DEFINING MOMENTS
THE INTERNET REVOLUTION

Kevin Hillstrom

Omnigraphics

615 Griswold, Detroit MI 48226

TABLE OF CONTENTS

PRIMARY SOURCES

Chapter Three

HOW THE INTERNET WORKS

The Internet is at once a worldwide broadcasting capability, a mechanism for information dissemination, and a medium for collaboration and interaction between individuals and their computers without regard for geographic location.

—Barry Leiner, "A Brief History of the Internet"

While many people use the Internet every day, few people understand how it works. The Internet is a vast network of computer networks that stretches around the world. Although each of the smaller networks that make up the Internet uses its own machines and software, all of these networks are able to communicate with each other through standard protocols, or rules, that define how the larger network operates. These protocols are like a common language shared by various types of computers and networks. Some analysts have described the communication between different types of networks that takes place on the Internet as being similar to a person from France and a person from Spain speaking to one another in English.

Ensuring the free flow of information among networks connected to the Internet requires a great deal of cooperative effort. For this reason, many people believe that the Internet must be under the control of a central organization or governing body. In reality, however, there is no centralized source of control over the functioning of the Internet. Instead, a number of different groups—including the Internet Society, the Volunteer Internet Engineering Task Force (IETF), and the World Wide Web (W3) Consortium—help guide its growth and development with input from computer scientists, software engineers, network managers, commercial enterprises, and other interested

A Michigan fourth-grader works on an assignment in his school computer lab.

parties. This culture of collaboration and sharing of ideas has existed since the earliest days of the Internet, when the graduate students who wrote the software for the first Interface Message Processors (IMPs) started the Request for Comments (RFC) system.

Since no one really owns or controls the Internet, some people wonder who pays to maintain it. The local networks that provide the basis for the Internet can be found in companies, universities, and government agencies around the world, and these various organizations generally provide the funding to support their own networks. In many cases, the organizations collect money from network users in the form of taxes, tuition, or fees for online services. It thus may be fair to say that users ultimately provide the funding for the Internet.

Packet Switching

One of the fundamental concepts underlying the operation of the Internet is packet switching. All information sent over the Internet is broken into

small pieces—less than 1,500 characters long—called packets. These packets are sent separately through any number of different routes to their final destination, where they are reassembled in the proper order. There is no single, dedicated connection that carries information from one computer to another over the Internet.

Some analysts have compared the packet-switching method of delivering information over the Internet with the operation of the U.S. Postal Service. When an individual customer sends a letter, it is mixed in with other letters and transported to the local post office. Then all the mail is sorted and sent to its final destination using the most efficient route. Similarly, an individual computer sends packets of information to a local network server or the host computer of an Internet Service Provider (ISP). Then the packets are mingled together with packets from other sources and transported through many different networks and computers. By the time the packets reach their destination, they may have been transferred through satellites, fiber-optic cables, Integrated Services Digital Network (ISDN) telephone lines, high-speed Digital Subscriber Line (DSL) connections, and wireless networks—all of which act like mail trucks and planes transporting information over the Internet.

The numerous small, local networks that make up the Internet are grouped together to form regional networks. These regional networks are connected together by Internet "backbones," or high-capacity lines that can carry huge amounts of data. Standing between these networks are powerful computers called routers. Routers read the "address" inside packets of information and decide how best to send each packet toward its final destination. In this way, routers work like a regional post office that sorts mail and places it on the appropriate trucks and planes. Though not all routers on the Internet have a direct connection to each other, every router has the ability to evaluate its available connections and determine the most efficient route to advance a packet toward its destination.

Protocols

All computers connected to the Internet operate using the same set of rules, or protocols. Transmission Control Protocol (TCP) is responsible for breaking down information into packets and then reassembling the packets in the proper order. Internet Protocol (IP) controls the process of addressing packets so that routers know where to send them. Together these protocols

Students study in the technology training center of an elementary school on the Spokane Indian Reservation in Washington state.

are known as TCP/IP. These days, all new personal computers come equipped with software that understands TCP/IP. The software is called Winsock in PC machines and MacTCP in Macintosh machines. This software functions like an interpreter between the computer and the Internet.

Every computer connected to the Internet has a unique IP address. This address consists of four numbers—each one lower than 256—separated by periods (an example of an IP address might be 66.139.79.225). Internet Protocol puts the IP addresses of the sender and recipient on each packet of information, sort of like an envelope, so that every other computer on the Internet will know where the packet came from and where it is going.

Though IP addresses are easy for computers to understand, human users of the Internet found it difficult to remember these long strings of numbers. Internet developers solved this problem by creating the Domain Name System (DNS), which made the Internet more convenient by allowing people to use

