1

The Challenge of the Internet

To say that the Internet, with its millions of end users operating 24 hours per day, is a complex system would be one of the greatest understatements of the 21st century. However, if we take a step back and look at the Internet from a systems engineering perspective, we can observe a number of subsystems that are much easier to understand. Let's begin our journey into this fascinating architecture by exploring its origins.

1.1 A Brief History of the Internet

The Internet with a capital I is one of the world's most interesting achievements in computer science and networking technology. It provides a worldwide mechanism for user-to-user, computer-to-computer communication that spans corporate and national boundaries. This achievement is even more amazing because the Internet is self-governing, run by committees comprised largely of volunteers. In the past, many government organizations, such as the U.S. Department of Defense and individual states, have subsidized the basic expenses. Much of the research into the Internet protocols is conducted at major U.S. research universities, such as the University of California, the University of Colorado, the University of Illinois, and the University of Texas. As the Internet has evolved, however, the expenses for Internet connectivity have been passed down to the individual users — those who benefit from access to this public resource. Let's see how this unique system evolved.

Today's Internet was born in 1969 as the Advanced Research Projects Agency Network (ARPANET) and was sponsored by the U.S. Defense Advanced Research Projects Agency (DARPA), now known as ARPA. The purpose of the ARPANET was to test and determine the viability of a communication technology known as packet switching. The contract to build the original ARPANET was awarded to a firm known as Bolt, Baranek, and Newman (now BBN Technologies, Cambridge, Massachusetts). ARPANET went online in September 1969 at four locations: Stanford Research Institute (SRI), the

University of California at Santa Barbara (UCSB), the University of California at Los Angeles (UCLA), and the University of Utah. The original hosts were Honeywell minicomputers, known as Interface Message Processors (IMPs).

The initial test was successful, and the ARPANET grew quickly. At the same time, it became apparent that nonmilitary researchers could also benefit from access to a network of this type, so leaders in the university and industrial research communities made proposals to the National Science Foundation (NSF) for a cooperative network of computer science researchers. The NSF approved funding for the Computer Science Network (CSNET) in 1981.

In 1984, the ARPANET was split into two different networks: MILNET (for unclassified military traffic) and ARPANET (for nonmilitary traffic and research). In 1984, the NSF established the Office of Advanced Scientific Computing (OASC) to further the development of supercomputers and to make access to them more widely available. The OASC developed the NSFNET to connect six supercomputing centers across the United States, using T-1 lines operating at 1.544 Mbps in 1987, and subsequently upgrading to the T-3 rate (44.736 Mbps) in 1990. NSFNET, with its higher transmission rates, was a resounding success; as a result, the U.S. Department of Defense declared the ARPANET obsolete and dismantled it in June 1990.

In the meantime, NSFNET connections encompassed a system of regional and state networks. The New England Academic and Research Network (NEARNET), the Southeastern Universities Research Association Network (SURAnet), and the California Education and Research Federation Network (CERFnet) were among the family of NSFNET-connected networks. Since these networks were designed, built, and operated, in part, with government funds, regulations called Acceptable Use Policies, or AUPs, governed the types of traffic that could traverse these networks. In general, traffic that was for "research or educational purposes" was deemed acceptable; other traffic was either discouraged or prohibited. Few, if any, accounts of the "Internet Police" apprehending an AUP violator were recorded, however, again testifying to the self-governing nature of the Internet.

The business community, seeing the new opportunity for electronic commerce, began looking for ways to support general business traffic on the Internet without violating these regulations. As a result of this opportunity, the Commercial Internet Exchange Association (CIX) was formed in 1991. CIX was a nonprofit trade organization of Public Data Network service providers that promoted and encouraged the development of the public data communications internetworking services. Membership in CIX was open to organizations that offered TCP/IP or OSI public data internetworking services to the general public. CIX gave these service providers a neutral forum for the discussion and development of legislative, policy, and technology issues. Member networks agreed to interconnect with all other CIX members, and to exchange traffic. There were no restrictions placed on the traffic routed between member networks. Nor were there "settlements," or traffic-based charges, as a result of these interconnections. CIX has since disbanded, and has been replaced by the United States Internet Service Provider Association (US ISPA).

