15.2 Open Systems and Protocols

Many protocols have been defined to assist in network communication. Some have gained a stronger foothold than others because of many reasons, often historical in nature. This section focuses on the protocols used for general Internet traffic. Before we discuss the details of particular protocols, however, it is important to put them in context by discussing the concept of an open system.

What is a protocol?
Protocol is defined as a code prescribing strict adherence to correct etiquette and procedure (as in a diplomatic exchange). Computing terminology has borrowed the word to describe the correct etiquette for computers to use when communicating with one another.

Open Systems

Early in the development of computer networks, commercial vendors came out with a variety of technologies that they hoped businesses would adopt. The trouble was that these proprietary systems were developed with their own particular nuances and did not permit communication between networks of differing types. As network technologies grew, the need for interoperability became clear; we needed a way for computing systems made by different vendors to communicate.

An open system is one based on a common model of network architecture and a suite of protocols used in its implementation. Open-system architectures maximize the opportunity for interoperability.

The International Organization for Standardization (ISO) established the Open Systems Interconnection (OSI) Reference Model to facilitate the development of network technologies. It defines a series of layers of network interaction. Figure 15.5 shows the seven layers of the OSI Reference Model.

Each layer deals with a particular aspect of network communication. The highest level deals with issues that relate most specifically to the application program in question. The lowest layer deals with the most basic electrical and mechanical issues of the physical transmission medium (such as types of wiring). The other layers fill in all other aspects. The network layer, for example, deals with the routing and addressing of packets.

Proprietary system A system that uses technologies kept private by a particular commercial vendor

Interoperability The ability of software and hardware on multiple machines and from multiple commercial vendors to communicate

Open system A system that is based on a common model of network architecture and an accompanying suite of protocols

Open Systems Interconnection (OSI) Reference Model A seven-layer logical breakdown of network interaction to facilitate communication standards

Number	Layer
7	Application layer
6	Presentation layer
5	Session layer
4	Transport layer
3	Network layer
2	Data Link layer
1	Physical layer

FIGURE 15.5 The layers of the OSI Reference Model

The details of these layers are beyond the scope of this book, but it is important to know that networking technology as we know it today is possible only through the use of open-system technology and approaches such as the OSI Reference Model.

Network Protocols

Following the general concepts of the OSI Reference Model, network protocols are layered such that each one relies on the protocols that underlie it, as shown in Figure 15.6. This layering is sometimes referred to as a **protocol stack**. The layered approach allows new protocols to be developed without abandoning fundamental aspects of lower levels. It also provides more opportunity for their use, in that the impact of new protocols on other aspects of network processing is minimized. Sometimes protocols at the same level provide the same service as another protocol at that level, but do so in a different way.

A protocol is, in one sense, nothing more than an agreement that a particular type of data will be formatted in a particular manner. The details of file formats and the sizes of data fields are important to software developers who are creating networking programs, but we do not explore those details here. The importance of these protocols is simple to understand: They provide a standard way to interact among networked computers.

The lower two layers in Figure 15.6 form the foundation of Internet communication. Other protocols, sometimes referred to as high-level protocols, deal with specific types of network communication. These layers are essentially one particular implementation of the OSI Reference Model and correspond in various ways to the levels described in that model. Let's explore these levels in more detail.

Protocol stack Layers of protocols that build and rely on each other

Transmission Control Protocol (TCP) The network protocol that breaks messages into packets, reassembles them at the destination, and takes care of errors

Internet Protocol (IP) The network protocol that deals with the routing of packets through interconnected networks to the final destination

TCP/IP A suite of protocols and programs that support low-level network communication

TCP/IP

TCP stands for **Transmission Control Protocol** and IP stands for **Internet Protocol**. The name **TCP/IP** (pronounced by saying the letters T-C-P-I-P) refers to a suite of protocols and utility programs that support low-level network communication. The name TCP/IP is written to reflect the nature of the protocols' relationship: TCP rests on top of the IP foundation.

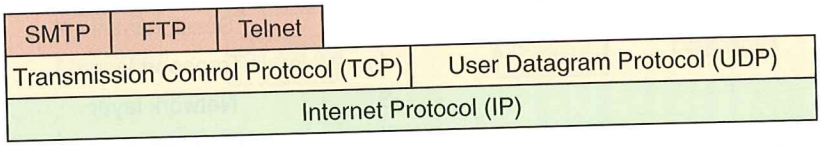

FIGURE 15.6 Layering of key network protocols

IP software deals with the routing of packets through the maze of interconnected networks to their final destination. TCP software breaks messages into packets, hands them off to the IP software for delivery, and then orders and reassembles the packets at their destination. TCP software also deals with any errors that occur, such as if a packet never arrives at the destination.

