should probably be shared by all identities, but some information, such as a list of the roller coasters I've ridden, may be kept in only one identity. Identities will also permit online negotiations. My authorized agents will be able to use my secure identity to conduct business on my behalf on the Internet. A distinct identity with a profile and role-specific information is the indispensable online surrogate for the real-world Glenn Ricart.

A manageable set of secure online identities will also help to get me out of my "too many passwords" problem. I have separate Web identities and passwords for arranging travel on Travelocity, Delta Airlines, American Airlines, and United Airlines. And these are separate from my e-mail accounts and my Hotmail passport and . . . well, I know I have dozens and dozens of online identities, but I have no exact idea of how many or where. Many have become useless because I've forgotten their passwords. We need to come to a universal and exchangeable idea of identity, strong authentication to real-world people, and profiles and role-specific, information-rich identities. It's also important to handle explicit anonymity.

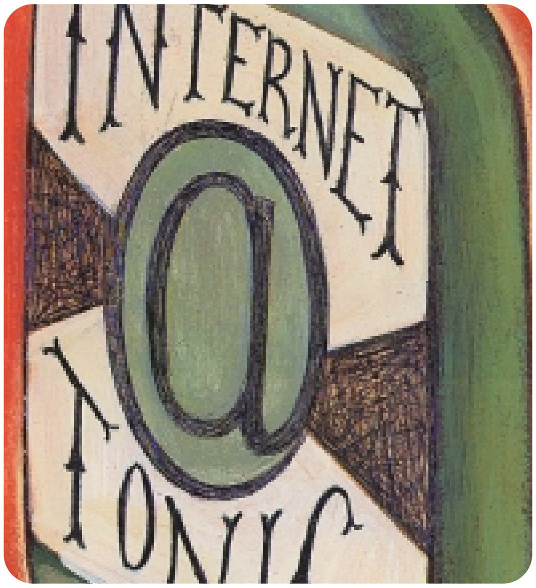


*Security*
The Internet has appallingly little security. This is not because we don't know how to make it secure. There are plenty of good, fast encryption algorithms and protocols. Part of the problem is inertia. For example, we know how to secure the Domain Name System, or DNS (the process that finds the IP address for a Web site) with an IETF (Internet Engineering Task Force) standard called DNSSEC. But most vendors are paying little attention because we haven't yet had a widespread, serious DNS spoofing attack on the Internet. But of course we shouldn't wait until the barn burns; we should install DNSSEC now.

Another part of the problem is a lack of trustworthy and guaranteed issuers of certificates. For example, we don't have a well-known, well-recognized, and convergent set of X.509 root certificates; as a result, we don't have widespread use of S/MIME secure e-mail. In the meantime, I can forge just about any e-mail to come from just about anyone and have it land in your in-basket. Higher education is the right group to take secure e-mail mainstream.

*Vouching for Information Quality*
Anyone with a modern word processor can save a document in html and publish it on the Internet. What's to be believed? What level of diligence has been expended to ensure the accuracy of the information and the soundness of the arguments? Higher education has always driven the creation of scholarly information. Although a growing amount is now freely available online, there are few ways to gauge quality. Higher education ought to create a rating scheme for the diligence, vetting, and refereeing involved in the creation of serious materials. The result could be expressed in a rating available to search engines, similar to the online PICS (Platform for Internet Content Selection).

## Disruptive Internet Time

The Internet has been one of the most disruptive as well as one of the most productive technologies in which higher education has been a significant driving force. Early in this century, Albert Einstein destroyed traditional Euclidean notions of space and time with his landmark general theory of relativity. Similarly, networking in general and the Internet in particular are changing our notions of space and time as they apply in the world of information and interaction.

The Internet can now transport information from nearly anyplace on the globe to nearly any other place in milliseconds. The distinct warp and weave of space now appears to be interconnected with a fast-growing number of Internet wormholes wherever backbones travel. Space is therefore much less distinct, and it's less valuable to us as a reference point.

Even more incredible, the Internet appears to be changing our notions of how fast things can happen. The phrase "we did it in Internet time" means it was done in a fraction of the traditional or expected amount of time. An apparent opposite to Einstein's time dilation, the Internet is a major cause of time compression.

I observe that Internet time is actually accelerating. The ratio of Internet days to calendar days is becoming larger. As speed has become a competitive advantage in itself, the Internet is allowing the newer leaders to increase that speed through higher bandwidth, better collaboration tools, and faster knowledge transfer and convergence.

Higher education will likely be the first to see and document the impacts of these effects. Like the explorers who discovered new continents, we are just beginning to discover the full impact of the Internet on information space and time. It's a great reason to look forward to the coming millennium.