Another outgrowth of the Internet expansion was the founding of the nonprofit Internet Society (ISOC) in 1992. The Internet Society, headquartered in Reston, Virginia, is an international organization that strives for global cooperation and coordination for the Internet. Members of the ISOC include government agencies, nonprofit research and educational organizations, and for-profit corporations. The charter of the ISOC emphasizes support for the technical evolution of the Internet, educating the user community in the use and application of the Internet, and promoting the benefits of Internet technology for education at all grade levels.

In 1993, the NSF announced that it would no longer provide the traditional backbone architecture, but instead would specify a number of locations, called *Network Access Points* (NAPs), where various ISPs could interconnect and exchange traffic. This concept was based on the CIX concept and specified four NAPs, located at San Francisco, Chicago, New York, and Washington, DC. In addition to these four NAPs, two Federal Internet Exchange (FIX) points, one on each coast, were developed: FIX-East and FIX-West. Other NAPs, operated as commercial entities, also exist today.

In April 1995, the existing NSFNET backbone was retired and replaced by a new architecture that provides for very high-speed connectivity. This new network is called the very high-speed Backbone Network Service (vBNS). The vBNS is based on both Asynchronous Transfer Mode (ATM) and Synchronous Optical Network (SONET) technologies, and is designed as a noncommercial research network to develop and test high-speed applications, routing, and switching technologies.

Figure 1-1 illustrates a timeline of significant events in the history of the Internet. Shown at the top of the figure are the various government organizations, starting with DARPA, that have participated in the development of the Internet. On the second line are the various organizations, starting with the ARPANET Working Group (WG) and then proceeding to the Internet

Society (ISOC) and the World Wide Web Consortium (W3C), that have provided technical oversight into the development. At the lower portion of the figure are some of the significant milestones and their associated dates [1-1].

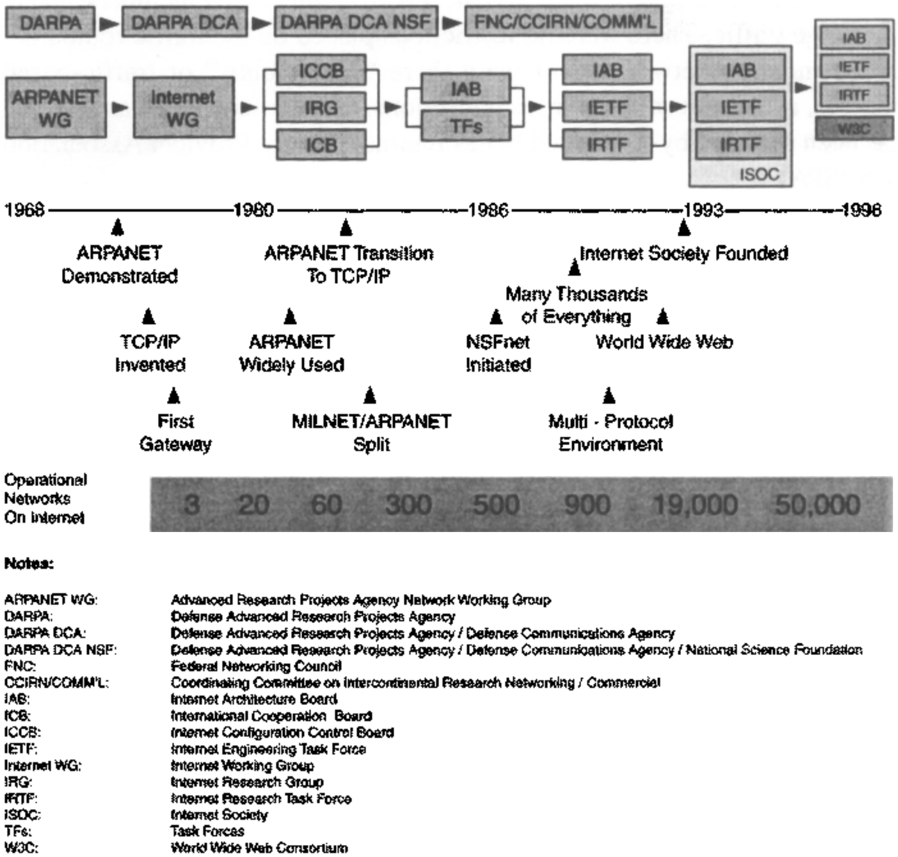

Figure 1-1. Internet Development Timeline.
(Source: www.isoc.org/internet/history/brief.shtml.)

