97-392 SPR
Updated December 24, 1997

CRS Report for Congress

Congressional Research Service • The Library of Congress

Internet Technology

Ivan P. Kaminow
and
Jane Bortnick Griffith
Specialist in Information Technology Policy
Science, Technology, and Medicine Division

Summary

The Internet is often described as a "network of networks" because it is not a single physical entity but, in reality, hundreds of thousands of interconnecting networks (hence the term Internet). This global connectivity is the result of the initial network design by Internet pioneers that created a simple, common standard for linking computers over telecommunications networks. The Internet continues to grow at a fast pace and new hardware and software provide enhanced capacity and new services. The most popular Internet applications today are electronic mail and the World Wide Web. Rapid growth and the prospect for further expansion of the Internet have raised many issues for policymakers. Knowledge of the underlying technology provides a basis for analyzing public policy issues related to the Internet.

Introduction

Growth of Internet Hosts (in millions)

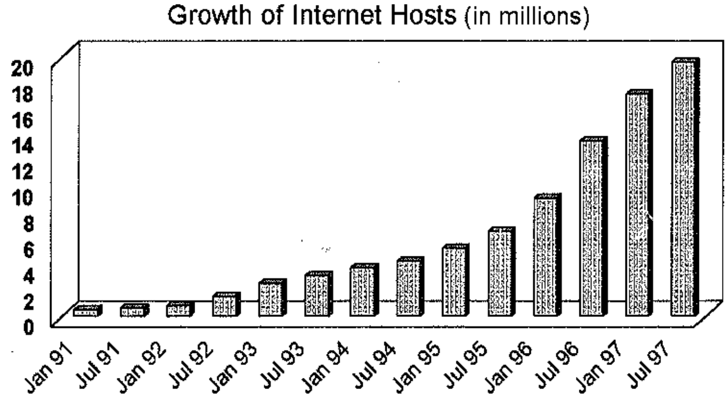

The Internet links millions of host computers around the world. Each host may be linked to smaller local networks, such as office networks, that typically link several computers, or to individual computers that connect over conventional phone lines. As a result, millions of computer users worldwide are capable of exchanging many kinds of information easily. The Internet continues to grow at an enormous pace as indicated in the accompanying figure.

Making the Connection: The Telephone

With the growing popularity of the Internet, computer and consumer electronics manufacturers, cable operators, and other service providers are competing to provide Internet access in a variety of ways at lower costs. The competition follows naturally from the convergence of the underlying technologies. Computer data, and increasingly voice and video signals are processed, stored, and transported digitally.[1] The next generation of television equipment, called Advanced Television (ATV), will be digital instead of analog. Hence, any digital network, appropriately designed, will be able to deliver many electronic information services.

The Internet originated with research funding provided by the Department of Defense Advanced Research Projects Agency (DARPA).[2] Its initial goals were to create a military network connecting critical computers and providing remote computer access to military commanders, which could survive enemy attack. As its use expanded, a civilian segment evolved under the National Science Foundation (NSF) and other science agencies. For many years, experimental networks linked research computers in university, industrial, and government labs. They provided a means for sharing resources, exchanging computer files, and exchanging messages electronically. When the Internet began to thrive, the government decided to leave commercial development to U.S. industry, turning over support of the U.S. Internet backbone to the private sector in April 1995. NSF and several other science agencies continue to fund advanced networking for the research and education communities.

The key technical idea that enabled the Internet to grow so successfully was *open architecture networking*. This allowed network service providers much greater flexibility to design their systems independently as long as they could link through a high-level interconnection protocol. That was a significant departure from the traditional circuit-switching approach used by the telephone network.[3]

Telephone and Internet Technology

Most people see the Internet in terms of how they use it for applications, such as sending electronic mail (*email*) or retrieving information from a World Wide Web site. These are described more fully below. The computer interface is designed to be "transparent," so that the user does not "see" the intervening physical network. Because the telephone network and the Internet are closely related, but have significant distinctions, analysis of legislative issues concerning appropriate government roles and regulatory frameworks requires a modest understanding of the underlying technical infrastructure.

---

[1] Digital signals convert information into binary digits or bits, which are ones or zeros used to represent different kinds of data. See Congressional Research Service, *Telecommunication Signal Transmission: Digital vs Analog*, by Richard M. Nunno, Report 96-401, 7 May 1996.

[2] For an overview of the history, growth, and costs of the Internet, see Congressional Research Service, *Welcome to Cyberia: An Internet Overview*, by Rita Tehan, Report 96-242.

[3] Leiner, Cerf, Clark, Kahn, Kleinrock, Lynch, Postel, Roberts, and Wolff, *A Brief History of the Internet*, Available online at: http://www.isoc.org/internet-history. For another useful history of the Internet see also: John Adams, "Geek Gods," *Washingtonian*, November 1996, p. 66.

