Chapter 1

Understanding TCP/IP Basics

In This Chapter

- Protocols in this chapter: IP, TCP, IPSec, PPTP, L2TP
- Introducing TCP/IP
- Defining a protocol
- Understanding RFCs — the protocol documentation
- Differentiating between intranets, extranets, and Virtual Private Networks (VPNs)
- Figuring out who's in charge of TCP/IP and the Internet
- Investigating different types of networks that rely on TCP/IP software

You bought or borrowed this book, or maybe you're just flipping through it to pick up some information and tips about TCP/IP and its pieces and parts. Transmission Control Protocol/Internet Protocol, or TCP/IP, is the internationally accepted software for networking in general and, specifically, for making the Internet's services possible.

As you read this book, you get a behind-the-scenes look at how TCP/IP makes the Internet work. You also see how to use TCP/IP to set up your own home, office, or even international network. This chapter gets started by defining a protocol in general and TCP/IP protocols specifically. Proposals known as Requests for Comment, or *RFCs,* document how TCP/IP should function. You may wonder who's in charge of defining these protocols that rule the Internet. The answer is: lots of people who join international committees. This chapter describes the main Internet governing committees and what they do.

The Internet is one giant worldwide network that consists of tens of thousands of other networks. We give you an idea in this chapter of the different kinds of networks that connect via TCP/IP into the Internet.

The TCP/IP pronunciation guide

Pronouncing TCP/IP is easy — you just say the name of each letter and ignore the slash (/). Ready? It sounds like this:

"Tee cee pee eye pee"

Skip the silly jokes, please. We've made them all. By the way, some people find five letters too much to pronounce, so they just say "IP" to refer to the whole thing.

Following Rules for the Internet: TCP/IP Protocols

A *protocol* is a set of behavior-related rules that people follow. Some protocols are formally defined. For example, when people meet and greet each other, they might say, "Enchante de faire votre conaissance" or "How do you do"? We also hear our niece, Emily, and her friends saying "Hey, dude!" All these examples are widely accepted behaviors for people to start communicating — they are protocols. The more formal greetings are written down in etiquette books. "Hey, dude" has become accepted (at least by people much younger than we are) because of its wide use. Common ways of connecting aren't enough, though. After you meet, you need a common language in order to communicate. Just as people connect and communicate in accepted ways, computers connect and communicate with each other and with you. In the world of computers and networks, TCP/IP is a common language used for both connection and communication.

Although TCP/IP sounds like it consists of just two protocols, it's a whole set of protocols for connecting computers to the Internet. This set of protocols is the TCP/IP *stack,* or *protocol suite.* We describe in Chapter 2 the most well-known protocols in the TCP/IP stack. Before we get to the protocols themselves, the following sections look at who's in charge of the Internet and who decides what gets to be a standard part of the TCP/IP protocol suite. You also get familiar with Requests for Comments (RFCs), the documents that describe TCP/IP standards.

Who's in charge of the Internet and TCP/IP?

You're in charge. Or, you might say that everyone is, and no one is, in charge of the Internet and TCP/IP. No one person, organization, corporation, or government owns or controls the TCP/IP protocols or the Internet. Moreover, no

one person, organization, corporation, or government finances the TCP/IP protocols or the Internet. To say that no one controls TCP/IP and the Internet doesn't mean, however, that protocols magically appear with no control or that the Internet just does whatever it wants.

This list describes some of the important organizations and committees that *steer* TCP/IP and Internet policies:

- **Internet Society (ISOC):** The Internet Society (`www.isoc.org`) guides the future of the Internet by overseeing Internet standards, public policy, education, and training. ISOC members include corporations, international and governmental organizations, and individuals. The Internet Activities Board (refer to third bullet), the Internet Engineering Task Force (refer to fourth bullet), and the Internet Research Task Force are all part of the ISOC.
- **Internet Corporation for Assigned Names and Numbers (ICANN):** The nonprofit corporation ICANN, at `www.icann.org`, is in charge of assigning Internet addresses. ICANN, pronounced "eye can," is run by an international board of directors and funded by the Internet community.
- **Internet Activities Board (IAB):** IAB, at `www.iab.org`, defines the architecture for the Internet. The IAB — just say its letters, "i-a-b" — also oversees the Internet's protocols (TCP/IP). The IAB contains subcommittees of volunteers who set standards and work on new solutions to Internet growth problems.
- **Internet Engineering Task Force (IETF):** IETF, at `www.ietf.org`, is a community of more than 70 informal committees responsible for keeping the Internet up and running every day. The IAB supervises the IETF, which is pronounced simply "i-e-t-f." You can join the IETF working groups to help draft and develop standards for TCP/IP protocols.

