Chapter 1

The Origin of TCP/IP and the Internet

Two people can communicate effectively when they agree to use a common language. They could speak English, Spanish, French, or even sign language, but they must use the same language.

Computers work the same way. Transmission Control Protocol/Internet Protocol (TCP/IP) is like a language that computers speak. More specifically, TCP/IP is a set of rules that defines how two computers address each other and send data to each other. This set of rules is called a protocol. Multiple protocols that are grouped together form a protocol suite and work together as a protocol stack.

TCP/IP is a strong, fast, scalable, and efficient suite of protocols. This protocol stack is the de facto protocol of the Internet. As information exchange via the Internet becomes more widespread, more individuals and companies will need to understand TCP/IP.

In this first chapter you'll look at the origins of TCP/IP. You will learn about:

- The features of TCP/IP
- ARPAnet
- TCP's method of moving data
- Requests for Comments (RFCs)
- The benefits of using TCP/IP

What Is TCP/IP?

TCP/IP is a set of **protocols** that enable communication between computers. There was a time when it was not important for computers to communicate with each other. There was no need for a common protocol. But as computers became networked, the need arose for computers to agree on certain protocols.

protocols
Rules or standards that govern communications.

Today, a **network administrator** can choose from many protocols, but the TCP/IP protocol is the most widely used. Part of the reason is that TCP/IP is the protocol of choice on the Internet—the world's largest network. If you want a computer to communicate on the Internet, it'll have to use TCP/IP.

network administrator
A person who installs, monitors, and troubleshoots a network.

NOTE

When multiple protocols work together, the group is collectively known as a protocol suite or protocol stack. TCP/IP is an example of a protocol suite (it describes multiple protocols that work together). The implementation of TCP/IP is described as a protocol stack. Both terms are used interchangeably, yet their definitions vary slightly.

Another reason for TCP/IP's popularity is that it is compatible with almost every computer in the world. The TCP/IP stack is supported by current versions of all the major operating systems and network operating systems—including Windows 95/98, Windows NT, Windows 2000, Windows XP, Linux, Unix, and NetWare.

Unlike proprietary protocols developed by hardware and software vendors to make their equipment work, TCP/IP enjoys support from a variety of hardware and software vendors. Examples of companies that have products that work with TCP/IP include Microsoft, Novell, IBM, Apple, and Red Hat. Many other companies also support the TCP/IP protocol suite.

TCP/IP is sometimes referred to as "the language of the Internet." In addition to being the official language of the Internet, TCP/IP is also the official language of many smaller networks. For all the computers that are attached to the Internet to communicate effectively, they must agree on a language. Just like every human language has certain rules so that the people involved in the conversation understand what the other is saying, a computer language needs a set of rules so that computers can effectively communicate. Some of the rules of a language that computers use to communicate include determining when to send data and when to receive data.

Features of TCP/IP

TCP/IP has been in a use for more than 20 years, and time has proven it to be a tested and stable protocol suite. TCP/IP has many features and benefits. In this section, you will learn about some of the most important ones.

host
Any device (such as a workstation, server, mainframe, or printer) on a network or internetwork that has a TCP/IP address.

Support from Vendors

As stated earlier, TCP/IP receives support from many hardware and software vendors. This means that the TCP/IP suite is not tied to the development efforts of a single company. Instead, the choice to use TCP/IP on a network can be based on the purpose of the network and not on the hardware or software that has been purchased.

Interoperability

One of the major reasons why the TCP/IP suite has gained popularity and acceptance so universally is that it can be installed and used on virtually every platform. For example, using TCP/IP, a Unix host can communicate and transfer data to a DOS host or a Windows host. A **host** is another name for a computer or device on a network. TCP/IP eliminates the cross-platform boundaries.

Flexibility

TCP/IP is an extremely flexible protocol suite, and in later chapters you will learn about some features that contribute to this flexibility. Examples of TCP/IP's flexibility include the latitude an administrator has in assigning and reassigning addresses. An administrator can automatically or manually assign an IP address to a host, and a TCP/IP host can convert easy-to-remember names, such as `www.sybex.com`, to a TCP/IP address.

Routability

A limitation of many protocols is their difficulty moving data from one segment of the network to another. TCP/IP is exceptionally well adapted to the process of routing data from one segment of the network to another, or from a host on a network in one part of the world to a host on a network in another part of the world.