Perhaps the most important aspect of the Internet history is the Internet's astounding growth in the last few years. Figures 1-2 and 1-3, taken from Reference [1-2], illustrate the growth in Internet Hosts and World Wide Web Servers, respectively. Note that in Figure 1-2 a new host counting mechanism, believed to be more accurate than the previous method, is shown for host statistics beginning in 1995.

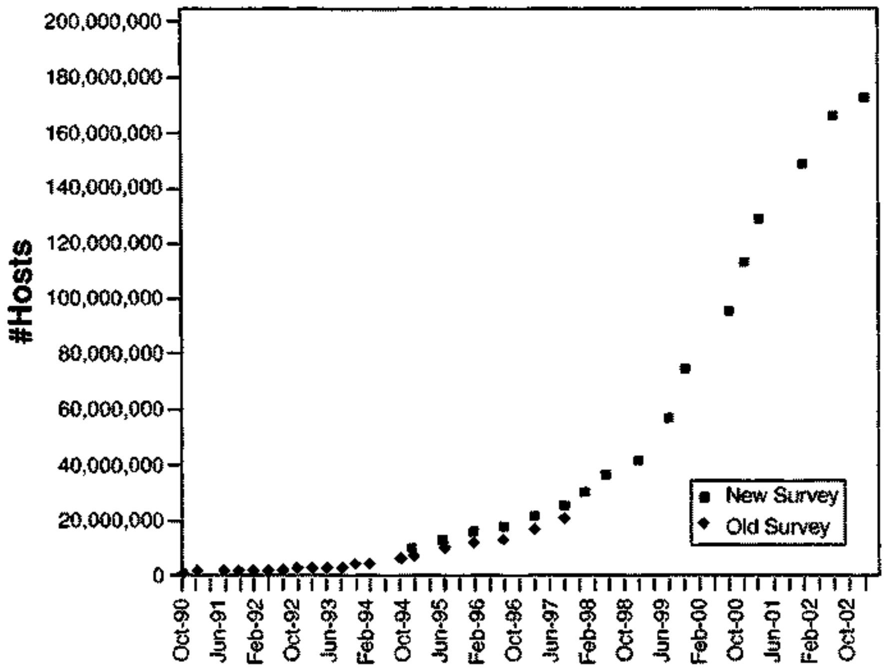

Figure 1-2. Internet Hosts.
(Source: Hobbes' Internet Timeline, www.zakon.org/robert/internet timeline/ Copyright 2003, Robert H. Zakon. Reprinted with permission.)

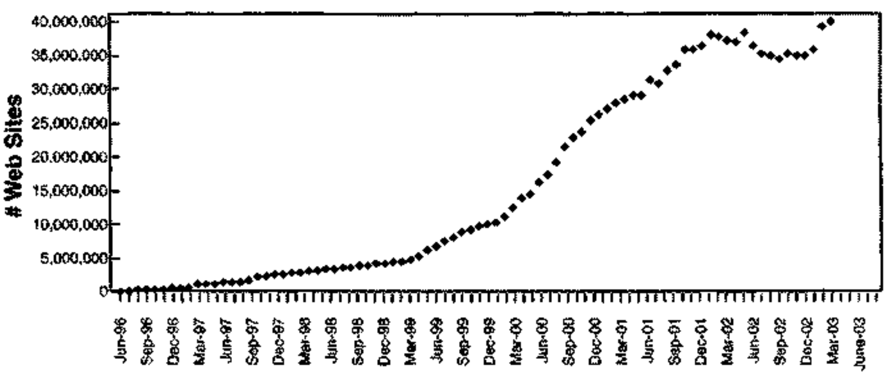

Figure 1-3. WWW Networks Growth.
(Source: Hobbes' Internet Timeline, www.zakon.org/robert/internet/timeline/ Copyright 2003, Robert H. Zakon. Reprinted with permission.)

