CRB *California Research Bureau, California State Library*

THE INTERNET: A CALIFORNIA POLICY PERSPECTIVE

Prepared at the request of
Assemblywoman Elaine White Alquist
Representing the Silicon Valley

By

Kenneth W. Umbach, Ph.D.

March 1997

California Research Bureau
California State Library
CRB-97-002

900 N Street, Suite 300
Sacramento, CA 95814
(916) 653-7843

CRB *California Research Bureau, California State Library*

THE INTERNET: A CALIFORNIA POLICY PERSPECTIVE

Prepared at the request of
Assemblywoman Elaine White Alquist
Representing the Silicon Valley

By

Kenneth W. Umbach, Ph.D.

Digest

The Internet, an international network of computer networks, began as a federal government initiative in the 1960s, but has now become a predominantly private and increasingly commercial system. Use of the Internet is rising exponentially, with no sign yet of leveling off. Growth over the last few years has been fueled by the rise of value-added, consumer-oriented networks such as America Online and CompuServe, by the graphics-oriented Worldwide Web, and by easy-to-use Web browser software--Mosaic, Netscape, and the Microsoft Internet Explorer.

Information available on the Internet encompasses practically every aspect of human activity, from classic literature to technical documents, city guides, news, statistics, K-12 lesson plans, casual real-time conversation, academic discussions, electronic mail, photographic and artwork archives, legislation, and government services at federal, state, and local levels. Information is placed online by individuals, businesses, organizations, publishers, and government agencies. Although there are no exact counts, users of the Internet have been estimated to number 40 to 50 million, with perhaps as many as 17 million currently in the United States, according to some estimates, and rising rapidly. California has been among the leading states in its citizens' and businesses' use of the Internet.

New technologies are bringing Internet access by way of cable, telephone lines, and wireless communications systems. New access devices promise to make access as simple as turning on a television set and navigating from site to site with a handheld remote control. "Smart cards" and other electronic payments methods could enable easy, secure online payment for purchases, thereby making the Internet an important part of the consumer economy.

Especially in view of its rapid and continuing growth, the Internet is important to California and raises issues of public interest for several reasons. These include:

- California state agencies are increasingly using the Internet as a means of providing services and information.

- The Internet serves as a means of communications among state and local agencies and between citizens and their government.

- Those state presences online raise complex and difficult security issues.

- Increasing use of the Internet for business and commerce is raising or complicating issues in the areas of regulation, trademark protection, taxation, and competitiveness.

- Demands are growing for public education to provide access to the Internet for use by students, teachers, and administrators, raising issues of how to pay for and implement

access, standards for use of the Internet, and changes that may be needed in curriculum and methods.

- California firms have an important role in the Internet industry in the areas of hardware, software, communications, and content.

- The privacy rights of Californians may be compromised by posting of personal information and by interception of private communications online, and some online activities may be hazardous to personal security.

- Concerns about equitable access to the Internet are growing as it becomes increasingly important in commerce, government, education, and entertainment.

This paper reviews the background and key terminology of the Internet, its current scope and use, and a selection of public policy issues. The paper also provides examples of the California state government and federal presences on the Internet and summarizes a selection of sites illustrating the myriad types of private-sector online services and information.

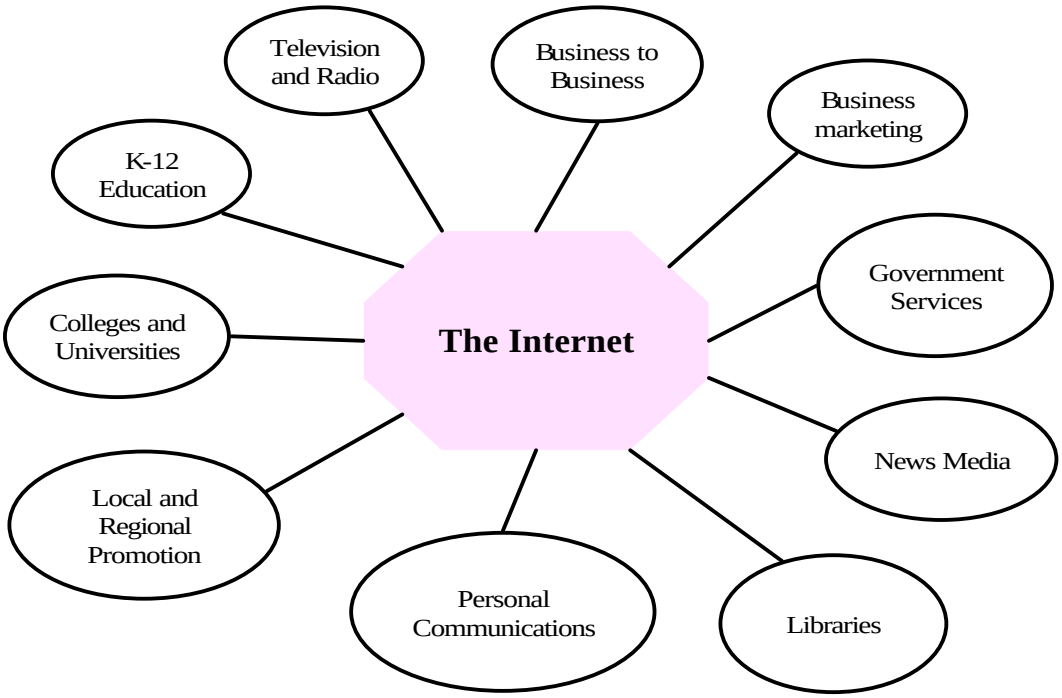

The Internet Provides a Shared Communications Channel

Contents

The Internet: A California Policy Perspective

> "This is the '90s, man, *anything* is possible." (Customer overheard at IHOP, Bakersfield, August 1996)
>
> "Computing is not about computers any more. It is about living." (Nicholas Negroponte, *Being Digital*)
>
> "[T]he medium is being oversold, our expectations have become bloated, and there's damned little critical discussion of the implications of an online world." (Clifford Stoll, *Silicon Snake Oil*)

INTRODUCTION

The rise and explosive expansion of the Internet and the Worldwide Web could have as profound an effect on society as did the rise of the automobile nearly a century ago. The Internet, the private preserve of a technical elite only a few years ago, is now used daily by millions of people around the world for work, research, education, entertainment, shopping, and personal communications. As the Net has opened to the general public, it has also blossomed into a graphically oriented and commercially driven environment unimagined by its originators. The Internet is a valuable tool, but it does raise a wide range of new policy issues and bring a new dimension to some old ones.

Why is the Internet Important to California?

Especially in view of its rapid and continuing growth, the Internet is important to California and raises public issues for several reasons. These include:

- Californians make heavy use of the Internet, and have been estimated to represent about a quarter of all U.S. users of the Internet.[1]
- California State agencies are using the Internet to provide services and information to the public.
- The Internet is a means of communication among State and local agencies and officials and between citizens and their government.

[1] Although exact counts do not exist, users of the Internet have been estimated to number 40 to 50 million worldwide, with perhaps as many as 17 million currently in the United States, according to some estimates, and rising rapidly. *Boardwatch* magazine, a computer/communications industry publication, is more conservative in its estimate, suggesting that about 10.4 million persons in the U.S., or about 3.9 percent of the U.S. population, had Internet access as of May 1996. *Boardwatch*'s worldwide estimate is 17 million, or about 0.3 percent of world population on line. ("Introduction to the Directory of Internet Service Providers," Spring 1996, posted at http://www.boardwatch/com.)

- State use of the Internet raises questions about the security of government Internet sites and of information posted on and sent to and from those sites.

- Use of the Internet for business and commerce is raising or complicating issues in the areas of regulation, trademark protection, taxation, consumer protection, and competitiveness.

- Public schools are using the Internet to serve students, teachers, and administrators, raising issues of how to pay for and implement access, standards for use of the Internet, and changes that may be needed in curriculum and methods.

- California firms have an important role in the Internet in the areas of hardware, software, communications, and content. Prominent California firms include Intel (microprocessors), Netscape (browsers), and Cisco Systems (internetworking hardware).[2]

- Information posted online poses risks to California residents' privacy and security.

Concerns about equitable access to the Internet are growing as it becomes increasingly important in commerce, government, education, and entertainment. The role of the State in assuring equitable (and safe) access is unsettled, especially in view of rapid changes in the technology and growth in the Internet.

Intended Audience and Purposes of this Paper

The intended audience for this paper is policy makers and the general public. This audience includes people with little knowledge of computers, communications, or networking *and* people with experience and knowledge in those areas. For that reason, I have tried to stay in a middle ground, introducing technical terms and explanations where necessary, but focusing more on concepts and on the policy issues raised by the growth of the Internet. In short, the paper seeks:

- To present an introduction to and overview of the Internet.

- To outline issues emerging from this exponentially growing technology.

