

LESSONS for the

FUTURE INTERNET

Learning from the Past

By Michael M. Roberts

Twenty-five years ago, the Internet was a small dog at the bottom of the telecommunications pile, fighting for recognition. Now, everyone is using the net, and using it in more and more interesting ways, some of them controversial. Members of the U.S. Congress are wrestling with issues, such as Internet gambling and porn, that they never had to think about before. This article describes a few highlights of Internet development in the past, analyzes some of the policy factors at work in that development, and suggests some avenues for academic contributions to the successful evolution of the future Internet.

Mike Roberts has been active in EDUCAUSE and its predecessor organizations for more than thirty years. In 1987, he opened EDUCOM's networking policy office in Washington, D.C., and helped coordinate successful efforts to develop federal support for Internet technology. He was one of the founders of Internet2 and subsequently was the first president and CEO of ICANN (Internet Corporation for Assigned Names and Numbers).

Many bright, dedicated people had to fight very hard to get the TCP/IP protocols, on which the Internet is based, into general use.

Let's start by asking: What makes the Internet different from earlier communications systems? First of all, the Internet is based on *simple network protocols* that assume a smart computer is out at the end of the network. The original design of the Internet reflected the U.S. Department of Defense's interest in robust communications and in systems that might have some degree of survivability in the event of nuclear war. The network was put into prototype operation in the 1970s and then became the basis of ARPANet in the 1980s. Suddenly, people began to sit up and say, "You know, there are possibilities for doing a lot more with this than we anticipated."

Another characteristic of the Internet is its *interoperability*. Unlike massive digital telephone switches, which require precisely engineered components, the Internet has always been able to connect to all kinds of devices at all kinds of bandwidths. That capability continues today, as Internet services are newly delivered via cellphones, iPods, PDAs, and televisions. The interoperability of the protocols allowed Internet traffic to *move on top* of other communications systems, which permitted millions of home users to connect inexpensively over their analog telephone lines.

The Internet is also very *scalable*. As late as 1985, the Internet was connecting fewer than 100 hosts; now it is connecting more than 500 million, on the same basic protocols. And the Internet is quite *reliable*. We assume that our e-mail will get through—and it does—even though there are more than one billion active IP addresses on tens of thousands of federated internets all over the world.

The Internet is a *layered system*, not a silo system. Earlier communications technology tended to be vertically integrated and based on proprietary designs. That is, the technology would support only the one application that it was designed to support. As a consequence, when new application possibilities came along, new silos had to be erected. The Internet completely reversed that vertical integration. Instead, it is layered horizontally, with

much flexibility between the layers and in the layers.

As a result of all these characteristics, the Internet has become a powerful platform for the development of applications. And because of its basic design, it is capable of much *personalization* for the user. We take it for granted that when we sit down at our machines today, we will have our browser preferences,

our e-mail preferences, and our Excel and PowerPoint preferences all right there. The machine feels comfortable to us because we've been able to customize and personalize it for ourselves. That's very uncharacteristic of earlier systems and is one of the main reasons the Internet has spread so broadly and is used so widely.

Stage One: Research and Academic Focus, 1980–1991

The 1980s was the first great era—stage one—of the Internet, with a research and academic focus. Federal research funding contributed to the rapid growth of the net. For four or five years in the mid-1980s, there was debate about which protocols were going to be used. Books have been written about this disagreement. Many bright, dedicated people had to fight very hard to get the TCP/IP protocols, on which the Internet is based, into general use.

The National Science Foundation (NSF) took a leading role in research networking, first with NSFNet I, the "supercomputer net," and then with NSFNet II, a generalized Internet that fostered the creation of over one thousand Internet nodes on U.S. campuses by 1990. One of the things that helped this development was the work of the Internet Engineering Task Force (IETF) in creating open standards for the use of the Internet. Thus, anyone who wanted to implement the Internet anywhere in the world could do so and could communicate with all the other Internet users just by going to the Request for Comments (RFC) standards documents. This was a complete break from the past, when standards were controlled by major corporations and when the destiny of every application

depended on negotiation with a corporate monolith.

Another key development associated with NSFNet II was the formation of regional academic networks to serve as the operational connection between the NSFNet backbone and—ultimately—thousands of campus networks. Campus CIOs, with financial assistance from the NSF, provided the leadership to form and operate these networks.

