ONE

Introduction to the Internet

The Internet Protocol (IP) has changed the networking landscape in the last few years and is being adopted as the de facto networking protocol. The success and explosive growth of the Internet in the last few years have led engineers and designers of networking protocols to embrace IP. The fact that IP is an open source protocol and is easily available in all the major operating systems has ensured its success among a very large community of developers, designers, and testers. A large knowledge base and strong development skills make IP a viable networking protocol.

What is it about IP that makes it such a powerful force? There is no single answer or explanation. Economics are obviously a strong force, and the fact that IP reduces the cost of building networks is a clear advantage. The multitude of applications and services that have been developed over IP is another.

Email, the Internet, and numerous applications based on Web technology have become an integral aspect of business and everyday life. The Internet has spawned an entire industry commonly referred to as *e-commerce*. Use of the Internet, as well as applications and services that run over it, is constantly growing.

What constitutes an IP network? In short, networks that are built using the Internet Protocol as the network layer protocol can be loosely defined as IP networks. However, a network that operates using the Internet Protocol does not necessarily become a part of the Internet. There are many private networks, such as networks that run in enterprises, that operate on IP but may not be a part of the Internet. This chapter discusses primarily the Internet, which can be considered the largest IP network in existence.

1.1 What is the Internet?

Although many consider the Internet a recent phenomenon, in reality it has existed and grown since the early 1980s. Above all, the Internet is a communications medium. It has transformed the way information is disseminated or exchanged, is changing the way we do business, access information, and view entertainment today, and holds the promise of becoming as indispensable as the telephone network.

The Internet is a global network of heterogeneous networks enabling computers of all kinds to directly and transparently communicate and share services throughout much of the world. Over 100 million hosts are connected to this network. The Internet is a constantly growing, evolving, and changing network. It is extremely difficult to define the structure of the Internet or explain its characteristics. Some have even referred to it as a living organism that constantly keeps growing at its edges, and as such has a life of it's own. The graph in Figure 1–1 shows the growth of the Internet through 2001 in terms of the number of hosts connected.

What started off as a research network with a few limited hosts has grown into a global-scale commercial network. The Internet has created a

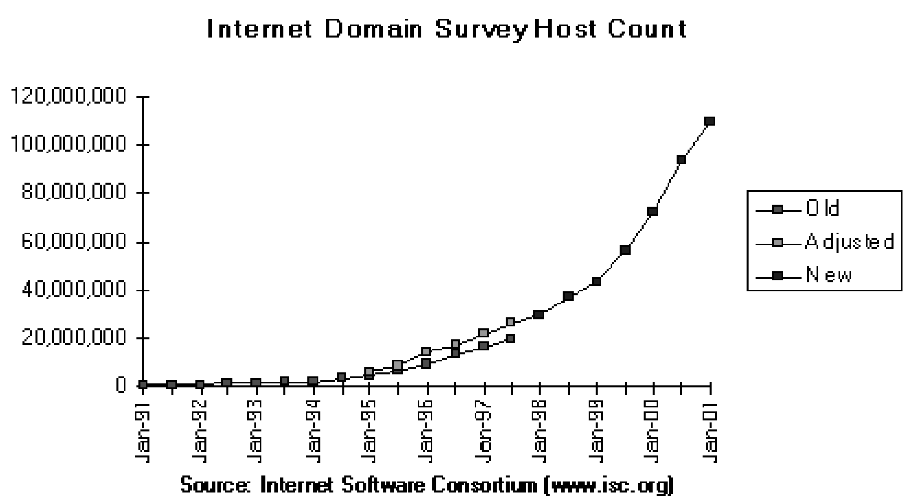

Figure 1–1 *Host count of nodes attached to the Internet.*

class of network operators called Internet service providers (ISPs). Over 6000 ISPs exist around the world today. Some of the bigger ISPs that offer services on a global scale include America Online, with about 30 million subscribers; UUNet, IIJ, NTT/Verio, Cable and Wireless, Earthlink, and Terra.

Fortunately, nobody really owns the Internet. There is no centralized control, and no one can turn it off. Since the Internet is formed by a collection of networks spread across the globe, there is no single company or government that has authority over it. As a result, we have a collection of networks that comprise university and research networks, commercial networks owned and operated by ISPs, networks belonging to governments, and non-government organizations (NGOs) as well.