Lessons learned from the ARPANET have had a significant effect on a number of data communication technologies, such as LANs and packet switching.

1.2 Governing and Documenting the Internet

An amazing characteristic of the largely volunteer Internet community is how smoothly this rather avant-garde organization operates. To quote from RFC 1726 [1-3]: "A major contributor to the Internet's success is the fact that there is no single, centralized point of control or promulgator of policy for the entire network. This allows individual constituents of the network to tailor their own networks, environments and policies to suit their own needs. The individual constituents must cooperate only to the degree necessary to ensure that they interoperate."

But no organization, avant-garde or not, can operate without some degree of structure. The Internet Society provides some of that structure, and one of the Internet Society's components is the Internet Architecture Board (IAB), chartered in 1992. The IAB consists of 13 members: 12 full members plus the chair of the Internet Engineering Task Force (IETF). The IAB's responsibilities include [1-4]:

- Appointing a chair of the IETF and its subsidiary Internet Engineering Steering Group (IESG)
- Oversight of the architecture for the protocols and procedures used by the Internet
- Oversight of the process used to create Internet standards
- Editorial management and publication of the Request for Comment (RFC) document series and administration of the various Internet assigned numbers
- Representing the interests of the Internet Society to other organizations
- Providing guidance to the Internet Society regarding Internet technologies

Two task forces report to the IAB. The IETF coordinates the technical aspects of the Internet and its protocols and ensures that it functions effectively. The Internet Research Task Force (IRTF) researches new technologies.

The IAB produces numerous protocol standards and operational procedures that require dissemination and archiving. Most of these documents are known as Requests for Comments documents (RFCs) and are reviewed by the appropriate IETF or IRTF members. Many of the Internet protocols have become U.S. military standards and are assigned a MIL-STD number. For example, the Internet Protocol is described in both RFC 791 and MIL-STD-1777. Internet documents are published according to two tracks. Off-track specifications are labeled with one of three categories: experimental, informational, or historic. These documents contain information that is useful but not appropriate for an Internet standard. Specifications that are destined to become Internet standards evolve through several levels of testing and revision, known as the standards track, described in RFC 2026 [1-5]. Three "maturity levels" are defined. The first level of maturity is called a Proposed Standard. This standard is stable, well understood, has been reviewed by the community, and has sufficient interest to be considered valuable. A Draft Standard is one from which at least two independent and interoperable implementations have been developed and for which sufficient operational experience has been obtained. After significant implementation and operational experience is obtained, a Draft Standard may be elevated to an Internet Standard. The various organizations that are involved in this process are described in RFC 2028 [1-6]. An interesting history of the RFC process is found in RFC 2555 [1-7].

Two sources of information on Internet standards and parameters are updated on a periodic basis. The Internet Assigned Numbers Authority (IANA) documents protocol parameters, assigned addresses such as port numbers, and many others [1-8]. The Internet Official Protocol Standards document (currently RFC 3300 [1-9]) describes the standards track process and lists recently published RFCs and the current standardization status of the various protocols. Information on how to obtain Internet documentation is given in Appendix B.

With the explosive growth of the Internet within the last few years, a wealth of information is available. Throughout this text, we will make extensive use of the RFCs and other documents in our study. Readers needing details on the Internet itself could benefit from the following: RFC 1207, "FYI on Questions and Answers — Answers to Commonly Asked 'Experienced Internet User' Questions" [1-10]; RFC 1402, "There's Gold in Them Thar Networks! Or Searching for Treasure in All the Wrong Places" [1-11]; RFC 1580, "Guide to Network Resource Tools" [1-12]; RFC 2664, "FYI on Questions and Answers — Answers to Commonly Asked 'New Internet User' Questions" [1-13]; RFC 1738, "Uniform Resource Locators" [1-14]; RFC

1935, "FYI on 'What is the Internet'" [1-15]; RFC 1983, "Internet Users' Glossary" [1-16]; and RFC 2151, "A Primer on Internet and TCP/IP Tools" [1-17].