In the following sections, you will learn about how these features of TCP/IP grew out of the military's need for a reliable, flexible networking standard.

The Origins of the Internet: ARPAnet

ARPAnet
The Advanced Research Projects Agency's super-network—the predecessor of the Internet.

Network Control Protocol (NCP)
The protocol used before TCP/IP.

Understanding the roots of the Internet will give you insight into the development of TCP/IP and many of its rules and standards. If you know why TCP/IP was created and how it evolved, the TCP/IP protocol suite is easier to understand.

The predecessor of today's Internet was **ARPAnet**, a supernetwork that was created by the Advanced Research Projects Agency (ARPA) and launched in 1969. This network was created in response to the potential threat of nuclear attack from the Soviet Union. One of ARPA's primary goals was to design a fault-tolerant network that would enable U.S. military leaders to stay in contact in case of nuclear war. By the standards of the time, this fault-tolerant network seemed to be almost science fiction. ARPA set out on a mission to create a network with what seemed to be impossible requirements.

NOTE

In the late 1950s, the United States Department of Defense (DoD), under the guidance of one of America's leading think tanks, the RAND corporation, formed the Advanced Research Projects Agency (ARPA).

The protocol, or language of choice, used on the ARPAnet was called **Network Control Protocol (NCP)**—TCP/IP had not yet been developed. As the ARPAnet grew, however, a new protocol was needed because NCP simply didn't fulfill all the needs of a larger network. The NCP protocol was similar to a human language that has only a few words. The language might enable a few people to communicate, but as you include more people who want to talk about many more subjects, you have to improve the language.

The ARPAnet project had some specific goals and requirements. To reach these goals and meet these requirements, some of the top computer minds worked in a collaborative effort with little financial or public glory. Many of the top computer minds that worked on the ARPAnet were affiliated with major universities. It was not the intention of the project leaders to create the worldwide network that exists today, but fantastic growth soon followed the ARPAnet's humble beginnings.

ARPAnet's Requirements

To fulfill the needs of the military, the new ARPAnet had to meet the following requirements:

No one point more critical than any other Because the network needed to be able to withstand a nuclear war, there could be no one critical part of the network and no single point of failure. If there were any critical parts of the network, enemies could target that area and eliminate communications.

Redundant routes to any destination Because any location on the network could be taken down by enemies in the event of a war, there had to be multiple routes from any source to any destination on the network. Without redundant routes, any one location could become a critical communications link and a potential point of failure.

On-the-fly rerouting of data If any part of the network failed, the network had to be able to reroute data to its destination on-the-fly.

Ability to connect different types of computers over different types of networks This network could not be tied to just one operating system or hardware type. Because universities, government agencies, and corporations often rely on different types of Local Area Networks (LANs) and network operating systems, interoperability among these many networks was critical. Connecting to the network should not dictate that a lot of new hardware had to be purchased; rather, the existing hardware should suffice.

Not controlled by a single corporation If one corporation had a monopoly on this network, the network would grow to boost the corporation instead of the usefulness and effectiveness of the network. This network needed to be a cooperative effort among many engineers who were working to improve the network for the sake of the supernetwork, not that of a corporation.

By December of 1969 the ARPAnet had four hosts. The ARPAnet consisted of computers at the University of California at Los Angeles, the University of California at Santa Barbara, the University of Utah, and Stanford Research Institute. The ARPAnet set the foundation for what would grow up to be the Internet.

Requests for Comments

Request for Comments (RFC)
A paper thoroughly describing a new protocol or technology.

Internet Engineering Task Force (IETF)
A governing body of the Internet.

To improve the technology that was being used on the ARPAnet, a system was designed to encourage and facilitate correspondence among the engineers who were developing this new network. This system, which is still in use today, relies on **Requests for Comments (RFCs)** to provide feedback and collaboration among engineers. An RFC is a paper that has been written by an engineer, a team of engineers, or just someone with a better idea, to define a new technology or enhance an existing technology.

The process of submitting RFCs was designed to be a "bulletin board" for posting technical theories. The old-school way of writing a thesis or book was too slow. RFCs provided an informal and fast way to share new technologies and ideas for enhancements. After an RFC is written and posted, it can be evaluated, critiqued, and used by other engineers and developers. If another engineer or developer can improve on the theory or standard, the RFC provides an open forum in which to do so. Many of these papers are long, painstakingly technical, and in most cases good reading material for someone with difficulty sleeping.