Figure 1-1 shows how these Internet management groups are organized.

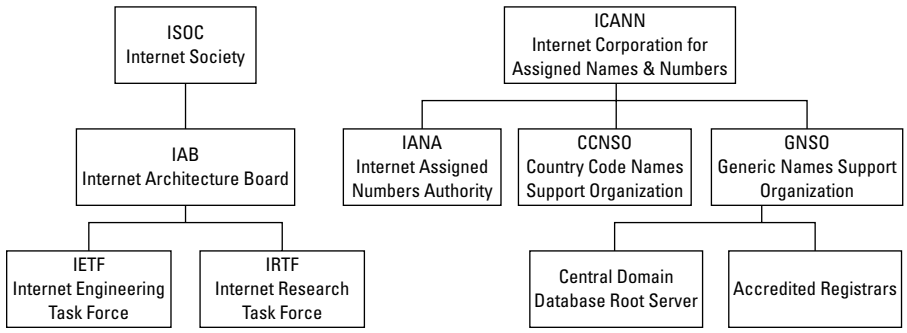

Figure 1-1: ISOC and ICANN are influential Internet steering groups.

Checking out RFCs: The written rules

TCP/IP protocols are written down in special Request for Comments (RFC) documents. An RFC (pronounced "r-f-c") document is available for everyone to read and comment on — it's part of the democracy of the Internet.

Toasting the RFC Editor

Surprise! The RFC Editor isn't just one person. It consists of a small group of people who work for the Internet Society. The RFC Editor Web site, at `www.rfc-editor.org`, keeps the official index of all RFCs ever written. You can find any RFC there. We find this site to be one of the most useful when we want information about what's going on with TCP/IP. You can search RFCs by number, author, title, or keyword. For example, click the link Search for an RFC and Its Meta-Data and then search for the keyword **security**. Notice how many pages it takes to display the results. And the list of results only grows — an RFC is never removed. It may be declared obsolete, but it stays available.

Knowing who writes RFCs

If you come up with an idea for a new or an improved capability for TCP/IP, you write your proposal as an RFC and submit it to an Internet committee for review. Working groups from various committees collaborate on most RFCs. You can join these working groups if you want to help but don't want to write a whole RFC on your own. For example, to join an IETF working group, send an e-mail to `Iptel-request@ietf.org`.

Understanding RFC categories

Three categories of RFCs are on the standards track:

- **Standard (STD):** An approved technical standard
- **Draft standard:** On its way to being adopted as a standard
- **Proposed standard:** On its way to being adopted as a draft standard

Here are some other RFC categories:

- **Best current practices (BCP):** Guidelines and recommendations, such as RFC 4107, "Guidelines for Cryptographic Key Management"
- **Experimental (EXP):** Part of a research or development project, such as RFC 5335, "Internationalized Email Headers"
- **Historic:** Refers to the fact that most historic RFCs are former standards that are now obsolete and have been replaced by more current RFCs
- **Informational (FYI):** Provides general information, such as RFC 4677, "The Tao of IETF — A Novice's Guide to the Internet Engineering Task Force"

If you have time and a sense of humor, check out the RFCs written on April 1, but *do not* take them seriously!

Examining Other Standards Organizations That Add to the Rules

Although the Internet corporations, committees, and groups listed in the preceding section specify the rules for using TCP/IP, other groups set standards for related technologies, as described in this list:

- **Institute of Electrical and Electronics Engineers (IEEE):** The IEEE (pronounce it "eye-triple-e") sets hardware standards, such as the hardware that connects Local Area Networks (LANs) and Wireless Local Area Networks (WLANs).
- **World Wide Web Consortium (W3C):** Although the Web is part of the Internet and follows TCP/IP standards, the W3C (say the letters and number "w-c-3") sets standards related to Web services.
- **International Organization for Standardization (ISO):** ISO ("eye-so") sets all kinds of standards, not just for networks. One of its standards indicates how the computers that run your car should interconnect.
- **Open Systems Interconnection (OSI):** The OSI ("o-s-i") sets networking protocol standards similar to TCP/IP, but different. At one time, OSI thought that its protocols would replace TCP/IP, but as hard as its members worked, it didn't happen.
- **Free Software Foundation (FSF) General Public License (GPL):** The FSF set up the GNU (pronounced "guh-new") project to create and distribute free software. GNU software, licensed under the GPL, is the reason that the Linux operating system is available for free or for a very low cost. GNU also provides lots of network tools and utilities as well as complete TCP/IP stacks.