1.3 The Protocols of the Internet

The Internet is based on an architecture that was developed by the Advanced Research Projects Agency (ARPA), which is funded by the United States government. There are literally hundreds of protocols that fit into the ARPA architecture and provide various functions within the Internet. In this text, we will concentrate our efforts on three broad categories of protocols. The first category could be called the *ARPA Core Protocols*, because they provide the underlying infrastructure that facilitates packet transport and end-user communications. These core protocols are shown in Figure 1-4, which will be explored in depth in Chapter 2. In general, these protocols would be visible (and known) to the end users of an Internet-connected network. Note that the associated RFC or MIL-STD reference is included with each protocol listed.

ARPA Layer	Protocol Implementation								OSI Layer
Process / Application	Hypertext Transfer	File Transfer	Electronic Mail	Terminal Emulation	Domain Names	File Transfer	Client / Server	Network Management	Application
	Hypertext Transfer Protocol (HTTP) RFC 2616	File Transfer Protocol (FTP) MIL-STD-1780 RFC 959	Simple Mail Transfer Protocol (SMTP) MIL-STD-1781 RFC 2821	TELNET Protocol MIL-STD-1782 RFC 854	Domain Name System (DNS) RFC 1034, 1035	Trivial File Transfer Protocol (TFTP) RFC 1350	Sun Microsystems Network File System Protocols (NFS) RFC 3530	Simple Network Management Protocol (SNMP) v1: RFC 1157 v2: RFC 1901-10 v3: RFC 3411-18	Presentation / Session
Host-to-Host	Transmission Control Protocol (TCP) MIL-STD-1778 RFC 793				User Datagram Protocol (UDP) RFC 768				Transport
Internet	Address Resolution ARP RFC 826 RARP RFC 903		Internet Protocol (IP) MIL-STD-1777 RFC 791			Internet Control Message Protocol (ICMP) RFC 792			Network
Network Interface	Network Interface Cards: Ethernet, Token Ring, MAN and WAN RFC 894, RFC 1042 and others								Data Link
	Transmission Media: Twisted Pair, Coax, Fiber Optics, Wireless Media, etc.								Physical

Figure 1-4. ARPA Core Protocols.

The second broad category of protocols could be called the *ARPA Control, Routing, and Address Resolution Protocols.* These protocols work behind the scenes to make sure that the packet transport functions operate properly and are kept up to date. The control, routing, and address resolution protocols are shown in Figure 1-5, and will be studied in greater detail in Chapter 3. In general, these protocols would not be visible to the end users of an Internet-connected network.