plain English words to locate individual computers on the network. The DNS designated a structure for host names, beginning with the specific identifier and followed by a period (dot) and a general classification or domain. The system originally divided Internet hosts into five "top-level domains" using three-letter abbreviations: "com" for commercial sites; "org" for nonprofit organizations; "gov" for government agencies; "edu" for educational institutions; and "net" for networking groups. A number of additional domains were added later, including two-letter country codes for international hosts. By dividing all hosts on the Internet into smaller domains, the DNS ensured that no names were duplicated and that every host could obtain the name of every other host. Every domain contains one or more "name servers," or computers that maintain updated lists of all host names and addresses within the domain. The name servers also handle the task of translating the Web site names into numerical IP addresses.
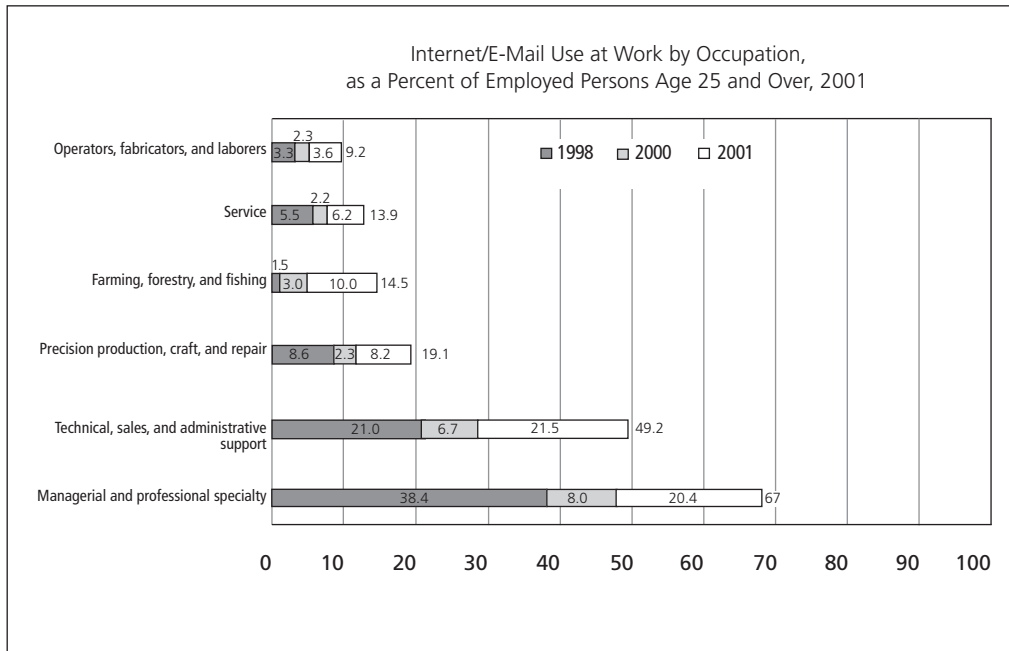
Communicating on the Internet

Electronic mail, commonly known as e-mail, is one of the most popular features of the Internet. E-mail messages are broken down into small packets and transmitted across the Internet in the same way as other data. In 2001 the Internet carried nearly 10 billion e-mail messages per day, and this number has increased every year since. E-mail provides a quick and relatively inexpensive way for people to keep in touch with friends, relatives, academic colleagues, and business associates in distant locations around the world.

An e-mail address typically consists of a user name and a host name, separated by "@"—the typographical symbol for "at." (An example of an e-mail address might be user99@aol.com.) A large number of individual computer users can access information and send messages using the same network server or host. In the case of home computers, the host is usually an Internet Service Provider—like America Online (AOL), Microsoft Network (MSN), or a local telephone or cable TV company—that provides access in exchange for a monthly fee. When a home computer is connected to the Internet through an ISP, it uses one of a shared pool of IP addresses belonging to the ISP. Each user of that computer can establish a unique e-mail address that will be recognized by the host, which will route messages to the home computer when it is connected to the Internet.

Most computers use e-mail software packages to manage mail. These programs allow users to compose and read e-mail, save and organize sent and

Internet/E-Mail Use at Work by Occupation,
as a Percent of Employed Persons Age 25 and Over, 2001

Occupation	1998	2000	2001
Operators, fabricators, and laborers	3.3	3.6 (2.3)	9.2
Service	5.5	6.2 (2.2)	13.9
Farming, forestry, and fishing	1.5	3.0	10.0 → 14.5
Precision production, craft, and repair	8.6	2.3	8.2 → 19.1
Technical, sales, and administrative support	21.0	6.7	21.5 → 49.2
Managerial and professional specialty	38.4	8.0	20.4 → 67

Source: NTIA and ESA, U.S. Department of Commerce, using U.S. Census Bureau Current Population Survey Supplements

received messages in folders, search through their messages, and maintain an address book of e-mail contacts. Many programs can also be used to create mailing lists containing e-mail addresses for a group of people with a shared interest, like members of a class or organization. Members of the group can then send messages to all names on the mailing list simultaneously. E-mail software also provides a means for users to attach text documents, pictures, and sounds to e-mail messages. The software on the sending computer automatically encodes such attachments as binary files so they can be "packetized" and transmitted over the Internet. Similar software on the receiving computer automatically decodes the binary files and displays them to the recipient in the original format.

Although e-mail can be a valuable method of communication, there are several problems associated with it. For example, regular e-mail is not secure because it passes through public lines and networks. In order to prevent snoopers and hackers from gaining access to e-mail messages, the messages must be coded using encryption software. In addition, destructive programs known as viruses can be attached to e-mail messages. When unsuspecting

recipients open the attached file, the virus becomes active and causes damage to their computers. Some e-mail viruses automatically send copies of themselves to every e-mail address in the recipient's address book, thus spreading the malicious program further across the Internet.

Another problem associated with e-mail involves the unsolicited junk messages known as "spam." Spam is typically sent out in mass mailings by commercial enterprises. The idea behind it is to attract people to a company's Internet site and convince them to purchase goods and services. In 2003 spam accounted for over one-half of all e-mail messages. Spam causes problems for recipients, who waste valuable time deleting dozens of unwanted e-mail messages. It can also clog servers and delay delivery of legitimate messages. Many ISPs and e-mail software packages contain filters to help users block out unwanted messages. Most of these filters are not very effective, however, because the largest senders of spam circumvent them by frequently changing or forging their return e-mail addresses. The U.S. Congress has passed legislation aimed at stopping spam, including the CAN-SPAM (Controlling the Assault of Non-Solicited Pornography and Marketing) Act of 2003, but such laws were expected to have little effect on spammers who route their messages through international hosts.

Two other popular forms of communication on the Internet are chat rooms and instant messaging. For many Internet users today, these services have revolutionized the way they communicate with others. Both services differ from e-mail in that they allow for real-time communication, meaning that recipients of messages are able to see words immediately. One of the most popular ways to chat on the Internet uses a protocol called Internet Relay Chat (IRC), which uses software like mIRC and xIRCon. IRC software is available for most types of computers. The software connects with an IRC server on the Internet, which is part of a global network of IRC servers. After logging in and selecting a "channel," or topic of conversation, participants can send and receive messages with other users of the channel all over the world. Instant messaging (IM) uses software like AOL Instant Messenger (AIM) and MSN Messenger. IM works in a manner similar to IRC, but chats are limited to a designated list of "buddies." The list is sent to an instant messaging server when a participant first logs in. The server then informs the participant when members of the buddy list are online and available to chat.