An RFC can be submitted for review to the **Internet Engineering Task Force (IETF)**. Engineers from the IETF review the papers that are submitted and assign a number to each. From that point on, the RFC number becomes the effective "name" of the paper. For example, the first RFC, which is about host software, is called RFC 1. RFC 1 was submitted in 1969 by a developer named Steve Crocker. There are currently more than 3,000 RFCs.

As the ARPAnet was growing and researchers and engineers were making improvements, they used RFCs as a tool to strengthen and ensure the network's foundation. TCP/IP is a child of the RFC method of development—no corporation makes money when you install TCP/IP. Using RFCs has been the method of growing the ARPAnet with the best network minds contributing.

TIP

It is possible for anyone to write and publish an RFC. Instructions on how to write and submit an RFC are detailed in RFC 2223. Today, RFCs are posted on many Web sites. Appendix D describes this book's companion Web site, which has links to RFC 2223 and other RFC Web sites.

The Birth of TCP/IP

As stated earlier, the "language" spoken by hosts on the ARPAnet in 1969 was called NCP. However, NCP had too many limitations and was not robust enough for the supernetwork, which was beginning to grow out of control. The limitations of NCP and the growth of the ARPAnet lead to research and development of a new network language.

Transmission Control Protocol (TCP)
The protocol describing communication between hosts.

In 1974 Vint Cerf and Bob Kahn, two Internet pioneers, published "A Protocol for Packet Network Interconnection." This paper describes the **Transmission Control Protocol (TCP)**, which is a protocol in the protocol suite that would eventually replace NCP.

Transmission Control Protocol/Internet Protocol (TCP/IP)
The suite of protocols that when combined create the "language of the Internet."

The TCP protocol describes the host-to-host portion of a communication. TCP explains how two hosts can set up this communication and how they can stay in touch with each other as data is being transferred. NCP did not resolve these issues to the extent that TCP was able to.

As you will learn in later chapters, TCP is responsible for making sure that the data gets through to the other host. It keeps track of what is sent and retransmits anything that did not get through. If any message is too large for one package, TCP splits the message into several packages and makes sure that they all arrive correctly. After they have arrived, TCP at the other end puts all the packages back together in the proper order.

By 1978, testing and further development of this language led to a new suite of protocols called **Transmission Control Protocol/Internet Protocol (TCP/IP)**. In 1982, it was decided that TCP/IP would replace NCP as the standard language of the ARPAnet. RFC 801 describes how and why the transition from NCP to TCP was to take place. On January 1, 1983, ARPAnet switched over to TCP/IP and the network continued to grow exponentially.

In 1990, the ARPAnet ceased to exist. The Internet has since grown from ARPAnet's roots, and TCP/IP has evolved to meet the changing requirements of the Internet.

Design Goals of TCP/IP

TCP/IP has evolved to its current state. The protocols within the TCP/IP suite have been tested, modified, and improved over time. The original TCP/IP protocol suite had several design goals that intended to make it a viable protocol for the large, evolving internetwork. Some of these goals included:

Hardware independence A protocol suite that could be used on a Mac, PC, mainframe, or any other computer.

Software independence A protocol suite that could be used by different software vendors and applications. This would enable a host on one site to communicate with a host on another site, without having the same software configuration.

Failure recovery and the ability to handle high error rates A protocol suite that featured automatic recovery from any dropped or lost data. This protocol must be able to recover from an outage of any host on any part of the network and at any point in a data transfer.

Efficient protocol with low overhead A protocol suite that had a minimal amount of "extra" data moving with the data being transferred. This extra data, called overhead, functions as packaging for the data being transferred and enables the data transmission. Overhead is similar to an envelope used to send a letter, or a box used to send a bigger item—having too much overhead is as efficient as using a large crate to send someone a necklace.

Ability to add new networks to the internetwork without service disruption A protocol suite that enabled new, independent networks to join this network of networks without bringing down the larger internetwork.

Routable Data A protocol suite on which data could make its way through an internetwork of computers to any possible destination. For this to be possible, a single and meaningful addressing scheme must be used so that every computer that is moving the data can compute the best path of every piece of data as it moves through the network.

The TCP/IP protocol suite has evolved to meet these goals. Throughout this book, you will learn how TCP/IP has met and surpassed these original design goals.