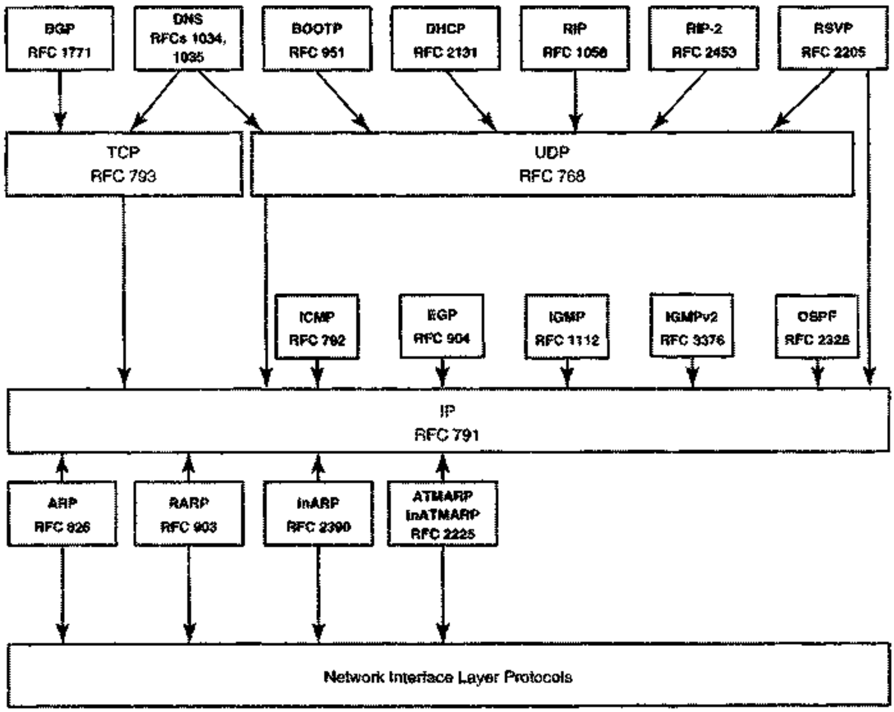

Figure 1-5. ARPA Control, Routing, and Address Resolution Protocols.

The third broad category of protocols could be called the *ARPA Multimedia Protocols.* These protocols facilitate the transmission of voice and video information over the Internet. The multimedia protocols are shown in Figure 1-6, and will be studied in greater detail in Chapter 9. In general, these protocols would also not be visible to the end users of an Internet-connected network.

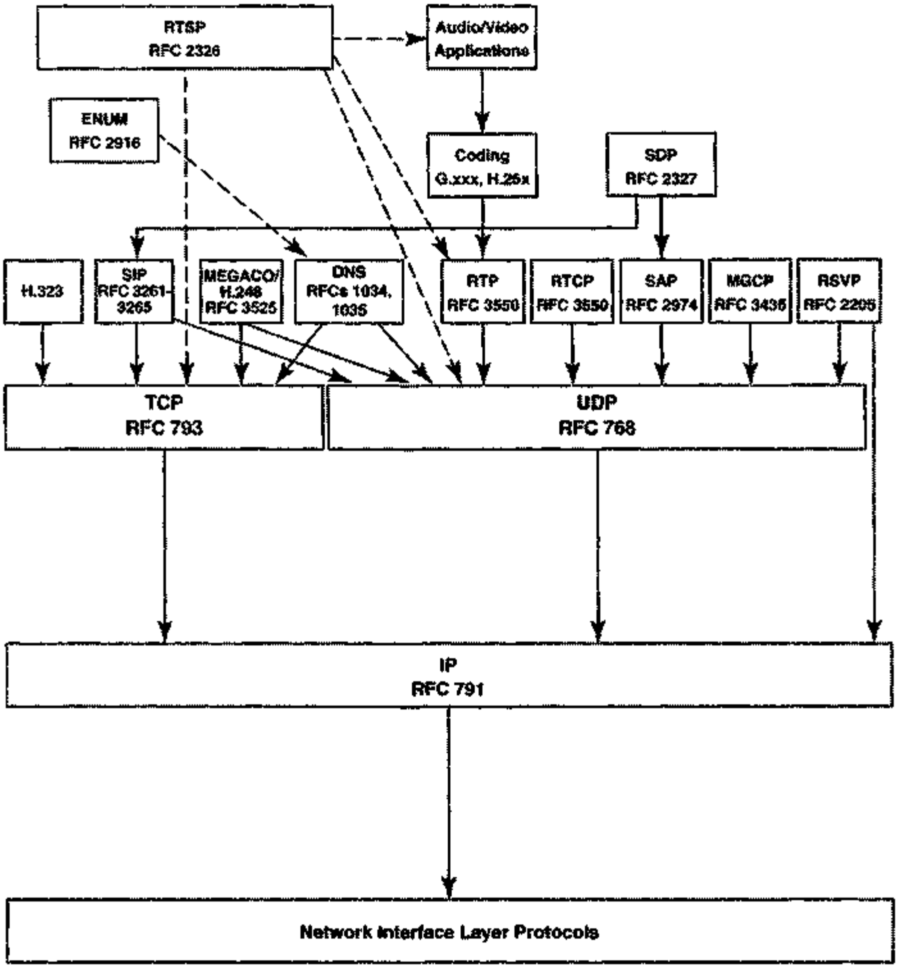

Figure 1-6. ARPA Multimedia Protocols.