Finally, the legitimacy process was completed when Congress passed and President George Bush signed the High Performance Computing Act of 1991, which put federal government research dollars squarely behind the Internet and declared that Internet technology was going to be the basis of federal research networking.

Stage Two: Early Public Internet, 1992–1997

By the early 1990s, the academic world was enthusiastically using the Internet and expanding its range. In 1991, the Federal Networking Council (FNC) made a decision to allow new companies, now known as Internet service providers (ISPs), to interconnect with federally supported internets. At about the same time, the NSF announced that it would gradually withdraw from support of NSFNet over the following several years, in anticipation of a transition to a largely privately funded and operated Internet.

Shortly afterward, the National Center for Supercomputing Applications (NCSA) at the University of Illinois adopted Tim Berners-Lee's work on the World Wide Web. Marc Andreessen and his team, who subsequently formed Mosaic and then Netscape, started us down the path to the browser environment of today. It was this watershed development that shifted the Internet from a command-line, e-mail, and file-transfer kind of user interface to the browser world of full-screen applications.

Starting in 1993, the Clinton White House adopted its Internet "Agenda for Action," implementing the High Performance Computing Act. It directed departments to use the Internet and to get the Internet deployed throughout the federal government, actions that also assured Wall Street that investing in the Internet was safe.

After years of debate, Congress enacted

the Telecommunications Act of 1996, which was both good news and bad news. On the one hand, the act didn't do anything to promote the Internet; on the other hand, it also didn't get in the way of network growth. The act was largely silent on the Internet.

For several years in the mid-1990s, campus CIOs debated their role in a post-NSFNet world, holding several national conferences. The significant conclusion from these debates was that higher education needed to maintain a leading-edge position in networking in order to ensure that advanced networking developments could be deployed in the service of research and education as rapidly as possible. In the fall of 1996, a group of more than thirty universities formed the University Corporation for Advanced Internet Development (UCAID), which subsequently became known as Internet2 and has grown to more than three hundred college and university, corporate, and affiliate members.

Stage Three: International Public Internet, 1998–2005

The Internet achieved both domestic and international critical mass in its third stage of growth, from 1998 to 2005. Actually, fueled by a giant speculative bubble in Internet stocks that peaked in 2000 and then collapsed, it reached more than critical mass. Stock market enthusiasm generated much progress in the development of new applications, especially Web-based applications, and in their underlying technologies. Of particular note were browser improvements, fiber-optic bandwidth improvements to gigabit-per-second levels, and price-performance improvements in personal computers. In short, the “bubble” years laid the foundation for broadband Internet applications and the accompanying integration of voice, data, and video services on one network technology base.

One consequence of the collapse of the Internet bubble was an oversupply of fiber-optic cable capacity in the United States. Networking managers in higher education soon realized that a “one-time only” opportunity to acquire operating rights to very-low-cost fiber lay before them. Rising to the occasion, several major leases of long-haul fiber were made for research and education purposes. One of

these efforts led to the creation, in 2003, of National LambdaRail (NLR), a nonprofit corporation whose purpose is to provide a national-scale infrastructure for research and experimentation in networking technologies and applications.

With the Internet suffering from growing pains in the mid-1990s, the U.S. government took a hard look at why research agencies were still handling the technical administration of the network. As a result, the government decided to move the technical administration into the hands of the Internet Corporation for Assigned Names and Numbers (ICANN). Incorporated in the United States, but with a board of directors selected on a global basis, ICANN operates under a “Memorandum of Understanding” with the U.S. Department of Commerce. Unfortunately, the good intentions surrounding the formation of ICANN as a private-sector organization designed to accomplish its technical functions through broad consensus mechanisms have largely not been realized. ICANN's work has been caught up in the rapidly increasing politicization of the Internet, a contest in which the struggling organization is outmanned and outgunned.

Stage Four: Challenges for the Future Internet, 2006–?

In its fourth stage of growth, the Internet has become a maturing, worldwide, universal network. The technology base continues to advance rapidly, with such recent developments as 100-gigabit transmission on Dense Wave Division Multiplexing (DWDM) optical fiber, Voice over Internet Protocol (VoIP), Internet service delivery to cellphones and other mobile devices, and IPTV, which delivers a broadband video stream to home computers. Rapid growth also continues, adding millions of addressable computers and devices every month.