Although information about the Internet is widely available, it is found in so many sources of such differing type and technical level as to be difficult to grasp. Further, the scope of the Internet's established and potential impact is so broad as to prohibit comprehensive

[2] In September 1996, the National Science Foundation awarded a $12.4 million grant to the USC Media Center, which will be matched by an additional $30 million in other funding, both corporate and governmental. It is expected that the result will be significant growth in related start-up companies in the area. Work of the Media Center, although not limited to Internet issues and technology, will have implications for the Internet on both the hardware and software sides. See Carol Troy, "USC's Media Center: Revenge of the West Coast Nerds," *New York Times* (online edition), September 27, 1996.

coverage of all the affected areas in a relatively brief paper. Instead, I have attempted to outline major issue areas and to summarize key concerns and policy options.

Almost every week during the research for this paper brought at least one significant new development or has revealed a prediction to be a *reality*, or nearly so, already. Some observers have claimed that time moves differently in cyberspace (the world of computer communications), that one "Web year" is the equivalent of five or even seven "human years." This is rather like the concept of "dog years." Those estimates may well be too conservative, at least in some respects and at the current stage of developments. The phenomenon brings to mind R.A. Lafferty's science fiction-fantasy story "Slow Tuesday Night." That story takes place in a world that moves with lightning speed, allowing financial empires and complex personal relationships to rise and fall in hours. Although I read that story long ago, it often comes back to me as I read daily of the changing world of the Internet.

Some of the issues touched on in this paper are only translations into the cyberspace setting of issues that long predate the Internet. I hope that bringing together many issues in one relatively short paper will give a new view that will help readers to understand the many issues raised directly and indirectly by the Internet and by its growth in content, application, and users. I have entirely omitted some Internet-related topics, such as how to select an Internet service provider, comparisons of Internet browsers, and explanations of Internet-access software. All of those topics are well covered in books listed in the bibliography.[3] Nor have I made any attempt at a comprehensive catalog of online resources, even within the scope of California State Government. There is simply too much available (and too much that is worthwhile) to permit that.

Also omitted from this paper (but planned for a follow-up analysis) is the issue of the telephone companies' "universal service" responsibilities and their potential impact on Internet access. The omission from this paper is in the interest of keeping the paper within manageable proportions and in recognition of the time that will be required to analyze the recent Federal Communications Commission recommended decision on universal service and related California Public Utilities Commission policies and proposals and to evaluate recent developments in the Internet industry in the light of those documents.

Policy options following chapters 2 through 6 are offered as suggestions for consideration only, not recommendations.

Note: A December 1996/January 1997 Nielsen survey estimated that there were then over 50 million Internet users in the United States, or almost one-fourth of all Americans over age 16. This figure was twice that found 18 months earlier. (Reported by Associated Press, March 12, 1997.)

[3] The best single overview volume at this time is Glee Harrah Cady and Pat McGregor, *Mastering the Internet* (San Francisco: Sybex, 1996).

CHAPTER 1: WHAT IS THE INTERNET?

> "[T]here is a never ending quest to have the largest, best, and most munificent network, made up entirely of the work and investment of somebody else, and offer it for sale to the public." (Jack Rickard, editor of *Boardwatch Magazine*, 1996.)

This chapter describes the Internet and the Worldwide Web. It includes an overview, capsule history, explanation of how it works, and basic terminology. The chapter also looks at the kinds of information available on the Internet.

One perspective on the Internet is the relatively tangible *physical* perspective of computer hardware, communications lines, and the software that manages their connections. The other is the *conceptual* perspective of the roles played by the Internet. Both of these perspectives are important to an overall understanding of the technology and its growing role. The next few sections attempt to outline the background and basic ideas from both perspectives.

Background and Buzz-Words

In a nutshell, the Internet is a global *network of computer networks*. This means that a user at a computer on one network can send e-mail (electronic messages) to, or access information on, a computer connected to another network linked into the overall Internet anywhere in the world. Specialized programs and devices that make links, route messages, and keep track of the host computers on the Internet manage the system. These processes happen out of sight of the user. The user need not know anything about how the information is actually transferred or the route(s) taken by the message. Only an occasional error message or the routing information appended to e-mail messages in transit might hint at the hidden workings of the system.

Where did the Internet Come From?[4]

The Internet began as an effort to design a computer communications system that could let messages continue to flow even if a link were disrupted by disaster or attack. The result of research was a packet-switching network. (This and other terms are defined below in this chapter, or in the glossary.) In 1970, the prototype packet-switched network connected UCLA, UC-Santa Barbara, Stanford University, and the University of Utah, Salt Lake City:

> This was the start of the Internet--four universities connected by a packet-switching network that was funded by ARPA [Advanced Research Projects

[4] This section (except for post-1994 developments) is based on *The Internet Unleashed* (Indianapolis, Indiana.: Sams Publishing, 1994), Chapter 1, although I have severely abbreviated the coverage. A more thorough background discussion may be found in Chapter 3 of Ravi Kalakota and Andrew B. Whinston, *Frontiers of Electronic Commerce* (Reading, Mass.: Addison-Wesley, 1996).

> Agency]. If any one link of the network failed, information could still be routed via the remaining links. This satisfied the original criteria for developing a computer network that could withstand hostile attack--or, as it actually turned out, attacks of nature and civil war.[5]

By 1972, the network had expanded to 40 sites connected to what was then called "ARPAnet." Researchers had created ways to send short messages through electronic mail (e-mail), to control one computer over the network from another computer on the network, and to transfer large files between computers through FTP--"file transfer protocol." With these developments, the core technology underlying the Internet had been established.[6] However, the Internet as it is known today was not yet in place.

By the late 1970s, other networks connected research centers that were not part of ARPAnet and the National Science Foundation helped to fund creation of the Computer Science Research Network (Csnet). At the suggestion of Vinton Cerf, who with Robert Kahn had invented the basic Internet communications methods called TCP/IP (Transmission Control Protocol/Internet Protocol), Csnet, as a collection of independent networks, was linked via a "gateway" to ARPAnet:

> It can be argued that this was the real birth of the Internet. Keep in mind that the Internet does not exist as a physical entity. You cannot reach your finger out and touch anything that can be called the Internet. The Internet is a collection of independent, free-standing networks that have come to an agreement about how to talk to one another. That is what Vinton Cerf envisioned when he suggested coupling Csnet to ARPAnet.[7]

More networks quickly appeared, as did new network and inter-network software and techniques. Those techniques included methods for online group discussions and electronic mailing lists on specific subjects. By the late 1980s, the National Science Foundation had launched NSFnet to link selected supercomputer centers across the U.S. Previously, Milnet (Military Network) took over the military-related portions of ARPAnet. By 1990, ARPAnet had given way to NSFnet and terminated. As these and other changes took place in the existing structure, new networks came online, including CompuServe, Prodigy, and America Online, as well as networks established by various businesses primarily to serve their internal needs.

In May of 1993, the National Science Foundation began a process that would end the NSFnet and replace it with a new system centered on a few Network Access Points (NAPs), "where private commercial backbone operators could 'interconnect' much as they had using the NSFNet backbone . . . In February, 1994, NSF announced that three

[5] *Ibid.*, p. 8.
[6] *Ibid.*
[7] *Ibid.*, p. 14.

NAPS would be built . . . On April 30, 1995, the NSFNet backbone was essentially shut down and the NAP architecture became the Internet."[8]

With that transition, the Internet could move quickly toward full commercialization and could complete the abandonment of old standards, long in decline anyway, prohibiting commercial activity on the precursors to the Internet. Those standards limited use to educational, governmental, and nonprofit purposes. From about 1993 to 1995, such commercial services as CompuServe, Prodigy, and America Online opened e-mail gateways, and eventually much more complete access, to the Internet. In essence they became part of the Internet in the process.

The feature that has done the most to propel Internet growth is the ***Worldwide Web***. The Web, or "WWW," is not something separate from Internet. Rather, it is one of the techniques by which information is presented *on* the Internet.

The Cyber Dictionary defines the Web this way:

> The graphical lane on the information superhighway.
>
> It's the hottest trend since the hula hoop. "The Web" offers text, graphics, sound, and video all in place on the Internet.
>
> It is called the Web because its many sites are linked together. You can jump from one site to another by clicking on a highlighted word or "hot button" (a little graphic item).[9]

The WWW was created by researchers at CERN, the European Laboratory for Particle Physics (the acronym comes from its name in French). There Tim Berners-Lee and colleagues, in 1989, developed the concepts and basic methods of the Web as a means of sharing research. They had no idea that sound, video, or graphics would be included in the system. By 1991, the Web had become part of the Internet suite of services. After that:

> The project quickly expanded beyond all imagination as others understood the potential for global information sharing. Hundreds of people throughout the world have contributed by writing and modifying Web software and documents. In a way never envisioned by the original project group, the project reached global proportions by the middle of 1993 with the introduction of the NCSA [National Center for Supercomputer Application]

[8] Jack Rickard, "Internet Architecture," http://www.boardwatch.com/isp/archit.htm, accessed August 29, 1996, pp. 3-4.
[9] David Morse, ed., *Cyber Dictionary: Your Guide to the Wired World* (Santa Monica, California: Knowledge Exchange, 1996), p. 304.