The Inventor of E-Mail

Ray Tomlinson

In 1971 a young computer engineer at Bolt Beranek and Newman (BBN) named Ray Tomlinson devised an ingenious—and now ubiquitous—new method of communication called electronic mail, or e-mail. Today, his invention is as familiar to many Americans and other people around the world as the television, the telephone, and other trappings of modern society.

Tomlinson's invention stemmed from his company's work on the ARPANET, precursor to the Internet. He has freely acknowledged that it originated not out of any officially sanctioned work, but rather out of his own conviction that the capacity to send messages electronically to far-flung destinations would be "neat." With that in mind, he devised a "send message" program that worked first on a local system, and then across all ARPANET nodes. He

The World Wide Web

The World Wide Web is essentially a sophisticated program that runs on the global computer network known as the Internet. Its introduction in the early 1990s helped change the Internet from a collection of separate networks into a web of interconnected servers. The World Wide Web consists of billions of interactive documents called Web pages. Web pages can incorporate words, pictures, sound, and other elements. They are generally built using Hypertext Markup Language (HTML). This coding system uses embedded commands or tags to identify the various elements in a document, but lets individual computers decide how to display these elements. In this way, HTML provides a shared format for documents that all computers can understand. As a result, anyone with basic computer skills can build a Web page using HTML and post it on a Web server using File Transfer Protocol (FTP) software.

40

later estimated that his e-mail program took him no more than six hours to complete, spread over a week or two "when I had a spare moment." He also indicated that he came up with the "@" symbol to designate remote mailbox destinations after "30 to 40 seconds of thought."

As it turned out, ARPA Director Larry Roberts liked Tomlinson's system so well that he began using e-mail for all his correspondence. Researchers eager for ARPA grants had no choice but to follow suit, and before long electronic mail was entrenched in the fast-growing ARPANET. And since that time, it has grown in tandem with the larger Internet phenomenon to the point that it is regarded as an essential tool in business and personal communication alike.

Tomlinson, who continues to work at BBN Technologies, acknowledges wistfully that it would have been nice if his innovation had brought him an infusion of personal wealth. But like many other pioneers in the development of the Internet, his efforts were not motivated by money, and he is proud of the enduring contribution he made to modern communications. "The cases where it has opened up new avenues of communication between people has been gratifying," he told *Darwin Magazine.* "I have received a number of e-mails from individuals who have found it to be a godsend in getting in touch with people and building a sense of community."

The term "Web site" usually refers to a collection of related documents or pages. The first or top page in a Web site, called the home page, generally acts as a table of contents to help guide users to information available elsewhere on the site. The various pages in a Web site are connected to each other using hypertext links—embedded HTML tags that identify the location of each page. On the computer screen, these links appear as highlighted words, symbols, or pictures. By clicking on a highlighted link, users can instantly jump to another Web page. Many Web sites also include links to other Web sites that feature related information.

Every Web page occupies a unique location on the Internet. Connected computers find and retrieve pages on the Web using a textual address called a Uniform Resource Locator (URL). The first part of each URL identifies the protocol that will be used to transfer information from the server—usually Hyper-

Internet pioneers Vinton Cerf, Lawrence Roberts, Robert Kahn, and Tim Berners-Lee pose for assembled media at a 2002 scientific award ceremony.

text Transfer Protocol (http). It is followed by the name of the Internet resource where the information resides—usually the World Wide Web (www). Next comes the IP address of the Web server or host computer that contains the page, including the three-letter domain abbreviation. An extended URL might also include a directory name, a document name, a date, and other information. (An example of a URL for a Web site might be http://www.isoc.org/internet/history.)

Individual computers interact with the World Wide Web using browser software, like Microsoft Internet Explorer or Netscape Navigator. These programs retrieve pages from Web servers, interpret the HTML codes, and display the information on the computer. Web browsers also automatically convert the text-based URLs into numerical IP addresses, since computers understand numbers rather than textual names.

Since the World Wide Web contains billions of pages, users need sophisticated tools to shift through the available information to find what they need.

The two main methods of searching for information on the Web are indexes and search engines. Indexes provide lists of Web sites broken down into categories. Users can select categories to narrow their search until they find the information they want. About.com is one frequently used Web index. Search engines are databases containing subject and keyword information for millions of Web pages. They are created using "spiders," special programs that collect information about the resources available on the Internet. Users begin a search by typing terms into the search engine. The search engine compares the terms to the information in its database and returns a list of documents that contain the terms. All of the major search engines create their databases and return search results in different ways. There is also meta-search software, which enables users to search through the databases of a variety of search engines simultaneously.

Benefits and Risks of Connecting to the Internet

The Internet is a valuable resource, facilitating communication and information sharing between people and organizations around the world. It allows users to access up-to-date news and weather forecasts, explore the collections of distant libraries and museums, sell goods and services, shop from the comfort of their homes, share interests through discussion groups and chat rooms, keep in touch with friends, collaborate on work projects, contact experts for help with problems, play adventure and role-playing games, and conduct business, academic, and personal research on thousands of topics.

For all its positive contributions, however, the Internet also exposes users to potential dangers every time they go online. The open, public nature of the network not only allows for the free exchange of information, but also leaves every computer connected to it vulnerable to hackers, viruses, and invasions of privacy. Hackers can go through high-speed Internet connections to gain access to home computers or corporate networks. They can take advantage of such access to steal information, damage data, or use the computer system for illegal purposes, such as sending spam or launching attacks on ISPs or Web sites. The main means of protection against Internet hackers are firewalls—combinations of hardware and software that create barriers to prevent unwanted traffic from flowing between a computer or network and the Internet.

The Internet is also a common source of viruses and worms—malicious programs that cause varying degrees of damage to computers and networks.

Students work in a school computer lab in Lahore, Pakistan.