1.2 History of the Internet

The Internet that exists today is the unintended outcome of the initial research objectives set by the Advanced Research Projects Agency (ARPA) of the United States Department of Defense (DOD) in the late 1960s.

The origins and history of the Internet go back as far as the early 1960s, and the ideas for the Internet as we know it today can be gleaned from memos written in August 1962 by J.C.R. Licklider, who discussed the concept of a “Galactic Network.” Leonard Kleinrock of the Massachusetts Institute of Technology (MIT) published the first paper on packet-switching theory in July 1961.

In the mid-1960s ARPA formulated a research interest in computer networks. ARPA is essentially an organization that does not directly do research by itself, but funds research work based on topics defined by it and manages these projects. In 1965 ARPA funded a study through Lawrence Roberts at MIT’s Lincoln Laboratory, largely as a result of Kleinrock’s earlier work at MIT. The study report, “A Cooperative Network of Time-Sharing Computers,” proposed to establish an experimental three-computer network. This proposal was successfully implemented a year later by Roberts and Thomas Merill when they connected the TX-2 computer at MIT to the Q-32 in California through a low-speed dial-up telephone line. A third computer located at ARPA was later connected to this network as well. This experimental network generated a great amount of interest in the computer research community and accelerated further developments. In late 1966, Roberts went to the Defense Advanced Research Projects Agency (former name of ARPA) to develop the computer network concept and came up with a plan for the ARPANET and published it in 1967. In 1968, ARPA invited prospective suppliers to build the network and in 1968 awarded the contract for ARPANET to Bolt Beranek and Newman (BBN). On September 1, 1969, the first Interface Message Processor (IMP) was shipped to the University of California at Los Angeles

(UCLA) and the first host computer was connected. The IMP at UCLA was connected to IMPs subsequently shipped to the Stanford Research Institute (SRI), the University of California at Santa Barbara (UCSB), and the University of Utah. In December 1970, Steve Crocker finished the initial ARPANET host-to-host protocol, called Network Control Protocol (NCP), and ARPANET sites completed implementing NCP during 1971–1972. In October 1972 a large successful demo of the ARPANET took place. Electronic mail was also introduced in 1972. Soon the ARPANET grew to include packet radio networks (ALOHA-net), packet satellite networks, and others.

The refinement of the ARPANET networking model was undertaken by Robert Kahn in 1972. The model of the ARPANET was to connect heterogeneous hosts via a homogeneous network. The concept that Kahn introduced was to allow the network itself to be heterogeneous. After joining DARPA and initiating the Internet program, Kahn enlisted Vinton Cerf of Stanford University to work with him on the detailed design of the protocol. Cerf, who had been involved earlier in the design of NCP, was already knowledgeable about interfacing to existing operating systems. Kahn and Cerf teamed up and came up with the specification for the Kahn-Cerf protocol, which was later named the Transmission Control Protocol/Internet Protocol (TCP/IP). The paper "A Protocol for Packet Network Interconnection," published in May 1974, described TCP and provided all the Internet's transport and forwarding services. The design intent was to have TCP support a range of transport services ranging from sequenced, in-order reliable delivery of data to a datagram service that was less than robust (lossy) and did not guarantee in-order packet delivery. Initial TCP implementations resulted in essentially a protocol that emulated virtual circuits. This model was fine for applications such as file transfer and remote login, but other applications such as packet voice demonstrated that it is not always necessary for the transport protocol to correct errors. That is best left to applications that obviously have their own requirements. This resulted in some changes to the initial proposal, and the outcome was two transport protocols: TCP and User Datagram Protocol (UDP). UDP was added as an unreliable datagram transport protocol for applications that did not need error correction, retransmissions, and congestion management.

DARPA funded Stanford (Cerf), BBN (Tomlinson), and University College London (Peter Kirstein) to implement TCP. Stanford produced a detailed specification from which multiple implementations were produced. Early implementations of TCP were for large time-sharing systems, but David Clark and his research group at MIT proved that the protocol could be adapted as well to small host computers and workstations. The 1980s introduced local-area networks (LANs), PCs, and workstations. Ethernet technology took off on a very large scale. This allowed Internet technology to take hold at the grassroots level. A bigger challenge was to make the technology widely available and have it adapted by multiple people. DARPA supported the University of California at Berkeley, which had developed a version of the Unix operating system. Berkeley developed the TCP/IP code and fit it into the Unix system kernel. Berkeley Software Distribution (BSD) Unix, which was quite popular with the computer science community, adopted the Internet technology and helped foster its growth even further.