1.4 Outline of This Book

With a brief glance at Figures 1-4, 1-5, and 1-6, you can see a large number of protocols that are required to support various Internet-based user applications. One of the challenges, therefore, is to sort out the various protocol functions and understand how they interrelate and work together. We will address this challenge by dividing this text into five major parts:

Part I — Introduction

Chapter 1: The Challenge of the Internet

Part II — Packet Transport

Chapter 2: Analyzing the IP Network
Chapter 3: Datagram Addressing and Delivery
Chapter 4: Routing and Intranetwork Communication
Chapter 5: End-to-End Reliability
Chapter 6: The Next Generation: IPv6
Chapter 7: Case Studies in Packet Transport

Part III — Application Support

Chapter 8: Data Transport
Chapter 9: Converged Networks and Multimedia Transport
Chapter 10: Case Studies in Application Support

Part IV — Network and Performance Management

Chapter 11: Network Management Architectures
Chapter 12: Network Management System Components
Chapter 13: Case Studies in Network and Performance Management

Part V — Reference Appendices

Appendix A: Acronyms and Abbreviations
Appendix B: Sources of Internet Information
Appendix C: Addresses of Standards Organizations
Appendix D: Trademarks

Of special note are Chapters 7, 10, and 13, which provide case studies, captured from live internetworks, that illustrate problems that can arise when using the Internet protocols and solutions that can address these problems. All of these case studies are illustrated with trace listings from the *Sniffer*® network analyzer from Network Associates, Inc.

1.5 Looking Ahead

In this chapter, we have considered a brief history of the Internet and examined some of the challenges that we will encounter from the study of the Internet protocols. In the next chapter, we will begin our work in earnest by considering the architecture of an Internet-connected network.

1.6 References

[1-1] Leiner, Barry M., et al. A Brief History of the Internet. Available from: http://www.isoc.org/internet/history/brief.shtml.

[1-2] Zakon, Robert Hobbes. Hobbes' Internet Timeline. RFC 2235. Current information is available from: http://www.zakon.org/robert/internet/timeline/.

[1-3] Kastenholz, F., and C. Partridge. "Technical Criteria for Choosing IP: The Next Generation." RFC 1726, December 1994.

[1-4] Carpenter, B., Ed. "Charter of the Internet Architecture Board." RFC 2850, May 2000. The IAB can be reached at www.iab.org.

[1-5] Bradner, S. "The Internet Standards Process — Revision 3." RFC 2026, October 1996.

[1-6] Hovey, R., and S. Bradner. "The Organizations Involved in the IETF Standards Process." RFC 2028, October 1996.

[1-7] RFC Editor, et al. "30 Years of RFCs." RFC 2555, April 1999.

[1-8] The Internet Assigned Numbers Authority (IANA) maintains an online database of protocol numbers and parameters at www.iana.org/numbers.html. This replaces the former "Assigned Numbers" document, which was last published as RFC 1700 in October 1994.

[1-9] Reynolds, J., et al. "Internet Official Protocol Standards." RFC 3300, November 2002.

[1-10] Malkin, G., et al. "FYI on Questions and Answers — Answers to Commonly Asked 'Experienced Internet User' Questions." RFC 1207, February 1991.

[1-11] Martin, J. "There's Gold in Them Thar Networks! Or Searching for Treasure in All the Wrong Places." RFC 1402, January 1993.

[1-12] EARN Staff. "Guide to Network Resource Tools." RFC 1580, March 1994.

[1-13] Plzak, A., et al. "FYI on Questions and Answers — Answers to Commonly Asked 'New Internet User' Questions." RFC 2664, August 1999.

[1-14] Berners-Lee, T., et al. "Uniform Resource Locators (URL)." RFC 1738, December 1994.

[1-15] Krol, E., and E. Hoffman. "FYI on 'What is the Internet.'" RFC 1935, April 1996.

[1-16] Malkin, G., Editor. "Internet Users' Glossary." RFC 1983, August 1996.

[1-17] Kessler, G., and S. Shepard. "A Primer on Internet and TCP/IP Tools." RFC 2151, June 1997.