Connection: The Telephone

A good illustration of underlying telephone technology is provided by following a long distance call connection. The dial tone indicates a connection to a local central telephone office. The dialed ten-digit number is forwarded through a switch in the central office to a regional office, which sets up a path dedicated to the call through switches configured to handle the traffic across the country. The path is determined by automatically updated "routing tables," which take account of nationwide traffic patterns and possible equipment outages to select the optimum path. The switches in the local central office and the regional office are owned by a Local Exchange Carrier (LEC), such as Bell Atlantic, serving the originating area code. The next switch in the sequence may be owned by an Interexchange Carrier (IXC), such as AT&T, selected by the paying party. The IXC switches continue across the country to the area code of the destination. The final switches are the regional and local switches of the destination LEC, such as Mountain Bell. After the telephone rings and is answered, the circuit from the calling telephone, through all the switches, to the called telephone is assigned to that two-way conversation for the duration of the call — even if there are long silences.

The sound pressure of the speaker's voice is converted by the telephone instrument into an electrical current that is an exact replica, or analog, of the sound. At the first switch, the analog signal is converted to a digital code that represents the same information by a stream of regularly spaced *bits* (binary digits), which are electrical pulses that may be present (as "ones") or absent (as "zeros"). Because of their simplicity, digital signals can be manipulated much more readily than analog signals to eliminate extraneous noise and distortions. In addition, many calls can be combined by intermingling their bits to provide more efficient transmission on a high-speed cable. In summary: (1) telephone circuits are switched over networks owned by several independent, regulated companies that have cooperative agreements for exchanging traffic, maintaining quality, and billing calls; (2) the circuits are dedicated to the call and not shared (i.e., they are circuit-switched); (3) the charge (but not the cost) is proportional to the time and distance of the call (i.e., metered or usage-based charging); and (4) except for the first and last segments (of about one mile each), signals are digital.

Making the Connection: The Internet

An Internet connection operates differently from a telephone connection. Generally, the connection can take one of two configurations. In an organization — a business, a university, a government agency — individual workstations often are connected through a local area network (LAN) to a host computer that serves the organization. As a user composes an email message, for example, the keystrokes are stored digitally in his workstation. When the message is transmitted, the mail is transferred to the host computer. The host can distribute email and other services to clients on the LAN, or route it to an Internet gateway computer.

For the user at home, a cheaper solution is to dial up an Internet Service Provider (ISP) by telephone when access is desired. A modem, which converts the digital signals to analog signals and vice versa, installed in the home computer provides a telephone connection from the computer through the telephone central office to one of a bank of modems at the ISP. Currently, most commonly installed modems operate at 28.8

kilobits/sec,[4] although modem speeds continue to increase. This is fast enough for email but can make accessing graphics exceedingly slow. Integrated Services Digital Network (ISDN) service provides digital data and voice lines at higher speeds, but at rather high cost.[5] A new digital technology, known as ADSL (Asymmetric Digital Subscriber Line), however, is stretching the bandwidth of some of the twisted pair of copper telephone wires that run to the home to a few megabits/second over limited distances. Many houses also have potential access to greater bandwidth through the coaxial cable provided by cable television service.

Increasingly, devices are being used to access the Internet other than the typical desktop computer. They may be simplified network computers (NC) that rely on connections to large remote computers for processing and memory capability and are therefore considerably less expensive and easier to use; televisions with built-in or set-top-box Internet access (so-called Web TV); cable modems that connect home computers to the Internet via the existing cable system; or other "Internet appliances" that are under development to provide both wired and wireless connections. One industry research firm, International Data Corporation, estimates that by the year 2000, 22% of Internet-access devices will be machines that are not personal computers.[6]

For Internet connections, the digital data stream from the workstation is broken up into *packets*, which are segments of about 1,000 bits per packet. Each packet contains both the digital destination and source addresses as well as a packet sequence number. The convention for forming and handling packets defines the Internet Protocol (IP), introduced at the inception of the Internet. The complete protocol is in fact a set of conventions. It includes the Transmission Control Protocol, so that the complete set is designated TCP/IP. These conventions allow different networks to connect to each other over the Internet.

The IP packets generally are sent from an organization over a leased telephone line, or from a home, over local phone lines, through the central office switch, to a commercial Internet Service Provider (ISP), whose network may be limited to a city, state, or region, or may be national. A *router*, which is a form of computer switch, stores each packet and reads its IP destination address and, using a routing table, directs it to the next router. The routing tables are automatically updated by communication with neighboring routers. As the packet makes its way across the country, hop-by-hop, it may pass through routers and cables owned or leased by several independent ISPs before arriving at the destination. As various routers and cables become congested, packets making up the same message may be sent along other paths, changing the order of arriving packets, so that the message must

---

[4] The bit-rate, or number of bits per second, is a measure of information transmission speed or bandwidth. A standard telephone call requires 64,000 bits per second, or 64 kilobits/sec. Other standard rates are 24 calls combined into 1.5 million bits per second, or 1.5 megabits/sec, which is commonly known as a T1 line. Thirty T1 channels can be combined into a 45 megabits/sec or T3 channel. The technology for expanding the availability of bandwidth continues to improve, while data compression techniques also increase the amount of information that can be transmitted over these channels.