> Mosaic--a multimedia front-end [easy-to-use interface] to all the information served by the Web.[10]

Hypertext documents include embedded links, enabling a user to navigate from one part of a document to another or from one document to another (a process called "hyperlinking"), even if the documents reside in different computer files or even on different computers far distant from one another. In a sense, hypertext and the Worldwide Web make all of the information on the Internet into a vast interconnected library.

We will return to the subject of hyperlinking after looking into some terminology.

NAPs, POPs, and Backbones

Although the purpose of this paper is not to provide a technical explanation of networking or internetworking, a brief explanation of some of the unavoidable buzz-words of the Internet will still be useful. Some of the terms that come up again and again in any discussion of the Internet follow. Additional terms are defined in the glossary.

A ***domain*** is the name and type of site on the Internet. The domain is part of every Internet address. There are many domains, and each in turn is divided into subdomains. At the top level are edu (education), gov (government agencies within the U.S.), org (non-profit organizations), com (commercial organizations), net (networking organizations), mil (U.S. military), and more than a hundred national domains, such as ca (Canada), us (United States), and nz (New Zealand).

For example, the U.S. Government Printing Office is on the Worldwide Web at www.access.gpo.gov. This indicates, reading from right to left, the (U.S.) government domain, Government Printing Office, host computer called "access," Web page. (It is becoming customary, but is not necessary, to use "www" as the first part of a domain name for a Worldwide Web site on the Internet.) ***Domain name servers*** (***DNSs***) are computers connected to the Internet that associate domain names with the numerical ***IP addresses*** used by the system, so that users can remember something simpler than a string of dot-separated numerals such as "133.34.111.87." IP addresses uniquely identify a specific host computer on the Internet.[11] DNS can also mean "domain name system," the system of domain-naming rules and the database (organized list) of assigned names.

The concept of the ***backbone*** is fundamental to the Internet. Internet backbones are very high-capacity communications lines connecting major cities across the country.[12] It might help to make use here of the concept of the "Information Superhighway," so often used in reference to the Internet. Think of the backbones as the equivalent of hundred-lane, high

[10] Ravi Kalakota and Andrew B. Whinston, *Frontiers of Electronic Commerce* (Reading, Massachusetts: Addison-Wesley, 1996), p. 228.
[11] To be more precise, the IP address "expresses the exact physical connection between a computer and the network on the Internet." Cady and McGregor, *Mastering the Internet*, p. 78.
[12] This discussion will be limited to the U.S.

speed freeways linking a handful of major metropolitan areas. In turn, smaller freeways (perhaps twenty lanes) connect smaller urban areas to these "backbone" freeways. Still smaller freeways connect still smaller cities and towns to the twenty-lane freeways.

Eventually, we reach the level of a single-lane road that allows an exit to your driveway. That is the equivalent of the home telephone line. Little roads (communications lines) connect to bigger ones, and those to bigger (and faster) ones, culminating at the backbone level. This, in turn, attaches to a Network Access Point and then works back down to the driveway level again. At every step of the way, connecting devices--think of on-ramps, off-ramps, and interchanges--must attach one road to another. The metaphor is probably not worth working out in too much detail, but it still helps to show how the network works.

The "little roads" that go right to the individual home are sometimes called the ***last mile***. Those connections, whether by way of the ***telco*** (telephone company), cable company, or even electric utility company, involve a huge investment that is difficult and expensive to upgrade or replicate. Any entity with a last mile system has leverage in the electronic communications game.

An article in *Boardwatch* magazine (a computer/communications industry publication) summarized the backbone concept this way:

> When we refer to a "national Internet backbone provider," we are describing a company that has physically located a high-speed TCP/IP [Transmission Control Protocol/Internet Protocol--an Internet communications standard] router in a number of cities, and then leased high-speed data lines from long distance exchange carriers to link the routers--thus forming a national "backbone" connecting those cities. By doing so, they can then sell access to many individuals and companies within each backbone city, and the traffic between cities moves over the leased lines of the backbone. The leased lines can actually be purchased from different long distance companies for each city, and in fact some backbone providers will lease several lines from different carriers to connect two cities so that if one carrier happens to have a backbone incident [that is, crashes] and the link is lost, they still have a connection through the other carrier--maintaining their backbone at a perhaps reduced data rate. This is termed "redundancy."[13]

There is not **a** single backbone to the Internet, nor is there a single communications company providing or controlling access. Relationships and interconnections are complex and sometimes quite strange. *Boardwatch*'s Jack Rickard remarked, "Deals are cut between operators at and across all levels of the Internet to the point of frenzied

[13] Rickard, "Internet Architecture," page 7.

confusion." He went on to surmise that "some operators are agreeing to buy things from themselves through several third parties in some cases."[14]

A ***NAP*** is a "network access point," a physical point where backbone operators connect to the Internet. "Currently," writes *Boardwatch*'s Jack Rickard, "the heart of the Internet remains the three 'official' network access points or NAPs in San Francisco, Chicago, and Pennsauken, New Jersey, near Philadelphia. (The latter is referred to as the New York NAP.) They establish the concept that interconnection is good, and that at [at] least these three points, anyone can in theory interconnect with the rest of the Internet."[15] "In this way," he writes, "anyone could develop a national backbone for the connection of LANs [local area networks], sell connectivity to it, and use the NAP as the physical point where they interconnected and exchanged traffic with all the other service providers."[16] Other "major interconnections" cited by Rickard bring the total to eleven. Major providers connect to more than one, and possibly to most, of these points.

Figure 1 very broadly illustrates the layout. Do not take the connections literally, as they are intended only to give a general impression of the top level of the Internet.

Figure 1

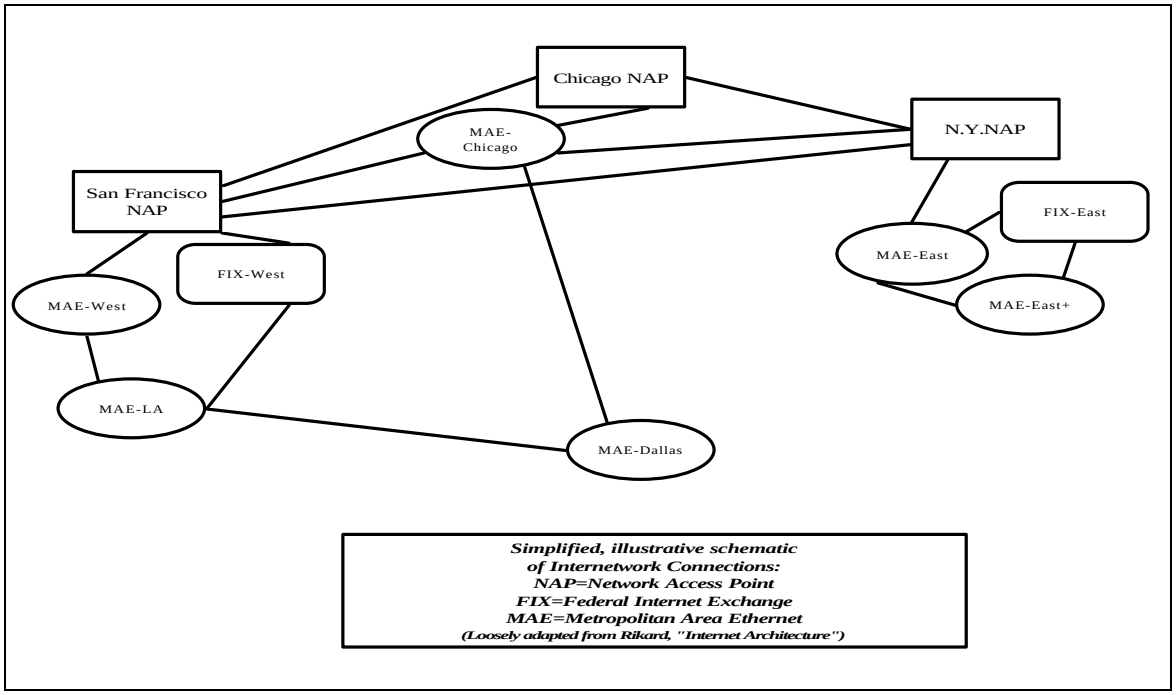

[14] *Ibid*., page 8.
[15] *Ibid*., page 4.
[16] *Ibid*., page 3. Elsewhere Rickard has speculated that the concept of national backbones and the role of NAPs could essentially disappear over time as local ISPs interconnect, moving to "hundreds and thousands of connections at lower levels in the hierarchy." Rikard, "Editor's Notes," *Boardwatch*, September 1996.

A ***POP*** is a more local version of a NAP. Again quoting Rickard: "Generally, we refer to all the nodes of the network owned by the national service provider as Points of Presence or POPS. Business customers then lease their own telephone line from the telco to this POP and so connect to the Internet."[17] That is, the business leases a line that connects to a POP, and the provider of the POP connects to a NAP, and once the connection is made to the NAP, the Internet has been accessed. It is almost as simple as NAP, crackle, and POP.

A typical home user of the Internet dials up an ***ISP*** (Internet Service Provider) through a modem in or attached to his personal computer. The user then has a telephone line connection to a vendor, who in turn connects to the Internet. ISPs (also known as Internet access providers) charge a fee (usually) for the user to connect, via them, to the ISP's Internet connection. In short, the home (or small business) user of the Internet connects via an ISP, which is in effect the retailer of Internet services. If an internet service provider (Netcom, for example) can be accessed by a phone call to a specific town, then that provider is said to have a point of presence--a POP--in that town.