Internet users can expose their machines to viruses by downloading files or opening infected e-mail attachments. Worms are usually disguised as helpful programs or files. When they are opened or run, however, they spread from computer to computer across a network. Worms can overwhelm systems with mail or unnecessary tasks to the point that they "crash." Likewise, viruses can corrupt data, programs, and system files so that computers behave oddly or stop working altogether. Antivirus software provides the best protection against viruses and worms. This type of program typically scans incoming e-mail and downloads and warns the user when malicious programs are detected. Most antivirus programs are able to quarantine or eradicate viruses and worms so that they are unable to damage the computer system.

There are still more potentially harmful programs. Spyware is a type of program that runs behind the scenes on a computer, usually without the knowledge of the user. Spyware is usually downloaded from the Internet by accident along with a useful program. Once installed, it tracks which Web sites the user visits and regularly reports back to a central spyware Web site. This site creates a profile of the computer user's surfing habits and takes advantage of the information to deliver self-launching advertisements (commonly known as pop-up ads) targeted specifically toward the user's interests. For example, a user who frequently visited Web sites devoted to financial planning might receive pop-up ads relating to banking and investments. Many Internet users find pop-up ads annoying and feel that spyware represents an invasion of their privacy. The only way to remove spyware from a computer is to employ software programs that are specially designed to kill it.

Tim Berners-Lee (1955-)
Creator of the World Wide Web

Tim Berners-Lee was born on June 8, 1955, in London, England. His parents, Conway and Mary Berners-Lee, were mathematicians who met in the early 1950s while working on the Ferranti Mark 1, England's first commercial computer.

Berners-Lee was fascinated by computers throughout his childhood and adolescence. As a student at Queen's College at Oxford University, he used a soldering iron, an old television set, and an assortment of electrical equipment to construct a primitive—but operational—computer. After graduating in 1976 with a bachelor's degree in physics, he worked for several British firms as a software designer and telecommunications research engineer.

Creating Enquire

In 1980 Berners-Lee spent six months as a software consultant at CERN (Conseil Européen pour la Recherche Nucléaire), the European Laboratory for Particle Physics in Geneva, Switzerland. During his stay at CERN, he cobbled together a computer program for his own personal use that linked together the information in his daily schedule planner, his address book, and papers and documents on which he was working. He called this experimental program "Enquire," after *Enquire Within Upon Everything,* a Victorian-era encyclopedia of sorts that he recalled from his childhood.

The Enquire program actually utilized a programming innovation devised nearly two decades earlier by a Harvard University programmer named Ted Nelson. Under this concept, dubbed "hypertext" by Nelson, any word or phrase in one electronic document could be outfitted with embedded codes that would enable it to link to another part of the document—or to an entirely different document, for that matter.

Building on this conceptual foundation, Berners-Lee established a series of electronic links, first between the various documents in his own computer,

and then to the computers of the hundreds of other scientists working at the vast CERN nuclear research facilities. These links created a veritable "web" of information through which Berners-Lee could move quite freely (at this point in time, the computer mouse had not yet been invented, so instead of the "point and click" system used by modern computer users, he had to manually key in the link identifiers). This program, quietly cobbled together over a period of a few months, would eventually become the conceptual foundation for the future development of the World Wide Web.

In 1981 Berners-Lee moved on to Image Computer Systems in England. He spent the next three years there, designing hardware and communications software for printers, but in 1984 he received an invitation to join the staff at CERN. He quickly accepted the offer and returned to Geneva, where he began working on developing new information systems for the institution.

Around the time that Berners-Lee returned to CERN, the research facility connected its computers to a growing worldwide network that became known as the Internet. Over the next several years, Berners-Lee became increasingly focused on finding ways to access information stored on remote computers that were also connected to the Internet. Returning to his Enquire program as a starting point, he began fleshing out the vital communication protocols needed for navigating the Internet and transmitting documents across the network. "I happened to come along … after hypertext and the Internet had come of age," he explained in *Weaving the Web*. "The task left to me was to marry them together." This work was never officially sanctioned by his supervisors at CERN, but they gave him quiet encouragement to pursue his ideas.

Weaving the World Wide Web

The fall of 1990 proved to be a particularly pivotal period in the development of the World Wide Web. It was during these autumn months that Berners-Lee invented Hypertext Markup Language (HTML), a coding system that identified various elements in a document, such as words, pictures, and sound. These codes became the language that allowed for the formatting of Web pages to include text, headings, graphics, and hypertext links. At the same time, Berners-Lee established a communication standard called Hypertext Transfer Protocol (HTTP), which allowed documents to be linked together across the Internet and retrieved in readable form from remote locations by special software programs (later known as Web browsers). Finally, he

invented the Uniform Resource Locator (URL) system, which gave each document a unique "address" on the Internet.

As he neared completion of his work, Berners-Lee recognized on some level that his combination of software programs and networking protocols had the potential to create an electronic "web" of information capable of spanning the globe. Thus, after considering and discarding several possible names for his creation, he decided to call it the World Wide Web.

In the summer of 1991 Berners-Lee unveiled the World Wide Web to the public. Rather than trying to make money from his invention, however, he posted the software on the Internet for anyone to download free of charge. Several months later he formally relinquished all rights to the HTML, HTTP, and URL programs he had created, even though he and many others sensed that the Internet was on the cusp of explosive growth. Berners-Lee has never expressed any remorse about this decision in subsequent years, however. In fact, he strongly defends his idealistic choice. "What is maddening is the terrible notion that a person's value depends on how important and financially successful they are, and that is measured in terms of money," he wrote in *Weaving the Web*. "That suggests disrespect for the researchers across the globe developing ideas for the next leaps in science and technology. If your requirement is to make a large amount of money, then your options in life are rather small."

From 1991 to 1993 Berners-Lee further refined the design of the World Wide Web, absorbing feedback from users across the Internet. Meanwhile, adherents to his system were literally remaking the Internet with each passing day, using his programs to build exciting new Web sites and browsers to access them. Berners-Lee watched the whole thing unfold with understandable pride. "I think the main intention was to make the thing fly," he commented in Stephen Segaller's *Nerds 2.0.1*. "When you're really attached to a dream of how things could be, then you puruse that dream and it's very, very satisfying to see it work. The fact that the World Wide Web did work, I find is just exciting for itself. Exciting that you can have an idea and it can take off and it can happen. It means that dreamers all over the world should take off and not stop."