Moving Data across the Network

Creating this new "super network" introduced many new concepts and challenges for the pioneering engineers. One of the most critical issues was how to move data across the network. Older communications protocols relied on a circuit-switched technology. TCP/IP, however, introduced a new way of moving data across a network. The protocol suite set a new standard for communications and data transport by using a packet-switched network.

circuit-switched network
A network on which all data in a communication takes the same path.

TCP/IP's method of moving data and information helped the protocol suite fulfill several of the requirements for the growing ARPAnet supernetwork. In the following sections, you'll learn about how circuit-switched and packet-switched communications methods work.

Moving Data on a Circuit-Switched Network

Historically, data has moved through a **circuit-switched network**. In a circuit-switched network, data moves across the same path throughout the entire communication. An example of a circuit-switched network is the telephone system. When you make a telephone call, a single path (also called a circuit) is established between the caller and the recipient. For the rest of the conversation, the voice data keeps moving through the same circuit. If you were to make a call and get a very staticky connection, you would hang up and try again. This way you could get a different circuit, hopefully one with less static. Early network data transmissions followed this type of pathway.

In the illustration below, notice that although the data could take multiple routes, all the data moves from the source to the destination along the same path. In a circuit-switched network, data communication moves along a single, established route.

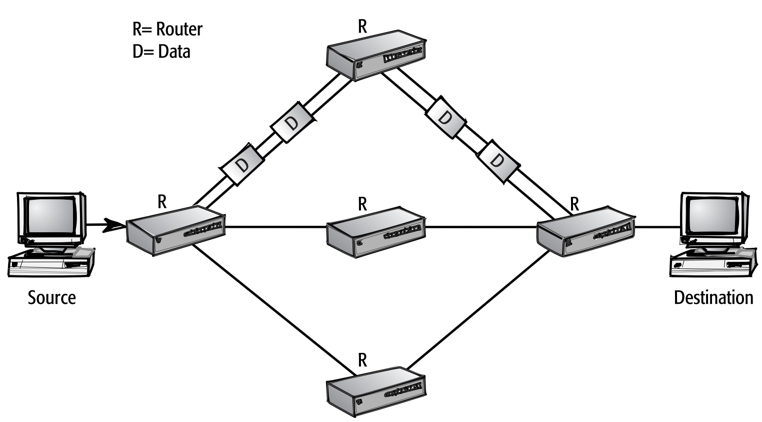

Moving Data on a Packet-Switched Network

packet-switched network
A network on which the data in a communication takes several paths.

packet
A unit of data that is prepared for transmission onto a network.

A circuit-switched network was unacceptable for both the ARPAnet and the Internet. Data had to be able to move through different routes so that if one circuit went down or got staticky, it didn't affect communication on the rest of the network. Instead, data simply would take a different route.

The Internet uses a **packet-switched network**. On a packet-switched network, the computer that is sending the data fragments the data into smaller, more manageable chunks. These chunks are called **packets**. Each packet is then individually addressed and sent to its intended recipient. As the several packets make their way through the network, each packet finds its own way to the receiver. The receiving computer reassembles the packets into the original message.

The illustration below shows how TCP/IP moves data. Notice that there are several routes that the data packets can follow from the source to the destination. Unlike the illustration on the preceding page, the data packets here use a variety of routes—some follow the same path, while others follow different paths. Each packet follows its own route, and data is reassembled at the destination. This is how information moves on a packet-switched network.

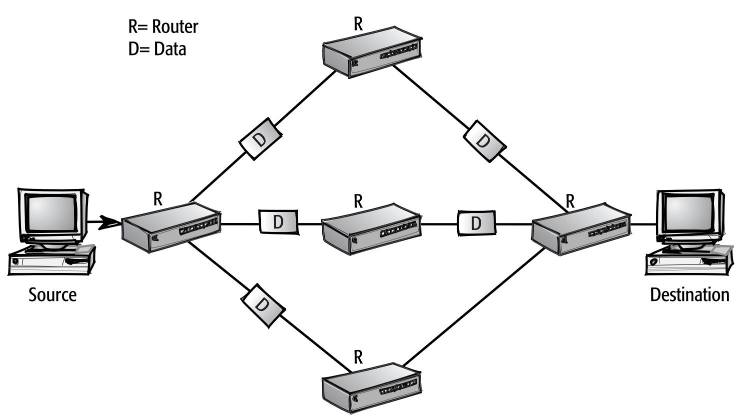

Understanding How a Packet-Switched Network Functions

To help you understand how a packet-switched network moves data, let's look at a similar real-world situation.