Distinguishing Between the Internet, an Internet, and an Intranet

Yes, we realize that you already know what the Internet is. But just so that we're all using the same definition, the *Internet* is the worldwide collection of interconnected computer networks that use the TCP/IP protocol. These networks reach every continent — even Antarctica — and nearly every country.

The Internet also consists of much more than its network connections. It's all the individual computers connected to those individual networks, plus all the users of those computers, all the information accessible to those users, and all the knowledge those people possess. The Internet is just as much about people and information as it is about computers and computer networks.

Although the Internet is public, many organizations (companies and universities, for example) have their own, private internets that may connect to it. *An* internet is built the same way as *the* Internet, except that an internet is private. You might even have *an* internet in your home.

Both the Internet and internets run on TCP/IP protocol software. In this book, we distinguish *the* Internet from *an* internet by capitalizing *the* Internet.

The difference between *an* internet and an intranet is just terminology. The term *intranet* is fairly recent. Old-timers (such as the authors of this book) grew up with "a*n* internet" and now we use both terms. The important concept is that all kinds of "nets" run with TCP/IP.

Extending Intranets to Extranets

Intranets are the building blocks of extranets. If part of your intranet is available to people outside your organization, such as customers and suppliers, the part you share with the outside world is an *extranet.* An extranet has these characteristics:

- It consists of multiple, interconnected intranets/internets.
- An organization's extended family of partners work together electronically.
- It might not exist physically — it's a *virtual* network.

Because an *intranet* is a private network within an organization or a department, you might find a few different intranets in a large institution. A university on the east coast, for example, might have one intranet for its medical school, another intranet for its college of liberal arts, and a third intranet for its business school. That university may also network those intranets into an even bigger intranet. Then, so that the university community can reach the rest of the world, the university intranet needs to be connected to the (capital *I*) Internet.

When that university needs to share data with a different university on the west coast, the two universities can link their respective intranets to create an *extranet.* Figure 1-2 shows how the east and west coast universities form an extranet.

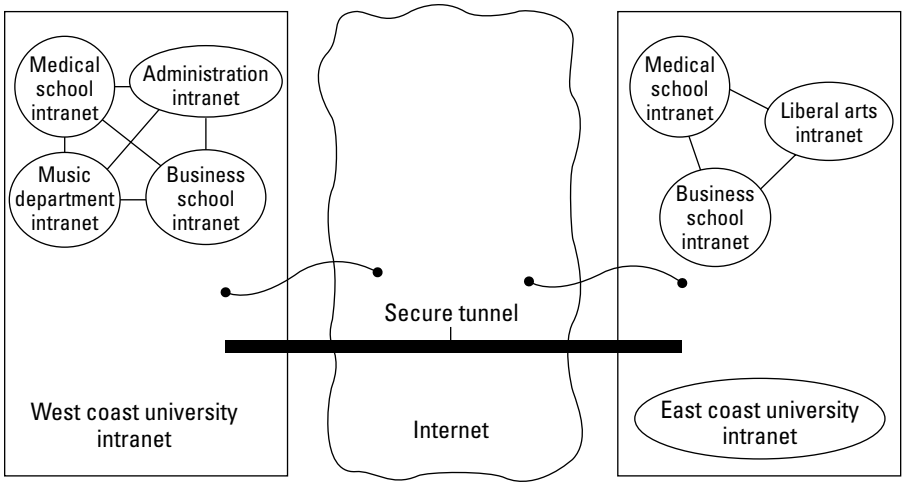

Figure 1-2: Intranets link to form an extranet.

An extranet consists of as many intranets as you need in order to communicate with your partners.

Introducing Virtual Private Networks

A *Virtual Private Network,* or *VPN* ("v-p-n"), is a private network that runs over public facilities, such as the Internet. Although it may seem like a contradiction to run a private network over the (very) public Internet, it works. In the olden days of computers (which is often six months ago, but we're talking as long as five years ago), if you wanted to work away from your office, you usually used a very slow modem to dial in across your phone line to the office computer. This method was slow and not secure because bad people could steal the data you were sending and receiving across the telephone lines.

Nowadays, most telecommuters connect to their offices through VPNs. They let you work as though you're on-site when you're not. You run VPN client software to establish a secure connection over the Internet to your organization's network. It's just like being in the office.

A VPN

- Is safe and secure because it scrambles *(encrypts)* data before sending it over the public lines
- Uses special tunneling and security protocols on the public network

  See the section about the IPSec, PPTP, and L2TP protocols in Chapter 22 for more information.

- Saves money for a large organization's networks because sharing the public Internet is cheaper than leasing private telecommunication lines
- Connects both intranets and extranets

The extranet shown earlier, in Figure 1-2, is also a VPN.