UDP stands for User Datagram Protocol. It is an alternative to TCP. That is, UDP software plays the same role as TCP software. The main difference is that TCP is highly reliable, at the cost of decreased performance, whereas UDP is less reliable, but generally faster. UDP is part of the TCP/IP suite of protocols. Because of the heavy reliance on TCP, and for historical reasons, the entire suite is referred to as TCP/IP.

An IP program called ping can be used to test the reachability of network designations. Every computer running IP software "echoes" ping requests, which makes ping a convenient way to test whether a particular computer is running and can be reached across the network. Ping officially stands for Packet InterNet Groper, but the name was contrived to match the term used when submarines send out a sonar pulse and listen for the returned echo. Because ping works at the IP level, it often responds even when higher-level protocols might not. The term *ping* is often used as a verb among network administrators: "Ping computer X to see if it is alive."

Another TCP/IP utility program called traceroute shows the route that a packet takes to arrive at a particular destination node. The output of traceroute is a list of the computers that serve as the intermediate stopping points along the way.

» User Datagram Protocol (UDP) An alternative to TCP that achieves higher transmission speeds at the cost of reliability

» Ping A program used to test whether a particular network computer is active and reachable

» Traceroute A program that shows the route a packet takes across the Internet

High-Level Protocols

Other protocols build on the foundation established by the TCP/IP protocol suite. Here are some of the key high-level protocols:

- Simple Mail Transfer Protocol (SMTP)—A protocol used to specify the transfer of electronic mail.
- File Transfer Protocol (FTP)—A protocol that allows a user on one computer to transfer files to and from another computer.
- Telnet—A protocol used to log into a computer system from a remote computer. If you have an account on a particular computer that allows telnet connections, you can run a program that uses the telnet protocol to connect and log in to that computer as if you were seated in front of it.
- Hypertext Transfer Protocol (HTTP)—A protocol defining the exchange of World Wide Web documents,

?

Sir Bill?

Bill Gates, co-founder of the Microsoft® Corporation with Paul Allen, is one of the best-known innovators of the PC revolution. He is consistently ranked as one of the world's wealthiest people and, as of March 2009, was ranked as "the" wealthiest. After his last full-time day at Microsoft in June 2008, he turned his attention to the Bill and Melinda Gates Foundation–the philanthropic institution he co-founded with his wife–which is currently the largest transparently operated charitable foundation in the world. In 2005, Gates received an honorary knighthood from Queen Elizabeth II in a private ceremony. He was honored for his charitable activities around the world and his contribution to the high-tech enterprise in Great Britain. Gates has received many honorary doctorate degrees, including one from Harvard University (2007), the university from which he dropped out in 1975 to found Microsoft.

which are typically written using the Hypertext Markup Language (HTML). HTML is discussed further in Chapter 16.

These protocols all build on TCP. Some high-level protocols have also been defined that build on top of UDP to capitalize on the speed it provides. However, because UDP does not provide the reliability that TCP does, UDP protocols are less popular.

Several high-level protocols have been assigned a particular *port* number. A port is a numeric designation that corresponds to a particular high-level protocol. Servers and routers use the port number to help control and process network traffic. Figure 15.7 lists common protocols and their ports. Some protocols, such as HTTP, have default ports but can use other ports as well.

Port A numeric designation corresponding to a particular high-level protocol

MIME type A standard for defining the format of files that are included as email attachments or on websites

MIME Types

Related to the idea of network protocols and standardization is the concept of a file's MIME type. MIME stands for Multipurpose Internet Mail Extension. Although MIME types do not define a network protocol, they define a standard for attaching or including multimedia or otherwise specially formatted data with other documents, such as email.

Based on a document's MIME type, an application program can decide how to deal with the data it is given. For example, the program you use to read email may examine the MIME type of an email attachment to determine how to display it (if it can).

MIME types have been defined for the documents created by many common application programs, as well as for data from particular content areas. Chemists and chemical engineers, for example, have defined a large set of MIME types for various types of chemical-related data.

FIGURE 15.7 Some protocols and the ports they use

Protocol	Port
Echo	7
File Transfer Protocol (FTP)	21
Telnet	23
Simple Mail Transfer Protocol (SMTP)	25
Domain Name Service (DNS)	53
Gopher	70
Finger	79
Hypertext Transfer Protocol (HTTP)	80
Post Office Protocol (POP3)	110
Network News Transfer Protocol (NNTP)	119
Internet Relay Chat (IRC)	6667