TCP/IP was adopted as a defense standard in 1980. The ARPANET was being used by both military and the research community. The ARPANET transitioned from NCP to TCP in 1983, and at the same time the network was split into MILNET (supporting the defense department) and an ARPANET (supporting the research community).

The usefulness of computer networking and applications such as e-mail was recognized by other communities as well, and as a result a host of networks began springing up. In 1985 the U.S. NSFNET (National Science Foundation Network—the Internet backbone in the U.S. which was supported by NSF) program announced its intent to serve the entire higher-education community. NSFNET also mandated TCP/IP for this network. Thus developed the NSFNET backbone. NSFNET, which was funded by federal funds, ended in 1995 and the research network transitioned into a commercial network. In an 8.5-year span the backbone had grown from 6 nodes to 21 nodes. The Internet also constituted more than 50,000 networks on all continents (with 29,000 networks in the United States). The ARPANET was decommissioned in 1990. TCP/IP had spread so vigorously that it made other networking protocols irrelevant or marginal and ensured the success of IP for the future.

1.3 The Internet Architecture

The Internet as it stands today is composed of more than 60,000 constituent networks. Each network is an autonomous network, and the only things that are common across these networks is the use of a common protocol, adherence to a common addressing scheme, and a common name structure. Hence it is hard to visualize the Internet architecture as beginning somewhere or having a center. Over the past 25 years since the ARPANET became operational, the Internet has increased by factors of thousands with respect to backbone speed and by factors of millions with respect to the number of hosts connected to the network. The network can be extended at any point and offer further connectivity. In a way the network is akin to the universe itself, which is ever expanding.

Is there an Internet architecture at all? Many in the Internet community believe that there does not exist any real architecture but instead there exists a set of guidelines. The architectural goal of the Internet is connectivity; the tool for achieving this is the Internet Protocol, and the intelligence is end to end rather than being embedded in the network [RFC 1958].

Important features of the original Internet architecture are as follows:

1. A connectionless packet-forwarding infrastructure (dumb network) that positions higher-level functionality at the edge of the network for robustness (fate sharing)
2. Addresses that are fixed-size numerical with a simple (net, host) hierarchy

The evolution of this architecture over the last 25 years has led to the following architectural features:

1. Subnetting, autonomous systems (AS), and the Domain Name System (DNS). These features were introduced as a direct result of the understanding obtained from deployments and as a measure to address scalability and growth.
2. Congestion control mechanisms were introduced in the late 1980s.
3. IP multicasting developed as a means to disseminate packets from a single node to a number of receiving nodes via an enhanced addressing scheme.

The Internet is structured in a hierarchical manner or at least can be visualized as such. However, it is not a clear hierarchy, but rather a mixed-up one. At the top there exist multiple backbone networks that connect to each other at network access points (NAPs) or exchange points (EPs). Backbone networks of large Internet service providers such as NTT/Verio, Sprint, and UUNet, which span the globe, carry the majority of the Internet traffic. Figure 1–2 shows a network map of UUNet's global network (http://uunet.com/network/maps/) as of June 2001.

Regional networks and ISPs connect to the larger backbones, and smaller ISPs, in turn, are connected to these regional networks. End users obtain connectivity to the Internet via ISPs in general. Peering between ISPs and service agreements between ISPs allow inter-ISP traffic. A model of this hierarchical model of the Internet is shown in Figure 1–3.

In reality, the Internet is not organized as a strict hierarchy. Peering networks could connect to each other and at the same time have connectivity to a larger network. The large backbone ISPs could also be the ones providing end-user connectivity. This type of a model is feasible when you have a relatively small number of ISPs, but as the Internet continues to grow in size and the number of users increases, so do usage and the number of Internet service providers. As a result, ISPs will need to continue to exchange increasing amounts of data traffic. One of the basic underlying mechanisms in the Internet today is the points where various ISPs exchange traffic. These interconnection points are called NAPs or EPs. Exchange points can be defined as