Let's say that I take my son's soccer team to an arcade and restaurant for a team party. I have the whole team outside of the arcade. My task is to get the team to the other side of the arcade, to my wife who is waiting for them in the restaurant. In this analogy, the team represents the complete file on one host, and each child represents a data packet. One of my goals is to lose as few of the kids as possible.

While we are standing outside, it is easy to put the team in order; all the children are wearing numbered jerseys. I tell the kids that we will meet on the other side of the arcade in a restaurant for pizza and that they should all move as fast as possible through the arcade and to the restaurant.

After I open the door and say, "go," the kids enter one at a time. Entering the arcade one at a time represents the fragmenting and sending of the file. Just as each of the kids has a numbered jersey, each packet has a number so that the receiving host can put the data back together.

Now picture a dozen six-year-olds moving through the arcade. Some of the children will take a short route; others will take a long route. Possibly, they'll all take the same route, though it is much more likely that they will all take different routes. Some will get hung up at certain spots, but others will move through faster. My wife is in the restaurant waiting to receive the team. As they start arriving at the restaurant, she can reassemble the children (packets) in the correct order because they all have a number on their backs. If any are missing, she will wait just a bit for the stragglers and then send back a message that she is missing part of the team (file).

After I receive a message that she is missing a child (a packet), I can resend the missing part. I do not need to resend the entire team (all the packets), just the missing child (packet or packets).

Please note, however, I would not go look for the lost child, I would just put the same numbered jersey on a clone of the lost child and send him into the arcade to find the restaurant.

Why Use TCP/IP?

TCP/IP offers many advantages over other network protocols and protocol suites. Here is a summary of some of the benefits of using the TCP/IP protocol suite:

Widely published, open standard TCP/IP is not a secret. It is not proprietary or owned by any corporation. Because it is a published protocol with no secrets, any computer engineer is able to improve or enhance the protocol by publishing an RFC.

Compatible with different computer systems TCP/IP enables any system to communicate with any other system. It is like a universal language that would enable people from any country to communicate effectively with people from any other country.

Works on different hardware and network configurations TCP/IP is accepted and can be configured for virtually every network created.

Routable protocol TCP/IP can figure out the path of every piece of data as it moves through the network. Because TCP/IP is a routable protocol, the size of any TCP/IP network is virtually unlimited.

Reliable, efficient data delivery TCP/IP can guarantee that the data is transferred to another host.

Single addressing scheme TCP/IP uses a single and relatively simple addressing scheme. You will learn about TCP/IP's addressing in Chapter 6. An administrator can transfer knowledge of TCP/IP to any TCP/IP network without relearning the addressing scheme.

The Internet has become a necessity for business, and it soon will be a necessity at home. Many businesses, large and small, are connected to the Internet and are using TCP/IP as the protocol of choice for their internal networks. As more and more homes connect to the Internet, those computers will also use the TCP/IP protocol suite. The commercial implications of the Internet have changed the dynamic of every business model that has ever been taught.

The Origin of TCP/IP and the Internet

TCP/IP is the standard for a communications protocol on the Internet. You cannot connect to the Internet without using TCP/IP. Whether you build a network at home with two hosts or you manage an **internetwork** at your business with 100,000 hosts, TCP/IP is a communications protocol that will work effectively. TCP/IP can scale to any size environment and is robust enough to connect different types of LANs.

These are a few of the many reasons why network administrators choose to use TCP/IP as the protocol on their networks.

internetwork
Several smaller networks connected together.

Review Questions

Terms to Know

- ❑ protocols
- ❑ network administrator
- ❑ host
- ❑ ARPAnet
- ❑ NCP
- ❑ RFC
- ❑ IETF
- ❑ TCP
- ❑ TCP/IP
- ❑ circuit-switched network
- ❑ packet-switched network
- ❑ packet
- ❑ internetwork

1. The Internet was originally called:

   ____

2. List three requirements that the military mandated of this new network.

   ____

   ____

   ____

3. Another name for a computer on a TCP/IP network is:

   ____

4. Describe packet-switched and circuit-switched networks.

   ____

5. What is an RFC?

   ____

6. What protocol did TCP/IP replace?

   ____

7. True or False: TCP/IP is one protocol.

   ____

8. What is IETF?

   ____

9. List four benefits of using TCP/IP.

____

____

____

____

10. What year was the change made from NCP to TCP/IP?

____