[5] ISDN connections for both voice and data provide up to 144 kilobits/sec.

[6] "The Race Is On to Simplify," *Business Week*, 24 June 1996, p. 7.

be reassembled using the sequence numbers. The host computer at the destination delivers the reassembled message when the recipient connects to the network.

A key advantage of packet switching is that computer and telecommunications facilities can be efficiently utilized by interleaving short packets of information from various sources to fill any available capacity. Thus, packet switching allows the sharing of resources at reduced cost. There is a price to pay, however, in terms of delays as the packets await their turn to be delivered. Those delays can be particularly troublesome for two-way, real-time applications such as Internet telephony. Host computers can be added anywhere in a packet network, and the network can be controlled by the routing tables, which send packets around congestion or equipment outages. In other words, control of packet switching is "distributed," whereas circuit switching is "centralized." The distributed nature of packet switching has contributed to the growth of the Internet, both because of ease of expansion and flexibility in developing a range of applications that adhere to the relatively simple TCP/IP protocol.

In contrast, the telephone network was designed for one application—voice telephony. In addition, the switches in the local telephone office have been configured specifically to handle voice traffic most efficiently. The average telephone conversation lasts three to five minutes. In contrast, Internet connections last considerably longer, with some users leaving connections active constantly. Congestion problems at local telephone switches have resulted from the difficulty of handling this type of data traffic. Companies are evaluating various technical solutions and business investment options to address this issue.

Bottlenecks, resulting in delays or lost packets at routers, can occur throughout the network as the local traffic increases. To make the Internet more suitable for commercial applications, national and global Internet providers are beginning to offer end-to-end service for their customers. These major network providers include the telephone interexchange carriers, AT&T, MCI and Sprint, as well as companies, like UUNET and PSInet, that have grown up with the Internet. Many smaller Internet providers lease services from the larger companies. The low economic and technical barriers to entry have allowed the appearance of thousands of ISPs.

Applications

Internet usage at home and in the office is growing rapidly. According to a recent CommerceNet/Nielsen survey, the number of Internet users has more than doubled since the fall of 1995. The survey found that 23% of all persons older than 16 years in the United States and Canada used the Internet during the month prior to the survey. It also found that the number of people shopping on the Internet is significantly increasing.[7] The more popular current network applications are described below.

---

[7] Press release located at: http://www.commerce.net/work/pilot/nielsen_96/press_97.html

Electronic Mail

Electronic mail is the application that initially popularized the Internet. It allows a user to type an address or a group of addresses, compose a message, and send it to another user, a group of users, or a computer that can respond with textual information. The destination may be in the next office or residence, on the same campus, or on the other side of the Earth. Each user is assigned an address by the administrator of the host network serving that office, or by the Internet Service Provider (ISP) providing access from home.

Email Addresses

The addresses consist of a user name and a domain name. Examples include: *user@mail.house.gov* for the House e-mail server, *user@aol.com* for a user connected to America Online as Internet Service Provider, or *user@lab1.ntt.com.jp* for a user at lab1 in NTT in Japan. The portion to the right of @ (the "at" symbol) is the domain name, which gets wider in scope as one moves to the right. The "top level domains" (TLDs) *.com*, *.org*, *.edu*, *.gov*, *.mil*, and *.net* represent company, organizational, educational, governmental, military, and networking enterprises, respectively, and *.jp* or *.us* represent country domains. The actual addresses (called IP addresses) are numerical and are translated by a domain name server from the name addresses, such as those above. Recently proposals have been made to expand the number of TLDs. For a more extensive discussion of domain names see: CRS Report 97-868.

World Wide Web

The key application that has produced the recent dramatic increase in Internet traffic is the World Wide Web (WWW). It is based on the traditional academic concept of providing the reader with further references in a text. In the Web application, instead of footnotes or references, certain author-designated words or phrases are highlighted as *hypertext*. Companies can offer a wide array of distributed information and can advertise products in this manner. A simple click on the hypertext brings the associated file to the screen. If multiple series of links from a home page to subsequent references and back to the home page are mapped out, the map resembles a web; hence the term "world-wide web."

The real explosion of Web usage began when user-friendly software or *browsers* became available that offered attractive color graphics for retrieving Web pages. To access the Web, the user types in an address that is similar to an email address. However, the addresses generally are prefaced with "http://", where *http* refers to the *hypertext transfer protocol* (or software convention), while the *hypertext markup language* (*html*) provides the linking capability built into the Web software on servers and browsers. Web addresses are designated as Universal Resource Locators (URLs).