Exploring Geographically Based Networks

Whether you're sending e-mail or browsing the Web, your data gets broken up into small pieces called packets. In other words, your data is "packetized" before it goes onto a network. Packets of data travel over many different kinds of geographical distances, ranging from local to global and beyond to space. TCP/IP doesn't care about earthly distance — just that your data gets where it's going. In this section, get ready for a lot of jargon-y terms that look a lot alike. If you aren't interested in network architecture, feel free to skip this section and save your brain from getting muddled.

Networks connected by wires and cables

Networks come in different shapes and sizes. Two main architectures for networks — LANs (Local Area Networks) and WANs (Wide Area Networks) — are usually based on these factors:

- The distance the network covers
- Architecture and connection media
- Speed
- Purpose

  (For example, does the network connect a city, a campus, or just a bunch of storage devices?)

Exploring LANs

Pronounce LAN as a word — "lan" (rhymes with "pan"). The computers and other devices in a LAN communicate over small geographical areas, such as these:

- Your home office — or even the whole house
- One wing of one floor in a building

- Maybe the entire floor, if it's a small building
- Several buildings on a small campus

Incorporating WANs

Imagine a company that has several buildings in different towns and provinces, or even in different countries. Does that mean that all the people who work in the company can't be on the same network because a LAN is limited by distance? Of course not. The Internet is worldwide and beyond, so you can even bounce data off satellites in outer space, to create a WAN.

A WAN ("wan") spans geographical distances that are too large for LANs. Figure 1-3 shows two LANs connected to form a WAN.

Wireless networks

You don't need cables and wires to connect the computers that comprise a network. You can go wireless, and cables can be expensive. (Air, a wireless connection media, is free — at least for now.) Just as cabled LANs and WANs exist, wireless LANs (WLANs) and wireless WANs (WWANS) also exist.

You pronounce WLAN as the letter *w* followed by the word *LAN:* "double-you-lan." Pronounce WWAN as the letter *w* followed by the word *WAN:* ("double-you wan").

Although the following network technologies differ, your packets of data can fly through the air faster than Superwoman:

- **WLAN:** Uses radio waves to connect computers and networks. It shows up in homes, cafés, malls — even whole cities.
- **WWAN:** WWANs are based on telecommunications (mobile cellular networks) and use Worldwide Interoperability for Microwave Access (WiMAX) technology. A WWAN lets anyone with a computer work anywhere within a mobile phone network.

The geography of TCP/IP

TCP/IP fits everywhere. Regardless of your geographical network technology, in the end it's TCP/IP that carries your data, such as e-mail or Web pages, to you.

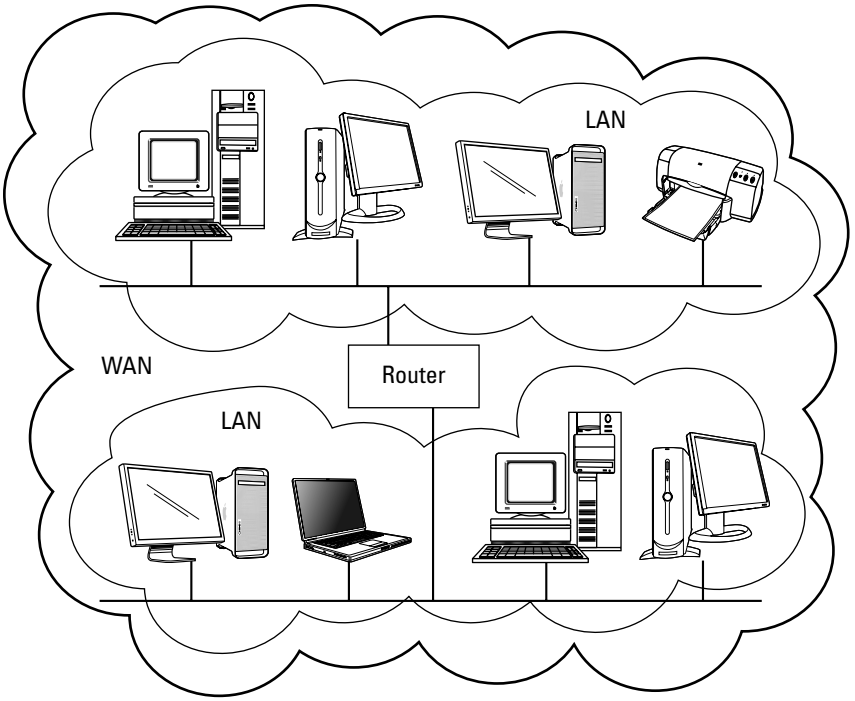

Figure 1-3: A special piece of hardware converts two LANs into a WAN.