Think of it this way: an enormous information freeway (Figure 2) connects San Francisco, Chicago, and New York (well, Pennsauken). That route is the heart of the system. Anyone who wants to can build a backbone (major information freeway) and connect it to one of those three points (NAPs), and in turn can have smaller connecting roads splitting off into more and more remote areas. Still smaller systems can connect to *those* systems. If you ("Mr. A," let's say) connect to any of the systems that ultimately connects to a NAP, then you can reach any point ("Ms. B," for example) connected to any other system likewise tied to a NAP. A system that does not somehow ultimately connect to a NAP is not guaranteed to have access to all of the Internet.

Routers are devices (encompassing both hardware and software) that forward messages from one network to another--that is, route them to their destination. Because the Internet encompasses thousands of networks, many routers are required to direct data from source to destination.

Packet switching refers to the method by which messages are ordinarily sent along the Internet. In short, a message is broken into packets (chunks, each with identifying information), which are relayed to the destination. The multiple packets into which one message might be separated will probably take different routes to their destination, and perhaps not even arrive in the order they were sent. The identifying information accompanying each packet allows their reassembly in the proper order at their destination. If a packet is lost en route, the originating computer is notified and resends it. This technology means that there is not a continuous, always-open connection from source to destination. Packet-switching facilitates efficient use of the lines, which can be filled with packets from and to countless destinations.

[17] Rickard, "Internet Architecture," page 7. Note that this kind of POP is not the POP that refers to "Post Office Protocol," a means for handling e-mail. A "node" is a point where equipment connects to a communications line.

Figure 2

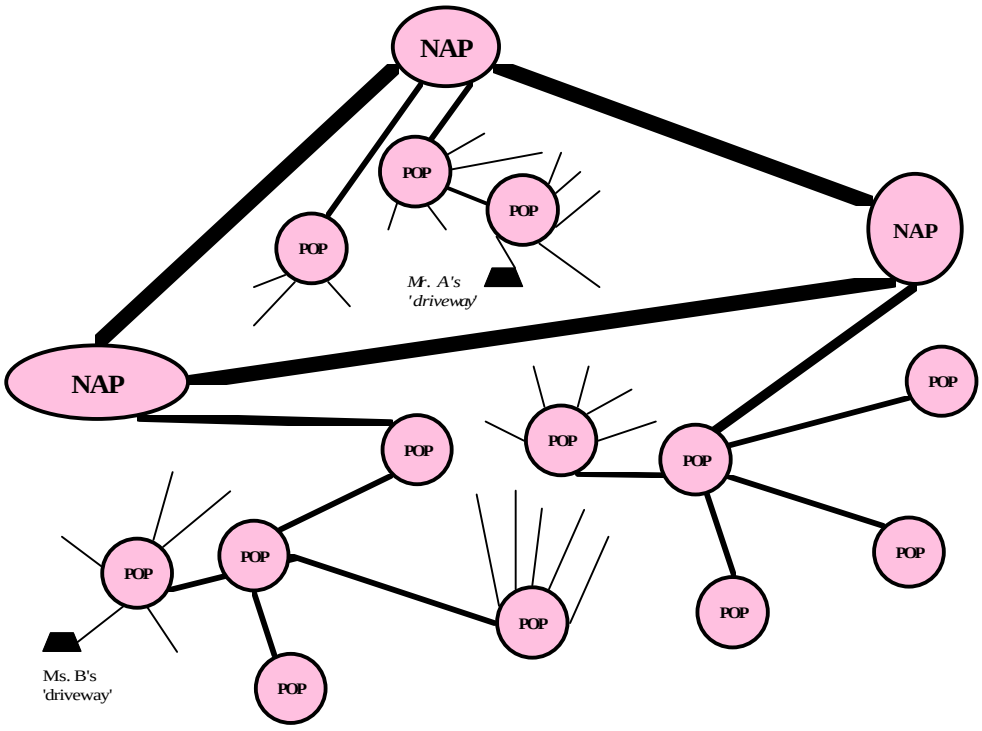

HERE A POP, THERE A POP . . .

Picture a busload of people (one large message) being delivered across country in dozens or hundreds of pieces (packets) by way of dozens of highways, selected on the fly depending on what is available and efficient at that instant, intermingled along the way with packets broken out from countless other buses, trucks, and automobiles, and reassembled at its destination, all in the blink of an eye. That is packet switching. It requires a lot of hardware and software to make it work, but is a very efficient way to use the telecommunications lines.

The system is also ***connectionless***, as there is not a continuously open connection from source to destination. That contrasts to a telephone call, during which a connection is made and kept open for two-way communications until the call ends. (Now you know what a *connectionless packet-switched network* is.)

Protocols are the "standard operating procedures" of networks. Protocols specify the technical details of how the intricate tasks of network communications are to be handled. Fundamental to the Internet is ***TCP/IP***, or Transmission Control Protocol/Internet Protocol, the instructions governing how packets are sent on the Internet and how errors are handled.

Bandwidth is the "size of the information pipeline," so to speak. It is a measure of how much data can be delivered across the line how fast, usually expressed in thousands or millions (or even billions or trillions) of bits per second. A bit is a single binary digit, a 1 or 0 ("on" or "off"). All digital data, however complex, are at bottom simply a stream of bits--zeros and ones. More bandwidth equals more capacity to transmit information within a given period of time.

In highway terms, a 100-lane freeway has ten times the bandwidth of a 10-lane freeway. In telecommunications terms, a T3 line, at 44.736 Mbits/s (megabits per second) has 28 times the bandwidth of a T1 line, at 1.544 Mbits/s.[18] Bandwidth has nothing to do with how fast bits move. They move at the speed of light *whatever* the capacity of the line. Fiberoptic cable (a highly polished glass or plastic cable that transmits light rather than electrons) has vast bandwidth, but not because the actual transmission speed is faster. Light (on fiberoptic cable) and electrons (on phone lines) both move at the same speed. The limiting factor in data transmissions is how quickly the medium can change signals.[19] To illustrate: in theory, digital data could be sent via smoke signals, but the bandwidth would be measured in bits per minute, not thousands or millions of bits per second because of the time needed to produce smoke puffs and to interpret them for retransmission along the next link. The light transmitting the puffs would, of course, move *at* the speed of light.

Three more terms often come up in discussions of bandwidth, **narrow-band**, **midband**, and **broadband**:

> Cables with limited bandwidth, used for text or voice transmissions, are called narrow-band circuit. Cables with more capacity, which carry images and limited animation, are "midband capable." Those with a high bandwidth, which can carry multiple video and audio signals, are said to have broadband capacity.[20]

Congestion occurs wherever the bandwidth is insufficient to handle the load at that point. The ***telco*** (telephone company) line into your home, designed for ***POTS*** ("Plain Old Telephone Service") has low bandwidth. By the time the Information Superhighway has gotten to your monitor, it has shrunk to an Information Bicycle Path (Figure 3). That fleet of 18-wheelers you ordered up on the Web will end up coming your way as a seemingly

[18] Yes, I know, if you do the arithmetic the ratio is much closer to 29:1. Nonetheless, a T3 line "supports 28 T1s," according to the *Dictionary of PC Hardware and Data Communications Terms*. A T4, in turn, carries six times the bandwidth of a T3. Backbones require high bandwidth, and therefore are moving toward *gigabit per second* (billion bit per second) levels. There is even talk of terabit levels (trillion bits per second). The potential traffic on higher bandwidths now under development, however, cannot be accommodated by current switching and routing technologies, so advances must take place across a broad range of technology. See Kalakota and Whinston, pp. 110-115.

[19] Douglas E. Comer, *Computer Networks and Internets* (Upper Saddle River, N.J.: Prentice Hall, 1997), pp. 21-22.

[20] Bill Gates, *The Road Ahead* (N.Y.: Viking, 1995), p. 31.

endless parade of bike-sized packets, virtual handle-bar tassels flapping in the electronic wind. Every one of those packets will follow an individually-scripted route designed on the fly, one router at a time.

In a system this complex, there are many places along the way where things may go wrong or where congestion may occur. A router might be misbehaving, DNS might be swamped, or the particular site one is trying to reach might be overwhelmed by many requests at the same time. Internet error messages are often ambiguous or misleading, but often mean the same thing: *it is too crowded right now somewhere along the way*.