With many of its technical and operational goals having been achieved or on the horizon, the Internet now faces renewed interest in its policy arrangements. As the reach of the net ex-

pands into more and more daily activities, the Internet has begun to mirror human society, with a great potential for both positive and negative consequences. The downside of its social impact has attracted national and international political attention. Issues such as porn, gambling, online fraud, and security have elicited sensible and also nonsensical political solutions.

The uses of the net now transcend conventional notions of time and space, of political and geographic boundaries, and of age and social class. So it is not surprising that politicians grounded in a previous era are having difficulty grappling with the challenges posed by a new transnational computer and communications system. What was once only a disruptive telecommunications technology is now a major force for change in economic and political affairs. Like any other disruption, this change is accompanied by both acceptance and resistance.

At least three important policy areas will receive attention in coming years: (1) the use of the Internet to achieve social goals; (2) the scope of government economic regulation of the net; and (3) the extent to which national security priorities should preempt network users' expectations for privacy. In each of these areas, legislation already exists, enacted in the days of conventional telephony systems and regulated monopolies. This creates a double challenge for legislators: what is the proper role for governments as the Internet continues to grow and change, and how do societies around the world

make the difficult transition from outmoded technology and obsolete laws to a new social, political, and technological equilibrium?

One currently debated policy issue—net neutrality—demonstrates the difficulty of aligning social and political arenas with the twenty-first-century reality of the net. Two of the few surviving U.S. telephone companies, Verizon and AT&T, have announced that they intend to levy special surcharges on broadband Internet traffic based on the

With many of its technical and operational goals having been achieved or on the horizon, the Internet now faces renewed interest in its policy arrangements.

The pioneering work of colleges and universities demonstrated that there was a viable market for Internet services and attracted investment capital for private-sector enterprises.

applications that the packets are carrying and on the company that is the source of the traffic. The two telcos, recently freed from regulation of broadband facilities by a Supreme Court decision, assert that their large investments in fiber-optic broadband facilities give them the right to recover these investments in any manner they see fit. The contrary view, shared by millions of Internet users, is that the extraordinary growth in functionality and value of the net could never have happened had there been proprietary, profit-based discrimination in managing packet flows—and that now is not the time to start accepting such discrimination.

But dissecting the arguments on both sides of net neutrality reveals more than a simple black-and-white controversy. Although some believe that “packets want to be free,” there is no doubt that Internet traffic does need management and that “discrimination” has always existed on the net, in the form of routing protocols, filters, and other devices intended to optimize traffic flows, for example. Likewise, the Internet has benefited from private investment for more than a decade and will continue to need generous amounts of capital as it grows. This leaves bystanders concluding that the net needs both discrimination and unregulated private investment, so what’s wrong with the telcos’ position?

The answer is that context is important. Both the telcos and the cable companies are still in the grip of a regulated monopoly mindset, where vertical silos and wide profit margins beckon. They know that commoditization of Internet access—the basic fiber-optic packet-transport facilities—is inevitable, along with accompanying diminished profit margins. But their response is not to invest in new applications that can create profitable business franchises in the upper layers of the Internet stack, which is what MSN, Yahoo, Google, eBay, and others are doing. Nor are they providing rich new medical, cultural, and educational resources for the public, which is what the American Heart

Association, the Library of Congress, the American Museum of Natural History, and colleges and universities and libraries across the nation are doing. Instead, the telcos and the cable companies are attempting to impose a silo view on their subscribers and extract profits through vertical integration and monopoly/duopoly control of access pricing.

Ultimately, of course, this approach cannot prevail. Silos may be fine for grain, but as a business strategy on the Internet,

they are headed for the trash heap. In the meantime, however, progress toward universal and affordable broadband access to all of the Internet’s services and resources (the United States stands sixteenth in the world in deployment of broadband) will be further delayed if the telcos’ well-funded campaign on Capitol Hill succeeds. Perhaps as important, this outbreak of bare-knuckles, profit-motivated lobbying strikes almost everyone—except the companies involved—as a corruption of the net for the sake of a doomed-to-failure business plan. In the world of politics, fairness still strikes a chord and may yet save the day for the Internet.

Where Does Higher Education Fit In?

Higher education in the United States played an important role in launching and expanding the Internet. Faculty teams did much of the original research; campus computing organizations prototyped early technology and later applied it in campus-wide production networks. Collaboration across campuses made the regional networks possible, and effective advocacy in Washington, D.C., served to legitimize the network and secured significant research and operational funding for national backbones and international connections. The wide availability of the Internet “platform” catalyzed multiple generations of creative discovery captured in new applications freely shared around the world. The pioneering work of colleges and universities demonstrated that there was a viable

market for Internet services and attracted investment capital for private-sector enterprises.