Director of the World Wide Web Consortium

In 1994 Berners-Lee accepted a position as a research scientist at the Laboratory for Computer Science (LCS) at the Massachusetts Institute of Technology (MIT). Shortly after arriving at MIT, he founded the World Wide

Web Consortium (W3C). Today, this consortium includes more than 300 companies that meet regularly to discuss Internet issues, coordinate Web development, and ensure that the Web's wealth of information continues to remain accessible to everyone. Berners-Lee continues to serve as the director of the W3C from his modest MIT offices.

In recognition of his enormous contributions to Internet technology, Berners-Lee has received numerous prestigious awards over the years. In 1999 *Time* magazine named him one of the "100 Most Influential Minds of the 20th Century," explaining that "unlike so many of the inventions that have moved the world, this one truly was the work of one man…. The World Wide Web is Berners-Lee's alone. He designed it. He loosed it on the world. And he more than anyone else has fought to keep it open, nonproprietary, and free." In December 2003 he was knighted by Britain's Queen Elizabeth II for his "services to the global development of the Internet," and in June 2004 Berners-Lee became the first-ever recipient of the Millennium Technology Prize, a $1.2 million cash prize presented by the Finnish Technology Award Foundation.

Berners-Lee has long been known as a semi-reclusive figure who avoids the trappings of fame. But while he is vigilant about protecting the privacy of himself and his family, he readily volunteers his thoughts on the future of the Internet and information technology. He expresses great optimism about the promise of new technological innovations, and speculates that today's Web will eventually evolve into what he calls the Semantic Web—a network in which data will have meaning to the computer, enabling it to make logical inferences and carry out routine research and analyses that currently blunt human creativity. "It's important that the Web help people be intuitive as well as analytical," he wrote in *Weaving the Web*. "If we succeed, creativity will arise across larger and more diverse groups."

Berners-Lee does, however, express concern about the potential erosion of the Web's longstanding universal access philosophy. He is troubled, for example, by growing speculation about the introduction of Web pages accessible only to people using certain browsers. "For me the fundamental Web is the Web of people," he declared in a 1999 address during the LCS's 35th anniversary celebration. "It's not the Web of machines talking to each other; it's not the network of machines talking to each other. It's not the Web of documents. Remember when machines talk to each other over some protocol, two machines are talking on behalf of two people. The Consortium has a

whole technical domain 'Technology and Society' which recognizes that, at the end of the day, if we're not doing something for the Web of people, then we're really not doing something useful at all."

Sources

Berners-Lee, Tim, with Mark Fischetti. *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor.* New York: HarperCollins, 1999.

Berners-Lee, Tim. Transcript of speech delivered at MIT Laboratory for Computer Science 35[th] Anniversary celebrations, April 1999. Available online at http://www.w3.org/1999/04/13-tbl.html.

Henderson, Harry. *History Makers: Pioneers of the Internet.* San Diego: Lucent, 2002.

Segaller, Stephen. *Nerds 2.0.1: A Brief History of the Internet.* New York: TV Books, 1998.

Bill Gates Discusses the Internet's Potential Impact on Education

Microsoft founder Bill Gates has emerged as one of the leading figures in the drive to spread Internet access to people across the United States and around the world. He believes that the Internet offers many advantages, particularly in the area of education. In this excerpt from a 1999 speech at the New York Institute of Technology, Gates describes some of the potential uses of Internet technology in education. He also suggests that the Internet will be such an indispensable tool for today's youth that they should be called "Generation I."

This is truly an exciting era. The major advances to be made in technology are transforming the way we do business. It's allowing us to work in more efficient ways and to reach out for new ideas and new products that wouldn't have been possible before. The news is full every week of new start-up companies that are taking advantage of these advances. But today I want to focus on what I think is really the most exciting use of the new technology. And that is to use it as an educational tool for the new generation. There's a lot that we need to do to achieve the potential here for bringing this tool into the classrooms. In fact, if we look at it in a generational time frame, we can say that the kids who have been born just recently, they will grow up with the Internet as a fact of life for them from the very beginning. If we give them the right opportunities, they'll show us how to take full advantage of what can be done with all these advances.

The PC era started 25 years ago. That's when the first microprocessor chip came out from Intel. That's when Microsoft was founded by myself and Paul Allen. And we had a vision then that the combination of microprocessor technology along with great software from thousands of companies would create a tool that would be valuable for everyone. Our vision was a computer on every desk and in every home. And during the last 25 years, we've made incredible progress towards that dream. Today, over half of U.S. households have a personal computer. And in the workplace it's become a standard tool for creating documents and communicating.

Now, a key element of this is taking those PCs and hooking them together using the standards of the Internet. If we look at the statistics there, we

have further to go. Only about 30 percent of kids at home today have both a PC and Internet access. Now, that's up very dramatically from just a few years ago, so the trend is in the right direction, but 70 percent of kids don't have that access today. We're quite certain these numbers will continue to grow as the cost of the hardware and the cost of the communications comes down, and also [as] it's viewed as more and more of a vital tool. The things that are really driving this phenomena are that more and more of the world's information is becoming easily accessible on the Internet, whether it's medical information, travel information, or just plain chat rooms where people are interested in getting together and talking about a topic that they care about.

If we look at schools, we're going to have to play the primary role in getting this technology out to kids. Some of the statistics give us hope that we can really fulfill this mission: 90 percent of schools have some type of Internet access. Now, what that means varies quite widely. In some cases, that means a single computer—if you had to wait in line, you'd get very impatient for an individual student to actually have a chance to go out there, and really browse the Internet. So that figure, in some ways, overstates where we are.