Figure 3

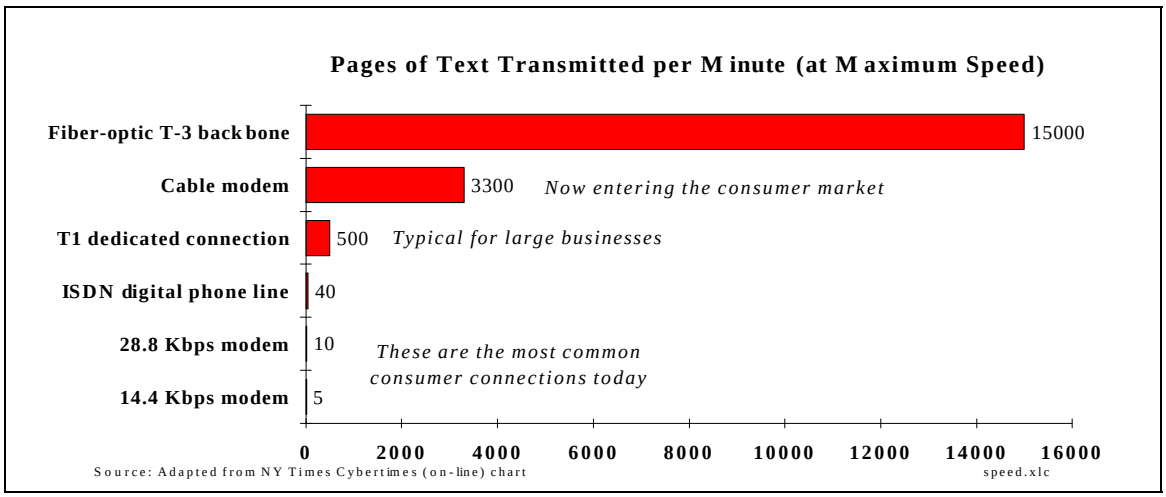

No One Controls the Internet

Controlling the Internet is almost as alien a concept as herding cats. It is just not in the nature of the beast to be run from a central location or organization. (But each individual *network* connected to the Internet does have its own administration.) Nonetheless, technical standards and various coordinative activities are the responsibility of specific organizations and are essential to the success of the Internet. In time, the Internet cats find that a particular direction is good, and they follow it (Figure 4).

Sharon Eisner Gillett puts it this way:

> Contrary to its portrayal as total anarchy, the Internet is actually managed. It runs like a decentralized organization, but without a single person or organization filling the manager's role. The system that allows 99% of day-to-day operations to be coordinated without a central authority is embedded in the technical design of the Internet. The manager's job--

handling the exceptional 1%--is performed by not one but several organizations.[21]

Figure 4

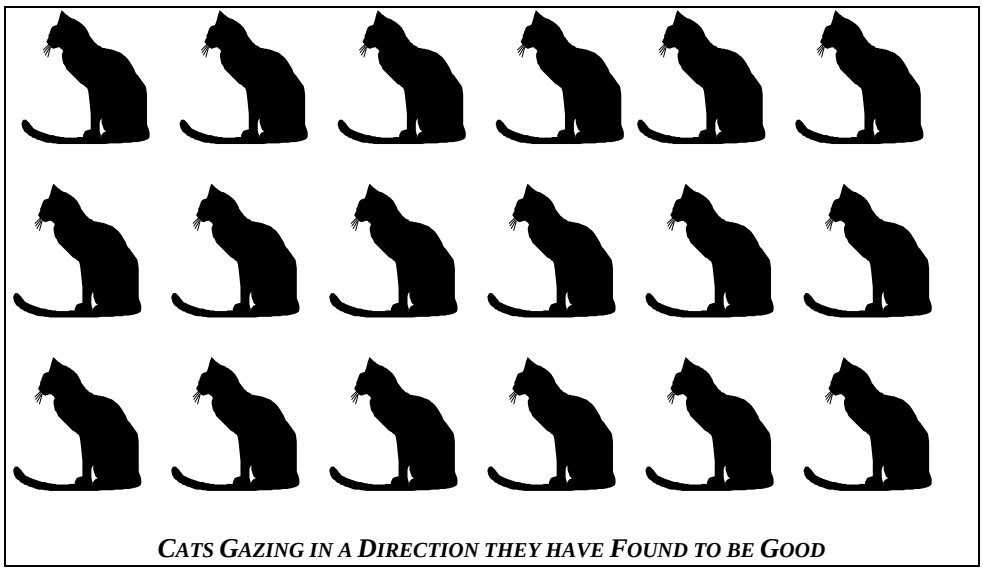

CATS GAZING IN A DIRECTION THEY HAVE FOUND TO BE GOOD

Internet organizations include:[22]

- InterNIC Registration Services, which registers domains, names, and addresses of computers connected to the Internet in the top-level domains *edu*, *gov*, *com*, *net*, and *org*. InterNIC charges a fee for this service. (Other registries handle other top-level domains, under the overall direction of Network Solutions, Inc.)

- The Internet Society (ISOC), an international group of volunteers, founded in 1982 in Reston, Virginia. Membership is open to all interested persons, although there is a fee. The Internet Society coordinates work of various other groups associated with Internet policy and practice.

- The North American Network Operators Group (NANOG), which handles North American network operations.

- Computer Emergency Response Teams (CERT), addresses security matters, coordinated by Forum of Incident Response and Security Teams (FIRST).

[21] "The Self-Governing Internet: Coordination by Design," prepared for Coordination and Administration of the Internet, Workshop at Kennedy School of Government, Harvard University, September 8-10, 1996. Posted at http://ccs.mit.edu/ccswp197.html. The quoted passage is from page 1 as printed.
[22] This list has been selected and adapted from *The Internet Unleashed*, *Mastering the Internet*, and *The Cyber Dictionary*, with an occasional glance at other sources.

- Internet Architecture Board (formerly Internet Activities Board), IAB, a technical group overseeing the development of the suite of Internet protocols (standard methods). IAB oversees the Internet Engineering Task Force (IETF) and Internet Research Task Force (IRTF). The latter is more theory-oriented than the former.

- The World Wide Web Consortium, W3C, "an industry consortium which seeks to promote standards for the evolution of the Web and interoperability between WWW products by producing specifications and reference software." This is strictly an industry group; in contrast to the Internet Society, there is no individual membership. Its focus is the Web, not the Internet generally. The group notes, "W3C . . . is vendor-neutral, and its products are freely available to all," despite its industrial member funding.[23]

"Standards on the Internet," states Ron Dippold in *The Internet Unleashed*, "are created in a wonderfully anarchic manner. Any person (or group) who thinks he has a good idea prepares an RFC (Request for Comments) detailing his proposal. If other people think the idea is good, it will be implemented. If not, it will languish, unused."[24] RFCs are posted online (InterNIC assigns numbers to them), and may be retrieved by anyone who is interested. They are also published in CD-ROM format.

Anarchic or not, the Internet as a whole is clearly not centrally controlled, but it *does* have a management system that develops, publicizes, evaluates, and implements standards in critical technical areas.

Navigation on the Net: Hyperlinking

A little further on, we will look at the kinds of information available on the Internet. For now, though, let's just acknowledge that many, many kinds of information are online, provided by an enormous array of individuals and organizations. Information resides on millions of computers linked to the Internet and communicating as described above.

A user of the Internet can visit a specific site online (or many of them, one after another) by typing in each site's address. By this means it would be possible, but awkward and slow, to go to a series of related sites, one at a time, each time entering the Internet address found in, say, *The Internet Yellow Pages* or some other reference.

Hyperlinks enable the user to select an embedded link (by clicking with a mouse or moving the cursor to the link and pressing the enter key) and in that way to go to the address the link points to *automatically*. This process can continue from one site to

[23] "About the World Wide Web Consortium," posted at http://www.w3.org/pub/WWW/Consortium/.
[24] *The Internet Unleashed* (Indianapolis, Indiana: Sams Publishing, 1994), p. 211. For another view of the management of the Internet, see Sharon Eisner Gillett, "The Self-Governing Internet: Coordination by Design," prepared for Coordination and Administration of the Internet, Workshop at Kennedy School of Government, Harvard University, September 8-10, 1996 (posted at http://ccs.mit.edu/ccswp197.html).

another, on and on and on. The user can move back and forth among links in this way, crafting a path meeting the immediate needs and interests.

Picture a student in a library reading room, at a table with a book--say, *Moby Dick*, by Herman Melville. If she wished to learn more about Melville, she would have to get up and go to a catalog to find a book on Melville, and then get that book, or would have to look him up in an encyclopedia. To find out more about the America of Melville's day, she would have to get another book or turn to another entry in an encyclopedia (maybe many entries). She might also search for magazine and journal articles, photographs, paintings, and other types of information source.

To find other books available from the publisher of this edition of *Moby Dick* (The Library of America, for example) would require writing to request a catalog. These traditional manual methods work, but are time consuming. Much of that time is taken up in simply moving around and physically retrieving documents--or in finding out where to get them.

Now imagine that student seated instead at a computer connected to the Internet, viewing an online copy of *Moby Dick* (Figure 5). If that version of the book includes hyperlinks (for the sake of discussion let's say that it does), the student could simply click on *Herman Melville* and the browser would go out on the Net and fetch the referenced biography of the author. Links embedded there might enable her likewise to jump to information on New York City (Melville's birthplace), Liverpool (where he went in 1837 as a cabin boy), or to a list of his other works (*Typee*, *Omoo*, and others). She might then jump to *Omoo* to see whether that is another book she wants to read. Similarly, she might jump to the publisher's Web site to examine its catalog, to a gallery of photographs of Melville, to some video clips about whales and whaling, or to any other topic for which links were provided. If the links were not sufficient, she could search for additional information via search engines or online libraries. All of this information would come via the Internet.[25]

Some critics, such as Sven Birkerts, cringe at the effect of this sort of information grazing on how people will learn, whether online or via a multi-media CD-ROM. They fear a disappearance of reading in depth and its replacement by endless lateral branching out. Others, however, see hyperlinking through the Internet and multi-media CD-ROMs as a means of learning that will broaden interests and expand knowledge by removing the barriers that now interfere with exploration of knowledge.