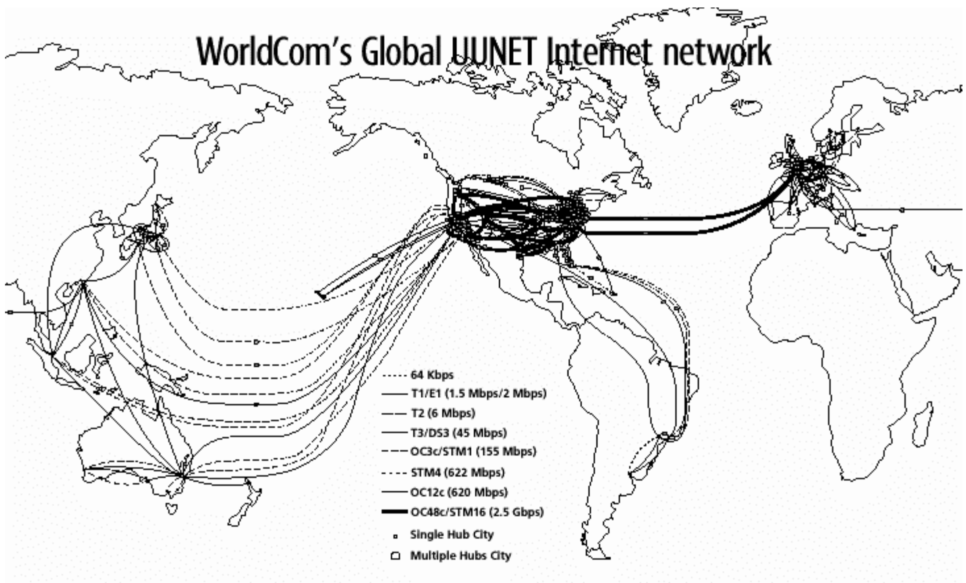

Figure 1–2 *An example global backbone network.*

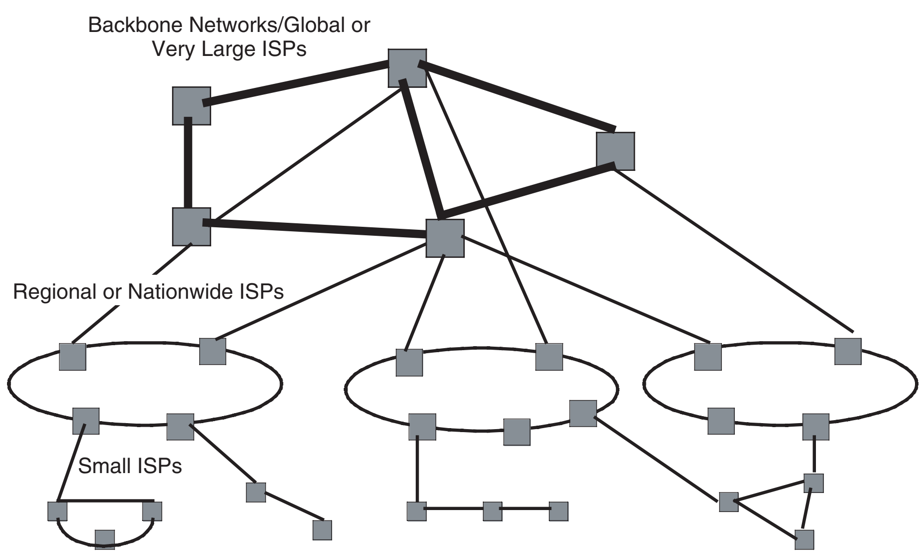

Figure 1–3 *Hierarchical view of the Internet.*

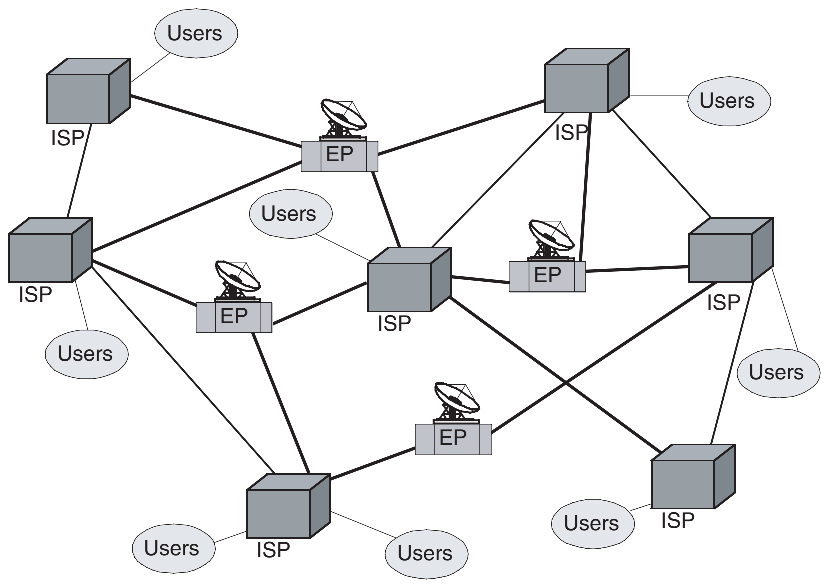

Figure 1–4 *Current view with exchange points.*

multiple-access networks allowing ISPs to exchange traffic and routing information with other ISPs. There are many exchange points across continents today. Some of the larger ones are Sprint NAP, MAE (metropolitan area exchange) East and West, and LINX (London Internet Exchange). A view of Internet architecture taking NAPs/EPs into consideration is shown in Figure 1–4.

The underlying telecommunications and transmission infrastructure that forms the Internet has grown significantly in the last few years. As the processing power of central processing unit (CPUs) of computers and handheld devices keeps increasing, so does the bandwidth of the networks.

1.4 Services on the Internet

There are many different uses of the Internet today, ranging from business communications to information access to entertainment and financial transactions.

Applications on the Internet today can be classified into two types: traditional and new age. Applications such as e-mail, file transfer protocol (FTP), bulletin boards, and Usenet can be considered traditional, while real-time voice, streaming audio/video, instant messaging, and network games can be thought of as the new-age applications.

Let's take a look at a traditional and very widely used application; e-mail or electronic mail. E-mail has become the communication medium of choice. It has become so pervasive that we see it in almost every walk of life today. It is just as common for someone to give their e-mail address as their point of contact as their phone number. The pioneers of the Internet, the ARPANET folks, had probably not envisioned e-mail as a key or "killer" application when the network was created. Another application with a similar scale of acceptance is Instant Messaging, the equivalent of which is Short Message Service (SMS) in the wireless world.

With the advent of faster processors and more bandwidth, the Internet has become a hotbed of innovation, and new applications are constantly coming online. The pace of innovation has been brisk over the last few years, and the future holds even more exciting and interesting applications and use.

1.5 The World Wide Web