Our goal should be that this Internet access should not just be in a lab somewhere, but it should be literally in every classroom. It should be part of the curriculum, and getting teachers enthused about bringing this in as a new tool. After all, every student starts out with amazing curiosity. They want to explore subjects, they want to find kids who have similar interests. And the Internet is fantastic for letting people go out and explore things on their own. And if it's done properly, you can bring back to the classroom the things that you found that were very interesting, and share those with all of the other kids.

Now, the number of Internet connections for each student in the United States is today about 14 students per Internet connection. That's much better than a year ago where it was 20, but still very far away from making it something that everybody gets a chance to use at leisure. The majority of schools report that teachers are really using the Internet, they're at least going out there to find materials that they can present in the classroom as well as involving in the kids in using the Internet. But the vision of what this should be is far different than we have right now. We talk about a "connected learning community" where the parents are going online and seeing what's being assigned. They're helping to research a topic, working together with their kids. And the curriculum itself has been redesigned to take full advantage of this.

178

And who are the kids that are going to have the full impact of this? Certainly the kids being born today, from 1994 on, they're a new generation, and nobody has really labeled this generation, so I would propose today that we think about calling this "Generation I." Of course, "I" for Internet. You know, these are kids who will always wonder why we talk about having records. To them, music will just be something you can get on your computer, and organized exactly the way you want and carry around with you however you want. To them, the idea that all the rich information should be easy to search and find, and that you should be able to find other kids in another country and speak to them about what their thinking is about that topic. They'll simply take that for granted. They'll think of buying as something where you can go out and get the best prices, or get the product reviews across the Internet. And so, they will think about the Internet in a far more profound way than most of us who grew up without it being an ever-present tool. And, in some ways, this is very, very exciting. These kids will be agents of change as they move out into their jobs. These kids include my own children. I've got a three-year-old and a four-month-old who are definitely going to be leading members of Generation I, and they have their high-speed Internet connection from the very beginning.

So, what will this lifestyle be like? Well, a lot of it will be saying that when you pull information together, you don't have to worry about writing it down on paper. You'll be able to go to the Internet and submit your homework that way. You'll be able to read books with very incredible screens. One of the things that is still ahead of us, but is definitely going to happen in the next three years, is to have flat panel screens that you can hold in your hands, just like the tablet, so it's comfortable to sit and read for long periods of time. That tablet display will have a resolution that makes it as comfortable as reading off of paper. I'm not saying that this will eliminate books, but it will give us a new flexibility to be able to call up material that otherwise wouldn't be easy to get to. And you can take textbooks and customize them for your class, take the portions that make sense from many different sources, and by bringing it together that way, the kids can browse it electronically. They can comment on parts of the text, send back something with a voice annotation or a handwriting annotation, whether they're confused about something they're reading or they want to make the comment on it.

That whole process of collaborating, letting kids show their ideas, will be very different than it's been—where we have to work strictly out of the

textbooks that are often out of date, and not really tailored to that particular child's interest. Kids will all have a little smart card, so whatever PC or tablet they see, they just will run their smart card through the reader, and it will immediately bring up the things that they care about. Their electronic mail, their schedule, and they'll have access to that information wherever they go.

These kids won't think about the phone and the PC as being two different worlds. Whenever you're using the PC and you go to a Web site, you'll be able to click and talk to people, and likewise the phone that you carry around more and more will have a pretty good screen on it. And so, calling up things about the weather or new messages, all of that will be available from the phone as well as the PC. Likewise, the TV will let you connect up, not just playing great games, or getting any video you want, but also navigating the world of information. So, we'll have one set of standards around the Internet, and all of these devices connecting up to that common network.

So, it will be a world where everything is online, and that's simply taken for granted. And these kids will, in a creative way, build the Web sites that will make the Web sites we have today look like really nothing, sort of in the same way that you look at the early TV shows, early radio shows, and realize that the medium was not being fully exploited, there was so much more that could be done by people who really grew up with it and thought about it as central to their life.

Having all these devices makes it a lifestyle activity. You won't be surprised when you look at your small screen and say, okay, what's the traffic looks like, or what are my friends up to. That will just be something you expect to do all the time. All the material that's in the library today is slowly but surely being digitized. All the new books, the periodicals, and even the books from the past that are hard to get when they're in paper form. In fact, a little later we'll see a great example of how digital access is letting people see the history and get back to their roots in a way that wouldn't have been possible without digital technology. So, the Internet, the power in the microprocessor, the miracle of great software, and other advances like the flat screen displays, will make this far more natural to use. The interface will include speech recognition. You won't have to type everything that you're interested in, you'll simply be able to talk to the device and ask it to help you find the information.

So, this will really be the most incredible tool that's ever been created. The way that we get the most out of it is, we connect everything together. We

have all the material that museums can put online, and have accessible for kids. We have people in the community, we have parents at home, and then we have the schools themselves all sharing across the Internet. So, the first time teachers will find that if they're teaching a subject in a particularly creative way, they can put up the material they're using and the approach they've taken onto the Internet, have other teachers find out, perhaps add things into it, and share it back to the original creator, and everyone else. In fact, I think there will be a lot of great awards that are given to teachers who go to the effort to take their good ideas, put them out on the Internet, and make them available to everyone else. So, whatever subject it is, whatever level it is, you'll be able to go out there and find neat, exciting ideas that really draw the kids in, that use examples that they can relate to by using the breadth of material that the richness of the Internet provides.

When I keep saying Internet, I don't just mean text pages, or even pages with just pictures. The Internet that we're talking about in the years ahead is one where audio and video are a full part of the experience. And so, if you want to have, say, an interactive experiment where the kid can play around with some of the different parameters, you can have rich software that can create a video experience showing you exactly what would happen when you change the different variables. So, if it's playing around with the physics equation, or the design of something, right there on the screen, you'll be able to get a lot of the feel of how the different inputs control the thing. And so you'll want to pursue that and go with your full curiosity to understand exactly how those things work. So the boundaries won't be the same as they've been today. And this connected learning community, the idea of everybody contributing, will be commonplace.