One thing is certain: the ability to hyperlink through knowledge is limited only by what is online. As the already astonishing scope of information online continues to expand, so too will the ability to follow unique paths of interest, whether to meet professional needs (technical, legal, or business information) or personal ones (such as health information, recipes, or world news).

[25] The same could be achieved on a CD-ROM. The difference is that the CD-ROM's content is static, limited to whatever was put on the disk, while the Web's content is dynamic and enables additional, *ad hoc* levels of linking and keyword searching for additional resources.

Hyperlinking through the Web is a remarkable process, and one that comes as a surprise to those who first encounter it. One manager in a State department had to advise his employees, recently provided Internet access for official use, not to lose months of work to the phenomenon, as "the Web can be very seductive."

Figure 5

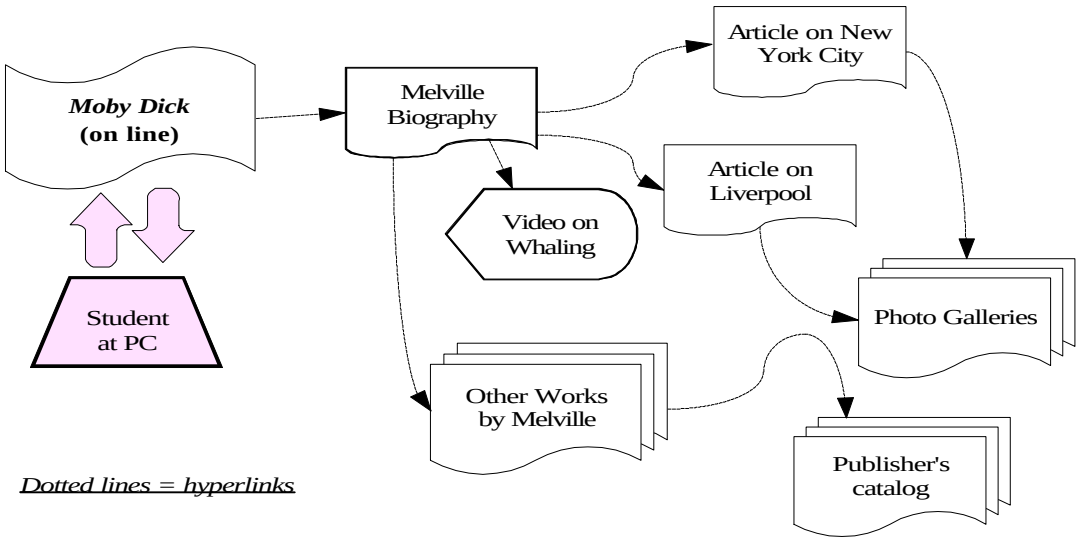

HYPERLINKING THROUGH A RESEARCH ASSIGNMENT

The Internet is Interactive

The Internet is an *interactive* medium. Although a user can simply look for information and read it online or download it to a personal computer, many users also post messages, chat online, and reply in various ways to what they find online. This makes the Internet fundamentally different from one-way communications media such as radio and television broadcasts and printed publications. Even the search for information on the Internet requires some interaction, as the user must navigate to a search engine, enter the topics desired, and select sites to visit. The user is sending instructions that result in the selected pages appearing on the screen. At the same time, the site the user is visiting may keep a record of the visit and even log certain information about the user, sometimes on the user's own hard disk. (This topic is discussed further in Chapter 5, below.)

The Internet is both Global and Local

One of the Internet's strengths at the same time adds to the complexity of issues it raises: it is *international*. Users from opposite parts of the globe can communicate as easily as can users who live next door to one another. The primary limitation on their ability to do so is in their own neighborhoods, reflecting the local communications infrastructure and

bottlenecks in national systems. Those systems vary *widely*, from virtually non-existent in many developing nations to advanced and broadly accessible systems in industrialized nations. Finland, for example, is regarded as one of the most connected countries on Earth.

Kalakota and Whinston estimated that over 140 countries have access to the Internet, and note that "the Internet has always been an international network." Each of the connected countries has its own approach to the Internet:

> Asian countries see the Internet as [a] way of expanding business and trade. Eastern European countries, longing for western scientific ties, have long wanted to participate but were excluded by government regulation. Since this ban was relaxed, development is progressing rapidly. Third World countries that formerly did not have the means to participate now view the Internet as a way to raise their education and technology levels.[26]

The Internet creates an international community by letting people in all of the connected countries talk to one another via e-mail and data files, and now even by live audio and video. This, then, lends an international aspect to matters that otherwise would be national or local. Issues related to the international nature of the Internet are touched on elsewhere in this paper.

Although the Internet is global, it may also serve local purposes. In a very loose sense, the Internet is like the daily newspaper. The newspaper includes articles about events in distant countries and it includes articles and advertising about large, international companies. But the paper also includes local news, advertising for local merchants, and classified advertising placed by businesses, organizations, and individuals in the local area. The reader may focus on the international coverage or may focus on purely local information.

Likewise, with the aid of Web sites that organize links to local information and that provide locally oriented content, an Internet user may focus on a local area. This sort of local use of the Internet is growing, and is the focus of competition between local interests and commercial entities that are seeking to develop many city and regional sites.

Microsoft, for example, has announced plans to develop city-oriented sites, a feature called "Cityscape." CitySearch, already established in a few cities, is developing similar plans nationwide, and Yahoo! (a search engine company) is also entering the field.[27] At the same time, however, locally-based interests are developing local-interest Web sites. The *Knoxville News-Sentinel* has such a site serving its area, with links to local event information, merchants, and news.[28] A local Web site serving Taos, New Mexico, offers

[26] *Frontiers of Electronic Commerce*, p. 115.
[27] Steven Levy, "Rise of the City Sites," *Newsweek*, September 30, 1996, p. 86.
[28] http://www.knoxnews.com/.

extensive information on local business, culture, events, and institutions.[29] Hyperlinking makes development of local sites simple, and as more and more businesses, organizations, and agencies are on the Web this is an area that will grow quickly.

Fast-Growing Medium for Communications and Learning

The Internet is a channel for communication and learning. A woman posted this comment to a philosophy-oriented mailing list on the Internet:

> When I first got this computer, I didn't even know where the on switch was. I was completely computer illiterate. But I KNEW that this machine could do wonderful magic things. I knew it could open up the world to me. If I just knew HOW to ask it to do it. When I finally learned a few of the secrets of this magic machine I was not disappointed in the wonders available to me.

This is a strong statement: *"I knew it could open up the world to me . . . I was not disappointed in the wonders available . . . ."* Anyone thinking about the Internet should bear in mind that beyond the computers, routers, bridges, T1 lines, and a host of arcane and acronymic devices lies a vast and multi-faceted community of *people*. As that community expands, so too does the potential power of the medium.

Indeed that community *is* expanding, and with an accelerating speed that almost defies comprehension (Figure 6). This is one reason why the Internet is such an important topic. It shows no sign of slowing or even leveling off in its growth. Even though its current rate of growth cannot continue forever, it need not continue for much longer before essentially saturating the potential market in the industrialized countries, especially the United States.

Figure 6

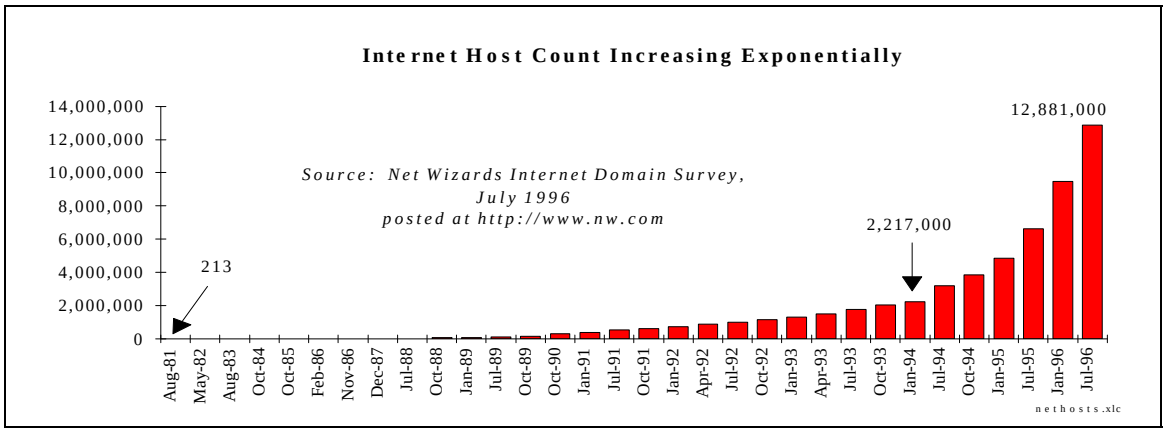

[29] http://www.laplaza.org/.