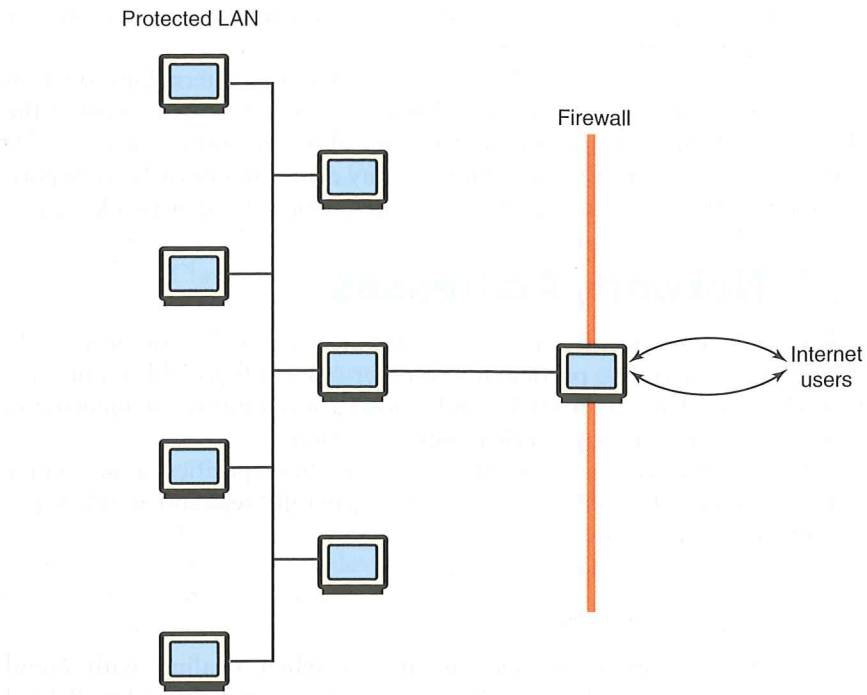

FIGURE 15.8 A firewall protecting a LAN

Firewalls

A firewall is a machine and its software that serve as a special gateway to a network, protecting it from inappropriate access. A firewall filters the network traffic that comes in, checking the validity of the messages as much as possible and perhaps denying some messages altogether. The main goal of a firewall is to protect (and, to some extent, hide) a set of more loosely administered machines that reside "behind" it. This process is pictured in Figure 15.8.

A firewall enforces an organization's access control policy. For example, a particular organization may allow network communication only between its users and the "outside world" via email, but deny other types of communication, such as accessing websites. Another organization may allow its users to freely access the resources of the Internet, but may not want general Internet users to be able to infiltrate its systems or gain access to its data.

The system administrators of an organization set up a firewall for their LAN that permits "acceptable" types of communication and denies other types. This policy can be implemented in a variety of ways, although the most straightforward approach is to deny traffic on particular ports. For example, a firewall could be set up to deny a user outside the LAN the

Firewall A gateway machine and its software that protects a network by filtering the traffic it allows

Access control policy A set of rules established by an organization that specify which types of network communication are permitted and denied

ability to create a telnet connection to any machine inside the LAN by denying all traffic that comes in on port 23.

More sophisticated firewall systems may maintain internal information about the state of the traffic passing through them and/or the content of the data itself. The more a firewall can determine about the traffic, the more able it is to protect its users. Of course, this security comes at a price. Some sophisticated firewall approaches might create a noticeable delay in network traffic.

15.3 Network Addresses

When you communicate across a computer network, you ultimately communicate with one particular computer out of all possible computers in the world. There is a fairly sophisticated mechanism for identifying specific machines to establish that communication.

A hostname is a unique identification that specifies a particular computer on the Internet. Hostnames are generally readable words separated by dots. For example:

» Hostname A name made up of words separated by dots that uniquely identifies a computer on the Internet; each hostname corresponds to a particular IP address

» IP address An address made up of four numeric values separated by dots that uniquely identifies a computer on the Internet

```
matisse.csc.villanova.edu
condor.develocorp.com
```

We humans prefer to use hostnames when dealing with email addresses and websites because they are easy to use and remember. Behind the scenes, however, network software translates a hostname into its corresponding IP address, which is easier for a computer to use. An IP address is usually represented as a series of four decimal numbers separated by dots. For example:

```
205.39.155.18
193.133.20.4
```

An IP address is stored in 32 bits. Each number in an IP address corresponds to one byte in the IP address. Because one byte (8 bits) can represent 256 things, each number in an IP address is in the range 0 to 255. See Figure 15.9.

It's tempting to assume that because both hostnames and IP addresses are separated into sections by dots, there is a correspondence between the sections. That is not true. An IP address always has four values, but hostnames can have a variety of sections.

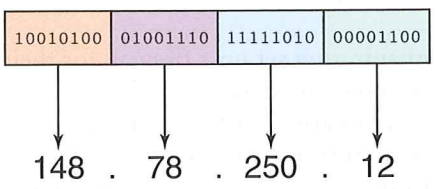

FIGURE 15.9 An IP address is stored in four bytes