So what are the key things we need to do here? We need to get great content out there. We need to have rewards for everybody who is putting up that content. We need ways of classifying it and linking it together, so that it's easy for people to find something that they might want to be interested in. We need to get teachers involved, so that the design of learning is built around this. We need to make sure that it's not just a small percentage of the kids who've got access. We don't want to have a divide here, where the kids who have the Internet at home are able to go and do wonderful things, and the kids who don't have that access in their home don't have a way to have the equivalent experience. We want to make sure that we're avoiding material online that is damaging, so the idea of how you control that, how you train kids to use this in a responsible way—that's very, very important.

In terms of content, the progress here is pretty fantastic. One of the things that Microsoft got involved in very early on is the idea of taking the encyclopedia and creating a digital version. And so we created over five years ago what we call Encarta. And that was delivered on a CD, and included not just the text and pictures, but also audio and video. In some ways the encyclopedia is a great example of how digital information can be more accessible. Now, I remember when I was a kid I had a copy of the *World Book,* and I thought, well, how am I supposed to read this thing? Well, I read it alphabetically. But, it's kind of strange, because you're reading about the 1600s, then the 1800s, then the 1900s, and it's sort of hard to get a feel for the subject when that's the only way that the information is organized. And when information is changed, you get the year book every year and you paste in the labels, but you really haven't revised all the things that should be changed there.

Well, moving this into digital form means that you can navigate the subjects in a far more natural way. Tell me all the articles about the kings in England. Give me all the articles about great scientists. And so it's far more natural to pursue your curiosity and go through the material that way. You can have timelines, music, all the things that paper form doesn't provide. Another key point here, and one that shouldn't be missed, is that it can be very inexpensive. Today, for $40 or $50, these electronic encyclopedias provide far more information than even the print version that was very difficult for a lot of households to buy, because just the printing costs and everything made those cost $300 or $400. So the results are quite dramatic here. This is the first paper document where the electronic form is far more popular than the paper form….

I've said that we're going to have to put a lot of investment into thinking about how to teach this Generation I. You know, how does technology fit in? Technology by itself is not the answer. The answer is teachers, great teachers, who are using technology in the best way. And, there's a lot of investment to be made in these teachers. Really helping them to feel prepared for the modern classroom. It's got to be intimidating when you've got kids who, in some ways, are ahead of you in using this tool. You have no idea how you're going to control their usage, and you have a curriculum that you need to teach that really hasn't been adapted to fit the Internet and what can be done there. And so, it was interesting in a recent survey, only 20 percent of teachers said that they really feel prepared to bring in the technology in the right way. That's quite a stark contrast to the fact that, yes, we are creating the connections,

and that's going to happen, but we haven't made the investment in the teachers to really bring them along and get them to drive this forward. And that's the only way it really can come together.

So, I think there is a clear call to action here. How do we make sure that the education curriculum is changed, how do we make sure that as teachers refresh their skills, they're really brought into this? How do we take those textbook budgets and really think not just about what's on paper, but also about what can be done on the Internet, and how those two things relate to each other?

So, there's a need for some massive training to be done, and this is a case where government at all levels, corporations, and philanthropists really need to come together to have this take place.

The kids today are anxious to have this opportunity. And so, the sooner we can get all the teachers to be enthusiastic about it, the better. There are many companies jumping in on this. We've done a lot of things ourselves with training labs. We have our Web site that we connect people up to training opportunities, and through our work, including the material we've created, over a million teachers have been trained. Now, compared to the need that's only getting part way there.

One of the programs that is most exciting that we've done is called "anywhere, anytime learning." And that's where you see that the best thing to do for a student is to actually give them their own computer. When they have their own computer, they have that sense of ownership, they don't have to stand in line to use it, they can go home at night, if there's something they're confused about they get as much time as they want, they can sit with their friends and talk about it. The teacher can know that every kid in the classroom has this tool. That's where you get the greatest impact. And we call that anywhere, anytime learning, because this is a program based on getting portable computers to all the kids, and then connecting up the network in the classrooms so they can print, and display their information. There are 500 schools around the world who've adopted this idea, and it's something that the results have really been quite fantastic….

Getting the prices of the machines down, so you can have them at home, coming up with special financing options for things like laptops, that's very important. The schools themselves, whether it's special levies, or allocation of resources towards technology, clearly that's a very central role. The libraries

are another place that can be part of this. That's one area where the foundation that I created has helped to make sure that over the next several years every library in the United States, all 16,000, will have a state of the art PC with an Internet connection. So kids who can reach the library can get in and take advantage of that. And the usage of that has been quite phenomenal.

We should also reach out to community centers, any clubs, any place where the kids go we should make sure that the technology is there. Now there is special work that needs to be done to make sure that the kids are using this tool in a constructive fashion, and that they understand some of the things they ought to know and be careful about. We've worked together with Boys and Girls Club to come up with what we call the Stay Safe Online program, it's the kind of things that teachers, or a community center, or parents can go through with children who are using this tool, and really give them guidelines about how it should be used properly. And so that's, we think, an important element of how this can move forward.

There's no doubt in my mind that we've got a fantastic opportunity here. The people at Microsoft who work on creating the software that's part of it, part of the reason they love their job is they know that these tools will be used in education in some exciting ways. We're really just at the very beginning of this, but I can say with great confidence that the Internet is going to change education as fundamentally as it changed when we had printed books. And that's something that it's going to be very exciting to be a part of.

Source: Gates, Bill. Remarks at New York Institute of Technology, October 28, 1999. Available online at http://www.microsoft.com/billgates/speeches/10-28GenI.asp.