Other indicators all show the same kind of trend: accelerating growth with no sign yet of leveling off. Estimates typically indicate that the Internet is doubling or more *annually*, by whatever measure one chooses. The Worldwide Web aspect is growing even more rapidly, climbing from little above zero "WWW-Prefixed Hosts" in mid-1994 to about 60,000 in December 1995, 150,000 in April 1996, and about 220,000 in July 1996.[30] Eventually the rate of increase has to level off, but when that will happen is anyone's guess. New methods of accessing the Internet and breakthroughs in ease-of-use could support the growth trend for years to come.

What Can You Find Online?

The United States District Court for the Eastern District of Pennsylvania found, "It is no exaggeration to conclude that the content on the Internet is as diverse as human thought."[31] For that reason, it might be more reasonable to ask what you *cannot* find online than to ask what you *can*. That list is becoming shorter by the day. When I started this section, the list appeared to include text of current, copyrighted books, the equivalent of "white page" listings of all Internet users, and . . . well, that seemed to be all. Within a week, both of those exceptions had to be deleted: announcements had appeared for online access to published books (including a way to pay online for viewing or downloading the contents) and for a white pages system (far from complete, but a start--and competitors have already appeared).

Plumbing the Depths

In a test early in September, 1996, I picked out of thin air a topic that seemed unlikely to be well represented on the Internet: *plumbing*. After about two minutes spent accessing a search engine (Altavista),[32] entering the search term "plumbing," and scanning the results, I was led to a 13-page illustrated paper on "The History of Plumbing in America." A moment later I moved to the home page housing that paper and found I could also choose from "The History of Plumbing in Babylonia," or Crete, or Egypt , or any of several other civilizations. Other choices on that page included plumbing supply companies, plumbing (and allied) contractors, an article on "Diane Sawyer on Plumbers," and too much more to list here. There was even "Toiletology 101: A Complete Course in Toilet Repairs"! This is pretty good success for the most unlikely of topics.

What else can you find online? Here is a sample, in no particular order:

- Magazines and newspapers (literally hundreds, if not thousands).

[30] In Network Wizards' definition, a "host," as used in the Internet Domain Survey, is a computer with an IP address and a connection (full-time or part-time, direct or dial-up) to the Internet. There are other, less inclusive definitions of "host." The WWW-Prefixed Host numbers are estimated from Network Wizards chart at http://www.genmagic.com/Internet/Trends/slide-6.html.
[31] Preliminary injunction issued June 11, 1996, in the case of *ACLU vs. Reno*, challenging the Communications Decency Act, in finding of fact no. 74 (text posted on line).
[32] Search engines are explained in the next section.

- Museums and scientific exhibits, including the Exploratorium (San Francisco) and the Monterey Bay Aquarium.
- Books, in electronic text (e-text) form, including classics in English and many other languages.
- Library catalogs, including those for the University of California and the Library of Congress.
- Weather information.
- Census data.
- State documents and information (California and other states).
- Businesses and business organizations.
- Commercial catalogs (retail, wholesale, services, and products).
- Agricultural data and advice.
- Special interest discussion groups on every topic imaginable.
- Job information, including job-search tips and job listings.
- College classes.
- Election data and political campaign information.
- Presidential and Congressional documents.
- Photographic archives and sound clip archives.
- Electronic mail ("e-mail") access to people all over the world.
- Real-time conversation ("chat"), likewise with people all over the world.
- Financial information.
- World, national, local, and special-interest news.
- Search engines for finding specific information on the Internet.
- Humor, parodies, and satires.

More appears online daily, added by individuals who do so as a hobby, by public and nonprofit organizations doing so as a public service, and by businesses seeking to make money by providing information or services or by selling goods. Of course, much of what is online is not worth the time it takes to look at it, just as much of what appears in any other medium is of little value.

Some Web sites provide customized information. For example, some travel-oriented sites provide local maps and even spell out street-by-street routes for getting from one place to another that you specify.

As more and more people see what is already online and find useful information and services, demand may grow for other information and services not already available, although it is becoming difficult to think of any type of information not already online in some form or to some degree.

Search Engines and the Hunt for Internet Information

The enormous number of sites on the Web and documents on Internet would be impossible to use without the aid of "search engines," services that find information at the user's request. Many sites offer search services. These include Yahoo!, Altavista, and Lycos. The user enters a word or phrase, submits the request, and after a while (seconds or minutes) receives a list of sites matching some or all of the search terms. Searches can be simple or complex. Each engine offers its own methods and options, and not all will find the same sites or find sites in the same order. Some sites offer subject-area guides to help the Web-surfer start through a topic without having to enter specific search terms. Many Web sites, especially those for libraries, include links to one or more search engines.

Unfortunately, search engines are not (yet, at least) very good at filtering what they find in order to present the most relevant or useful documents. The user may have to scroll through many trivial or off-the-mark links before finding one that is useful. This is a result of the ability of the search engines to find and index millions of documents on the Internet. Nonetheless, search engines reduce an impossible task--finding and searching among millions of sites--to the merely difficult task of examining a list of perhaps 200 links, many of which can immediately be dismissed as not helpful.

Information Hierarchies on the Web

One of the most striking and useful developments on the Web is the evolution of link pages that serve as tree-structured directories of information. Hyperlinking allows sites to be indirectly joined in many ways, reflecting different ways of viewing the information in and purposes of the sites. For example:

- Local government link pages ease access to local government sites across the country.

- City and regional pages join a variety of sites serving their respective geographic areas--businesses, schools, newspapers, radio and television stations, non-profit organizations, churches, and government agencies.

- Topical pages gather links to sites around the world concerned with a particular topic--for example, technology, insects, oceanography, movies, broadcasting, or social services.

Such link-oriented sites are themselves in turn gathered into more inclusive Web sites, creating a branching tree structure that starts at a higher, more general level. For example, a site might bring together links to all known city or regional Web link sites, so that the user could browse among listings of cities and regions across the country, if not around the world. In turn, even *higher* level sites could offer links to *those* sites. (These are not speculations. Such Web sites do exist.)

The ultimate site may be The WWW Virtual Library,[33] a catalog of subject-based links (when printed, filling four single-spaced pages, as of September 24, 1996), to sites literally from A (Aboriginal Studies and Agriculture, among others) to Z (Zoos). This site alone contains links to more information than any one person could absorb in a lifetime.

Frankly, this all sounds very confusing until one looks around at the choices on the Web. It becomes much clearer with experience.

Active Agents can Aid Internet Users

"Active agents" are computer programs that act on behalf of the user to find, filter, and deliver information. This is an area that, like so much else in the world of Internet, is still in its infancy, and like any infant may offer many surprises as it matures. Considering the overblown and largely disappointed expectations for "artificial intelligence" a few years ago, hopes for active agents should be restrained. They will not work miracles, but they can do useful work by easing, automating, and expediting tasks online.

A simple example of an agent is the Amazon Bookstore notification service. A customer may sign up for notification of new books on specified topics (for example, African History, Chemical Engineering, or Electronic Commerce). When a book is added to the store's catalog in the specified subject, e-mail goes out to all those who have requested notification, advising of the new arrival. In this way, the customer need not keep checking back for new listings, nor go though complete listings to look for new arrivals.

An agent might search the Web daily to find new articles mentioning an individual, company, or topic (much like a clipping service), arrange the articles according to some criterion of relevance, and send the text to the user for later reading or printing off line. Or an agent might be assigned to search commercial sites for the best price on Size 7

[33] http://www.w3.org/pub/DataSources/bySubject/Overview.html.

Widgets, educational sites for college courses in C programming, or government and legal resource sites for laws and regulations on driver's license renewal. Such actions would combine the processing ability and speed of the computer with the access to information provided by the Internet.

New Internet Access Methods

Internet-access technology is rapidly developing and cross-breeding with other technologies (Figure 7). The various access methods are being designed to meet needs and preferences of market segments not necessarily well served by the familiar PC, modem, and SLIP/PPP connection. They are also being designed to use existing infrastructures (such as cable systems) and devices (such as TV sets and video game machines).