Some of the current success of the Internet is undoubtedly because of the emergence of the World Wide Web in the last decade. The phenomenal success of the Web has made the Internet accessible to so many people that it has become synonymous with the Internet. Of course, the Internet itself has been around for a while and was used by mostly engineers, researchers, and technical people before the introduction of the Web browsers and the World Wide Web. For a complete history of the World Wide Web, refer to the W3C Web site (www.w3c.org). W3C (World Wide Web Consortium) is the body that is defining the standards and ensuring the evolution of the Web. But is the World Wide Web the Internet? Not really; the Web is merely an application on the Internet. A good definition of the Web comes from W3C, which defines the Web as the universe of network-accessible information (available through your computer, phone, television, or networked refrigerator).

The concept of the Web originated at CERN (European Organization for Nuclear Research) in Switzerland in 1989 and was the idea of Tim Berners-Lee, who is considered the inventor of the Web. The proposal of using hypertext for information distribution/linking and management was the basis from which the Web was developed. Over the course of the last 10 years, the concept has been refined and has been developed into a global-scale application. Since the Web is essentially an application on the Internet, some of the design principles that have been taken into account are based on the principles on which the Internet has been established:

- *Interoperability:* Specifications for the Web's languages and protocols must be compatible with one another and allow (any) hardware and software used to access the Web to work together.

- *Evolution:* The Web must be able to accommodate future technologies. Design principles such as simplicity, modularity, and extensibility will increase the chances that the Web will work with emerging technologies such as mobile Web devices and digital television, as well as others to come.
- *Decentralization:* Decentralization is without a doubt the newest principle and most difficult to apply. To allow the Web to "scale" to worldwide proportions while resisting errors and breakdowns, the architecture (like the Internet) must limit or eliminate dependencies on central registries.

We can think of the Web as a set of cooperating clients and servers. The window to the Web for most people is via graphical user interface applications known as Web browsers (such as Netscape or Internet Explorer). The World Wide Web is different things to different people. Some think of the Web as a source of information; others think of it as a place to shop or conduct financial transactions; still others think of it as a source of entertainment, including music, audio, and video.