IMPORTANT PEOPLE, PLACES, AND TERMS

Advanced Research Projects Agency (ARPA)
Agency formed by the U.S. Department of Defense in 1957 to develop new technology for military and scientific use

Andreessen, Marc
Co-inventor of Mosaic, one of the first successful Web browsers, and founder of Netscape Communications

ARPANET
The world's first computer network, developed by Advanced Research Projects Agency (ARPA) researchers in 1969, and a precursor to the modern Internet

Barlow, John Perry
Writer, legal expert, and co-founder of the Electronic Frontier Foundation, and a leading proponent of free speech on and open access to the Internet

Berners-Lee, Tim
Computer expert and inventor of the World Wide Web

Bolt Beranek and Newman (BBN)
Technology consulting firm based in Boston, Massachusetts, that built the first Interface Message Processors (IMPs) for the ARPANET

Brin, Sergey
Co-inventor of Google, the first Internet search engine to rank Web sites based on the number of links they receive from other sites

Broadband
Various types of Internet connections—including Digital Subscriber Lines (DSL), cable modems, and Wi-Fi—that offer greater bandwidth and faster download speeds than traditional dial-up telephone access

CHRONOLOGY

1957

The Soviet Union launches *Sputnik,* the world's first artificial satellite, into orbit. *See*

1958

U.S. President Dwight Eisenhower creates the Advanced Research Projects Agency (ARPA) within the Department of Defense.

1961

Leonard Kleinrock, a computer science expert at the University of California-Los Angeles (UCLA), publishes the first paper on packet-switching theory.

Researchers Paul Baran at the Rand Corporation and Donald Davies at the National Physical Laboratory in England develop packet-switching theory independently of Kleinrock.

1962

J.C.R. Licklider becomes director of ARPA's Information Processing Techniques Office (IPTO).

Licklider publishes an academic paper describing his vision of a "Galactic Network" of interconnected computers that would allow people to access information from any site around the world.

1966

Bob Taylor becomes director of IPTO and launches the ARPANET project.

1968

ARPA managers hire Bolt Beranek and Newman (BBN), a small consulting firm based in Cambridge, Massachusetts, to develop key components of the ARPANET.

1969

BBN engineers deliver the first Interface Message Processor (IMP) to Leonard Kleinrock at UCLA.

Some of Kleinrock's graduate students—including Vinton Cerf, Steve Crocker, and Jonathan Postel—write the software program that connects UCLA's mainframe to the IMP.

SOURCES FOR FURTHER STUDY

Gralla, Preston. *How the Internet Works.* 7th ed. Indianapolis, IN: Que, 2004. This informative text provides an overview of the technical aspects of the Internet's operation, using language and format accessible to students. It includes details about the inner workings of routers, wireless networks, search engines, instant messaging, viruses, spyware, and many other aspects of the online experience.

Hafner, Katie, and Matthew Lyon. *Where Wizards Stay Up Late: The Origins of the Internet.* New York: Simon and Schuster, 1998. A fascinating account of the early development of the Internet, beginning with the conception of the ARPANET in the 1960s and ending with the decommissioning of this first computer network in 1990.

Internet Society. "All About the Internet." http://www.isoc.org/internet. The online home of one of the Internet's governing organizations, this site includes a wealth of information about the development and growth of the network, as well as links to histories, biographies, and important documents.

National Museum of American History. "Birth of the Internet." http://smithsonian.yahoo.com/birthoftheinternet.html. Compiled by the Smithsonian Institution and Yahoo!, this site offers information on important people, companies, and events in the history of the Internet and World Wide Web.

Richards, Sally. *FutureNet: The Past, Present, and Future of the Internet as Told by Its Creators and Visionaries.* New York: John Wiley and Sons, 2002. Such famous Internet pioneers as Vint Cerf and Leonard Kleinrock reminisce about the challenges of bringing the Internet to life and speculate about the possible directions the technology might take in the future.

Segaller, Stephen. *Nerds 2.0.1: A Brief History of the Internet.* New York: TV Books, 1998. As the title implies, this book focuses on the technological pioneers who contributed to the development, growth, and commercialization of computer networking. It provides an entertaining yet informative account of the history of the Internet through the personal stories of its creators.

BIBLIOGRAPHY

Books and Periodicals

Abbate, Janet. *Inventing the Internet.* Cambridge, MA: MIT Press, 1999.

Abram, Stephen. "Twenty Reasons to Love IM." *Information Outlook,* October 2004.

"The Abridged History of the Internet." *Internet Magazine,* January 2004.

Angel, Karen. *Inside Yahoo! Reinvention and the Road Ahead.* New York: John Wiley, 2002.

Berners-Lee, Tim, with Mark Fischetti. *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web.* New York: HarperCollins, 1999.

Brody, Herb. "Net Cerfing." *Technology Review,* May-June 1998.

Ceruzzi, Paul E. *A History of Modern Computing.* 2nd ed. Cambridge, MA: MIT Press, 2003.

Conner-Sax, Kiersten, and Ed Krol. *The Whole Internet: The Next Generation.* Sebastapol, CA: O'Reilly and Associates, 1999.

Conrades, George. "The Future of the Internet: Predicting the Unpredictable." *Vital Speeches of the Day,* April 1, 1998.

Cothran, Helen, ed. *The Internet: Opposing Viewpoints.* San Diego, CA: Greenhaven Press, 2002.

Dolliver, Mark. "The Digital Divide Just Isn't What It Used to Be: 'Quality of Access.'" *Adweek,* September 27, 2004.

"The Dot-Com Bubble Bursts." *New York Times,* December 24, 2000.

Egendorf, Laura K., ed. *The Information Revolution: Opposing Viewpoints.* San Diego, CA: Greenhaven Press, 2004.

Ferren, Bran. "The Future of the Internet." *Discover,* November 2000.

Gates, Bill. *The Road Ahead.* New York: Viking, 1995.

Gonsalves, Antone. "Internet Posts Fastest Growth in Political Spending." *Internet Week,* August 19, 2004.

Graham, Ian. *The Internet: The Impact on Our Lives.* Austin, TX: Raintree Steck-Vaughn, 2001.

Gralla, Preston. *How the Internet Works.* 7th ed. Indianapolis, IN: Que, 2004.

Grossman, Lev. "It's All Free!" *Time Almanac,* 2004.

Hafner, Katie, and Matthew Lyon. *Where Wizards Stay Up Late: The Origins of the Internet.* New York: Simon and Schuster, 1998.

Henderson, Harry. *History Makers: Pioneers of the Internet.* San Diego: Lucent Books, 2002.

INDEX

(ill.) denotes illustration