Maturing the net now requires not only continuing investment in network research and technology transfer but also a sharper focus on the social and economic impact of the Internet. Thus far, the search for institutional stability within the Internet community to guide future growth has largely failed. With no successful models to show the way, it is too early to know whether institutional failure is a fundamental characteristic of an iconoclastic and anarchistic medium or simply a temporary phenomenon.

The tradition of scholarly openness and collaboration in academia gives colleges and universities a special opportunity to contribute to the future evolution of the Internet. It isn’t necessary for a grand political design to be put in place for us to identify and address many pressing issues that are affecting our ability to expand Internet use within our scientific, research, and educational mission. Some of these issues involve work within the higher education community itself, whereas others involve focused advocacy to motivate work and funding by industry and government.

A sampling of areas where help is needed includes the following:

- **Basic Research.** Among developed countries, the United States is unique in its policy of allocating a large share of federal research funding to universities, awarding the funds on a competitive basis to individual faculty researchers. It is vitally important that the government continue its support of university research in networking in future years. Increases in the breadth and size of the network require constant updating of its underlying technologies, and this can be accomplished best by the bold and unfettered thinking of academic minds.
- **Advanced Network Facilities.** Since the early days, academic internets have provided a testbed for prototyping new network technology. Recently, the capability of such nets has expanded to include direct control of national-scale fiber-optic facilities, such as those operated by National LambdaRail. The

ability to test and develop very-high-capacity hybrid transmission facilities is vital to serving the rapidly growing needs of science and research and will be necessary on the public Internet in the not-too-distant future when the integration of voice, video, and data produces a worldwide bandwidth crunch.

- *Universal Affordable Broadband.* More than half of the students in higher education in the United States commute to school. From their homes, these men and women are dependent on the public Internet for access to campus networks and the wide range of learning resources now hosted and delivered by campus-based servers. With few exceptions, this access today is neither genuinely broadband nor affordable on students' budgets. Higher education has a direct stake in the development of a national policy for affordable broadband access to serve these millions of individuals. Although there are many

voices advocating such a step, Congress has been slow to adopt new policy goals and to revise existing telecommunications statutes so that the generously funded universal service programs of the telephone era are redirected toward broadband access. Concerted advocacy for broadband by the academic community is needed.

- *Middleware.* At the initiation of Internet2, it was widely assumed that "if you build it, they will come," replicating the enthusiasm with which NSFNet had been welcomed on campus in the 1980s. Unfortunately, it turned out that the dynamics of university research teams had shifted, at least in part as a result of the greater complexity of broadband research networks and their associated high-performance servers and computation engines. The result was an urgent need for the creation of a variety of development programs, known collectively as *middleware*. These

efforts are designed to provide a set of commonly needed software tools that interpose themselves between applications code and network facilities, thus the term *middleware*. Like other instances of necessary infrastructure, middleware has difficulty finding sponsors and funding. There always seem to be good reasons for letting "someone else" do the hard work. There also is occasional conflict over the compromises necessary to achieve a widely usable product, especially in the diverse academic community. This is an example of an important program priority in which the responsibility for progress lies largely within higher education and its corporate affiliates.

- *Preservation of the Internet Commons.* Those of us in higher education need to be thinking about the work we do and how it can move into the broader society. We are the ones that others rely on to be masters of technology and at the same time to be very concerned about openness and fairness. To the extent that there are a large number of individuals who truly believe the Internet represents a new, worldwide commons of society, the heart of that commons is in higher education. It is not by any means exclusively in higher education, but certainly the strongest voices are in the college and university community.

Preserving a commons is always hard. Those who take up the mantle of working for the common good generally encounter more detractors than supporters. Many times, social consensus about the value of a cause arrives long after the contributions of a dedicated few have been made. That there is an Internet commons at all, instead of a strictly commercial marketplace, is a tribute to the values of the original designers and developers. Having created a powerful instrument for advancing democracy and improving the quality of life—core values in higher education—we all share a responsibility for keeping the vision alive. *e*

Note

This article is based on a presentation delivered at the EDUCAUSE Net@EDU Annual Member Meeting in Tempe, Arizona, in February 2006.