Figure 7

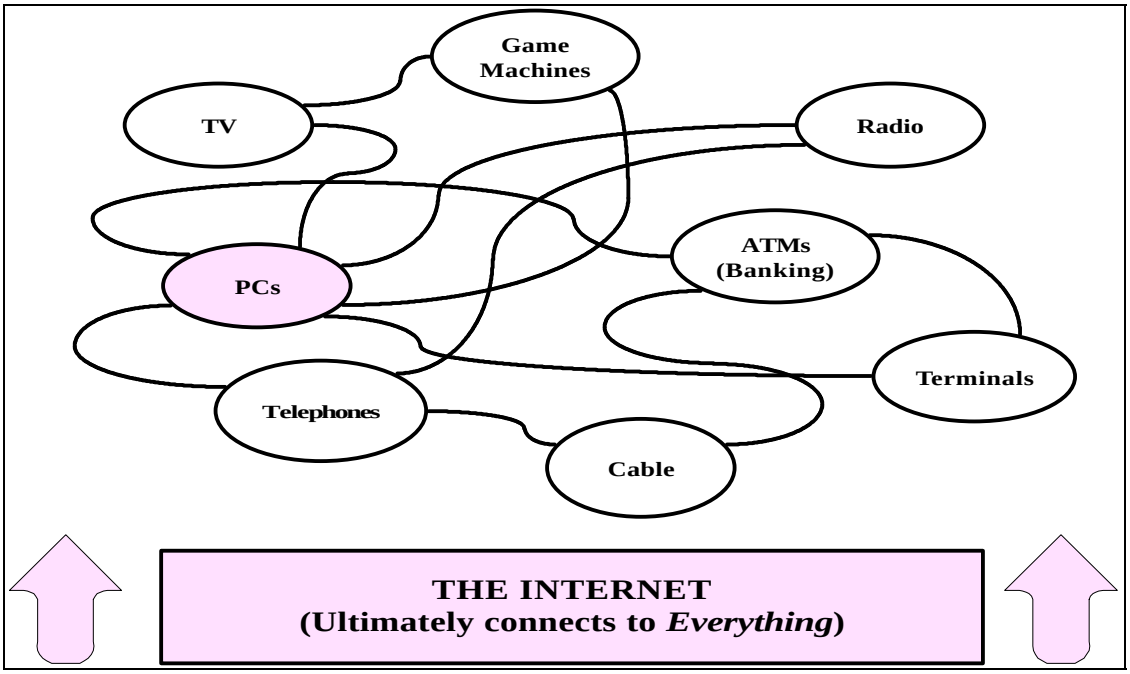

"WebTV" Might Make Mouse Potatoes of Us All

WebTV [34] is an easy-to-use combination of Web access and television introduced in 1996. Other systems of this type are likely to follow,[35] and some new television sets will come

[34] There is a dispute over the ownership of the trademark "WebTV." I use the term to refer to the products and services of a partnership headquartered in Palo Alto, California, having licensing agreements with Sony Electronics, Inc., and Philips Consumer Electronics Company. Information on the partnership and its products may be found at http://www.webtv.net. The other claimant of which I am aware is on the Web at http://www.webtv.com. I neither express nor imply an opinion as to the rightful ownership of the trademark.

[35] RCA will produce and market something similar starting in the spring of 1997.

"Internet ready." Versions of the WebTV device made by Sony and by Philips started appearing in stores in October 1996, at a cost of $329 plus a $19.95 monthly charge for access to the WebTV service.

The WebTV system consists of a TV-top box that houses a special-purpose computer to make the television set into a Worldwide Web screen. The service includes access to the Web through an on-screen menu, e-mail, and other features. A handheld remote control operates the box, and an optional keyboard is available. The device allows the TV set to display Internet sites, taking the place of a standard computer monitor, and to receive regular television signals (but not both simultaneously unless the user has a TV with "picture in picture" capability).

If WebTV or something like it succeeds in the market despite its initial lack of printer capability and lack of disk storage, that success could generate a spurt in sites *designed* for that type of system. Those sites would feature immediate entertainment or quickly grasped information not calling for printing or saving--sound and video clips, brief news summaries and announcements, graphics, interactive applets (small embedded software applications), games, chat, and so on. A growth in such sites could then spur further interest in WebTV and similar Internet access devices.

Cable Companies Seek a Role as Internet Access Providers

Cable TV systems have enormous downstream (toward the user's home) bandwidth. That enormous bandwidth offers an opportunity for cable operators to put their last mile infrastructure to a new use: Internet access. The key limitation is that the *upstream* bandwidth is relatively small, posing a bottleneck in the system.[36] It would, however, be convenient for the consumer to have Internet access over the existing system (no additional bill to pay, as the same company is providing both regular cable and Internet access).

Cable's downstream speed allows Web pages to appear almost instantly on the screen, in contrast to the slow, chunky, and crawling appearance usually experienced over POTS lines. This speed offers a large advantage over WebTV and similar devices that access the Net through modem and phone lines. (Cable Internet access requires a "cable modem," but does not use the consumer's phone line.)

Tests of cable-based Internet access have already begun, but the service is still new and available only in limited areas. As of September 1996, according to a report in *Time*, "America's two largest cable operators, TCI and Time Warner, launched the nation's first commercial cable-modem services in Fremont, California, and Akron, Ohio, respectively."[37] Other locations are expected to follow soon.

[36] Cable systems are designed to send large amounts of information to the consumer but need only receive small amounts back from the consumer. The connection is therefore not symmetrical.
[37] Michael Krantz, "Wired for Speed," *Time*, September 23, 1996, p. 54.

Cable-based Internet access might raise some interesting regulatory questions, as cable companies are regulated, but Internet access providers as such are not (although there are regulated companies, such as Pacific Telesis and AT&T, which also are or encompass Internet access providers). Exploration of such regulatory questions is beyond the scope of this paper, but could be addressed in a follow-up analysis.

Game Machines Meet the Web

In another example of the cross-pollination of technologies, video game maker Sega has announced a device to connect its game machines to the Internet. (A telephone connection is required.) This could provide another entry-level option for families that want to use the Web but do not want to buy a personal computer to do so. Like WebTV, the Sega device is an option that could start to demonstrate its viability in the market by January 1997. The device is to sell for around $450 (including the game machine itself), and requires an Internet access account. Sega has arrangements with Concentric, an Internet service provider, for accounts at $19.95 per month for users of the device. According to Sega's announcement, posted online:

> The Sega Saturn Net Link modem plugs into the cartridge slot on the Sega Saturn, providing an online connection to the Internet. Included with the modem is CD-based, HTML 2.0 compatible web browser software developed specifically for the NTSC TV display standard and offering a similar feature-set to Netscape 1.1 and other popular browsers currently in use.

In another venture, Atari and Chuck E. Cheese pizza chain founder Nolan Bushnell and partners have announced a system, to be installed in bars, restaurants, and similar locations, to link game-players via the Internet. Multi-player games could involve competing teams at one or many locations.

High-Tech Phones Offer Access the Internet

New telephones, on sale as of September 1996, will allow e-mail, financial transactions, and several other functions without a personal computer.[38] The phones include a small keyboard and a liquid-crystal display screen, making them the equivalent of a small, special purpose computer connected directly to the phone line.

Because e-mail is one of the most-used features of the Internet, even at a list price of around $300 such phones could quickly become widely used for that purpose. (They cannot yet receive e-mail, but that capacity is to be added within a year.) They also offer other features, such as caller identification, the ability to send text messages to pagers equipped to receive them, and a "cash card" slot to facilitate payments for goods and services.

[38] David Elrich, "High-Tech Phones Give Home Computers a Run for Their Money," *New York Times*, September 5, 1996 (posted on *New York Times* Web site).

Connected Applications: Hybrids of Local and Remote Resources

A "connected application" combines information flowing over the Internet with information residing on the user's local computer. This approach can radically reduce the need to transmit data (especially large files such as graphic images). For example, one use of the Internet is to play games among a group of users at remote locations. If the images (scenery, room layouts, backgrounds), sound files, and text are stored on each player's local hard disk (or other appropriate storage and access device, such as a CD-ROM disk), then all that must be transmitted is coordinates or brief text entries.

If this technique were limited to games, it would be of relatively narrow interest. However, there is no reason why the same approach could not be used with commercial catalogs, maps, reference materials, educational resources, and other types of data. America Online and other commercial networks routinely store graphics on the user's machine rather than transmitting images online every time the user invokes a new function. (America Online is notorious for the frequency with which it adds artwork online, but once the files have been sent, they may be called up on the user's machine, saving time online and speeding responses.)

Intranets and Corporate Computing

The Internet and the software it runs on have in turn spawned "intranets." An intranet is an internal corporate version of the Internet, with a similar user interface (browser), and with protective "firewall" and a gateway to the Internet.[39] Microsoft, Netscape, Sun, and others are very interested in this market, and are actively seeking to make their software (and in Sun's case, also hardware) the standard for intranets.[40]

Unfortunately, vendor documents on intranets tend to argue the merits of their respective approaches and to be fogged in techno-babble. The essential point, though, is that the same sorts of tools that work on the Web can work within an organization to make its documents easily available to staff and to facilitate communications and in-house transactions.

By means of a link to the Internet (through appropriate security systems, as shown in simplified form in Figure 8), corporate intranet users also have access to the entire scope of the Internet *and* outsiders may be granted controlled access to corporate information.

[39] A firewall is a system of hardware, software, or both that limits access between connected computer networks. Sometimes internetwork firewalls are called "proxy servers."

[40] Reportedly, profits in the browser business depend predominantly on these corporate intranet markets, not on individual users of Netscape, Microsoft Internet Explorer, and so on. For an extensive discussion of Intranets, see Ryan Bernard, *Corporate Intranet: Create and Manage an Internal Web for Your Organization* (N.Y.: John Wiley & Sons, 1996). In June 1996, both Microsoft and Netscape both posted white papers on their intranet strategies on their respective Web sites.