What ties the Web together is the use of a common protocol, which in this case is HTTP. From the stack perspective, HTTP is an application-layer protocol that runs over TCP. HTTP is a text-oriented request response protocol. Another protocol for interacting with the Web has been developed for the wireless industry. The protocol, known as WAP (Wireless Application Protocol), allows wireless devices such as mobile phones and PDAs to access the Internet. WAP is discussed in much further detail in Chapter 13.

The Web and the Internet have become a utility service that is as ubiquitous as the telephone, and the next expansion in this area is essentially in the area of wireless networks and users connecting to the Web via devices that are no longer tethered.

1.6 Internet Registries

Three regional Internet registries are responsible for the assignment of IP addresses and autonomous system numbers globally (other organizations are responsible for the assignment of domain names):

- ARIN—American Registry for Internet Numbers
- APNIC—Asia Pacific Network Information Centre
- RIPE—Réseaux IP Européens

ARIN is a nonprofit organization established for the purpose of administration and registration of IP numbers for the following geographical areas: North America, South America, the Caribbean, and sub-Saharan Africa.

APNIC represents the Asia Pacific region, comprising 62 economies. It is a not-for-profit, membership-based organization whose members include Internet service providers, national Internet registries, and similar organizations.

RIPE is an open collaborative community of organizations and individuals operating wide area IP networks in Europe and beyond. The objective of the RIPE community is to ensure the administrative and technical coordination necessary to enable operation of a pan-European IP network. RIPE has no formal membership, and its activities are performed on a voluntary basis.

1.7 A View of the Future

As voice and data continue to converge, the Internet is becoming the focal point of attention and is taking on the role of the one single global network providing access to information as well as different types of services. But the Internet is not about to replace the existing voice networks, such as the public switched telephone network (PSTN) and other data networks, such as the ATM and frame relay networks, that are deployed on a global scale as well. The one thing that the Internet has yet to achieve is the reliability and security of the PSTN. But it is definitely on its way to achieving that factor of reliability and security. The Internet in its first generation has been mostly connected to fixed and wired end hosts. But as the world becomes increasingly wireless and people are more mobile and need constant connectivity, the Internet as we know it today will change. The Internet will grow in areas that make information access ubiquitous from any place on the earth.

Exponential growth over the years has put some strain on the original design, and as a result a new Internet protocol, IPv6, will be deployed over the next few years. New types of applications, multimedia services, games, and a host of other things will make the Internet an integral aspect of daily life.

1.7.1 Internet2

The Internet was initially a network that was primarily used by research institutes. At one time, use of the Internet by commercial organizations required these organizations to agree to a policy that constrained them to use the Internet only for commercial purposes. The Internet as we know it today is primarily a commercial network.

As a result of this change in philosophy, (i.e., from being a research network to a commercial network), there has been an impact on the network in terms of congestion and availability. Research organizations decided to build a next generation Internet exclusively for research purposes. This initiative resulted in Internet2 (www.internet2.org).

The primary goals of Internet2 are to:

- Create a leading-edge network capability for the national research community.
- Enable revolutionary Internet applications.
- Ensure the rapid transfer of new network services and applications to the broader Internet community.

At the time of this writing, Internet2 has also deployed IPv6 in its backbone.

1.8 Summary

This chapter covered some of the high-level aspects of the Internet. A description of Internet architecture and history, as well as the standards bodies that define the protocols that makes the Internet work, were described. A great amount of information and detail has been omitted because this book focuses mainly on the use of IP in wireless networks. In order to understand the reasons for wireless networks adopting IP and moving toward IP-based architectures in the next generation, it is important to understand the Internet itself.

References

Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Lawrence G. Roberts, Stephen S. Wolff, *The Past and Future History of the Internet.*

RFC 1958, *Architectural Principles of the Internet,* B. Carpenter, Editor. Available on-line at www.rfc-editor.org/rfc/rfc1958.txt.

Robert Braden, David Clark, Scott Shenker, and John Wroclawski, *Developing a Next-Generation Internet Architecture.*

Licklider, J.C.R. and Clark, W., *Online Man–Computer Communication,* Aug. 1962.

Roberts, L. and Merrill T., “Toward a Cooperative Network of Time-Shared Computers,” *Proceedings of the Fall AFIPS Conference* (Oct. 1966).

Internet Software Consortium (2001). “Internet Domain Survey, January 2001.” Available on-line at www.isc.org/ds/www-200101/index.html.