4

HOW COMPUTER NETWORKS ARE ORGANIZED

In order to exchange files and messages through a network, all the computers connected to that network must use the same set of rules (known to network designers as a *protocol*). The rules that control the Internet are called *transmission control protocol/Internet protocol (TCP/IP).* Even if you don't plan to connect your network to the Internet right now, you should use TCP/IP for at least two reasons: First, TCP/IP is built into the Windows, Macintosh, and Linux operating systems and most inexpensive networking equipment; and second, you would waste a lot of time and money finding equipment that works with one of the other, older network protocols.

This chapter offers a relatively simple explanation of the TCP/IP protocols and how your network uses them.

TCP/IP Networks

TCP/IP is really a suite of protocols. The most important are TCP (transmission control protocol), which controls the way commands, messages, and files are broken into packets and reassembled at the other end, and IP (Internet protocol), which provides the rules that guide each data packet through different kinds of networks to the correct destination.

Your computer handles transmission control automatically, so you don't have to devote a lot of attention to individual data packets and their contents. The information in Chapter 2 of this book provides as much detail as most users ever need. But the Internet protocol is another matter; you should understand how your network (and just about every other network connected to the Internet) uses names and addresses for individual computers and other network nodes and how to use some of the standard software tools that are included in every network computer.

Fortunately, internal routing through the Internet is automatic; if you enter a valid address in your web browser, email client, or other program, the Internet will almost always find a path to the computer with that address. If it doesn't, the `ping` and `traceroute` commands described in "Network Tools" on page 41 will help you find the source of the problem.

Names and Addresses

An "addressing convention" sounds like an event where people attend speeches and workshops about house numbers and receive awards for sending out five million pieces of junk mail without an error. The formal sessions are often boring, but the after-hours parties are great. In networks, *addressing conventions* are actually the rules that everybody uses to identify the computers and other devices connected to a network and the people who use them. Every computer connected to a network has a unique name and address within that network, and every network connected to the Internet has its own unique numeric Internet address known as an *IP address*.

Numeric Addresses

The technical committees, international standards organizations, and government agencies that manage the Internet have all agreed on a 32-bit numeric address format shown as four numbers between 0 and 255, separated by periods, like this:

192.168.3.200

When you read an IP address out loud, you pronounce each digit separately and each period as "dot." So you would read this sample address as "one-nine-two dot one-six-eight dot three dot two-zero-zero."

You can think of an IP address as similar to your telephone number. Every computer connected to your LAN and every device or network connected to the Internet has a different address.

The agency responsible for assigning numeric IP addresses on the Internet is the *Internet Assigned Names Authority (IANA)*. Some formal contracts with the US government are involved, but the real reason IANA can provide this service to the worldwide Internet community is that everybody agrees to respect their assignments.

As the owner of a small LAN, you will never deal directly with IANA. Your Internet service provider controls a block of numeric addresses, and it will assign you one address (or more) when you set up your new connections.

Reserved Addresses

As Chapter 3 explained, your LAN communicates with other networks through a router. As far as the networks connected to that router are concerned, the router is just one more network connection with an IP address. Therefore, as Figure 4-1 shows, a router has two different IP addresses: one for its connection to the LAN and the other for the WAN or the Internet. The router presents a single address to the Internet that represents all the computers and other devices on your LAN; it performs a function called *network address translation (NAT)* that converts your public address to the addresses of individual network devices. One of the benefits of this system is that you can use the same IP addresses within your LAN as your neighbor across the street (or a LAN on the other side of the world), and the addresses won't interfere with one another.

In order to make this system work properly, IANA has reserved several blocks of IP address numbers for LANs; when a router receives a packet with an address in one of these ranges, it does not relay the packet to the Internet. If you use these addresses for the devices in your LAN, you can be certain that your packets (and the commands, messages, and files that make up those packets) won't end up at the reading room of the National Library of Ecuador when you wanted to send them to your assistant across the corridor.

The reserved IP addresses are:

10.0.0.0 to 10.255.255.255

169.254.0.0 to 169.254.255.255

172.16.0.0 to 172.31.255.255

192.168.0.0 to 192.168.255.255

Fixed and Dynamic Address Assignments (DHCP)

The computers and other nodes in your LAN can obtain their numeric IP addresses in one of two ways: The person who sets up the network connection can assign a permanent address, or a router or other network control device can automatically assign an address every time the device connects to the network. A permanent assignment is called a *fixed* or *static* IP address; an automatic assignment is a *dynamic* address.

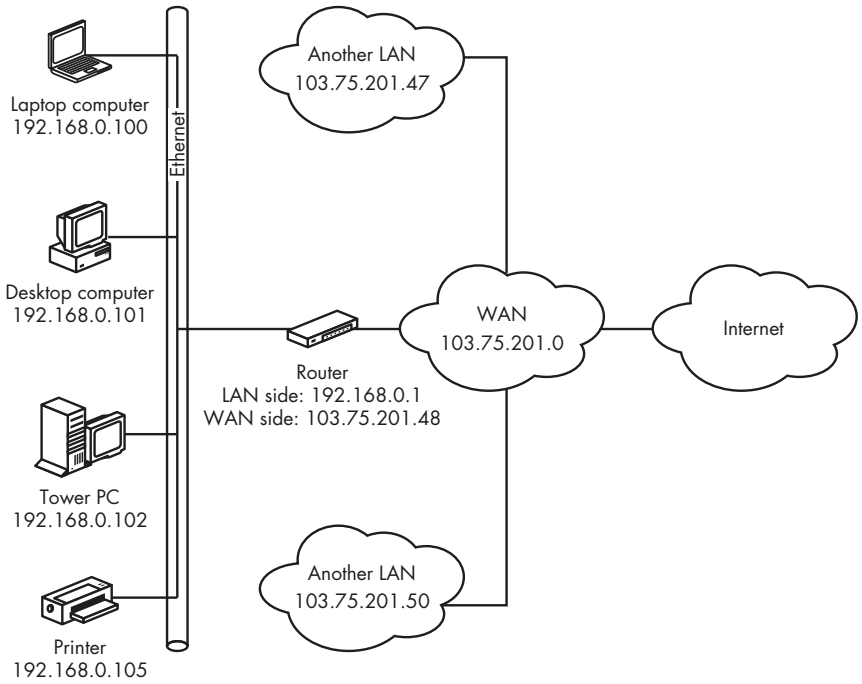

Figure 4-1: A router presents separate IP addresses to each network.

The method for assigning dynamic IP addresses is called *Dynamic Host Configuration Protocol (DHCP)*, so the device that makes the assignments is a *DHCP server*. In a LAN, the DHCP server uses numbers from the reserved range; on the Internet, the servers use numbers from a range provided to your ISP by IANA.

Both fixed and dynamic IP address assignments can work equally well, but all the devices on the network must use the same system; otherwise, more than one device might use the same number at the same time.

NOTE *If your LAN includes laptops and other portables that connect and disconnect from the network, DHCP is the better choice because it allows the network to assign an address automatically when a user connects and to re-use the same address after the first user has disconnected.*

Some Internet service providers and corporate network managers assign static IP addresses to each user, whereas others use DHCP to generate addresses. Chapters 10 and 11 explain how to set up your own computer and LAN to use either method.

The Domain Name System

Computers have no trouble handling long strings of numbers, but people often do. Addresses in the form of words rather than numbers are generally easier to remember and use. That's why the Internet and just about every LAN use names for each computer connected to a network. In a LAN, each

computer reads the name of every other device on the same network automatically; on the Internet, a computer called a *Domain Name System server (DNS server)* converts names to numeric addresses; when you type the name of a website into a browser, a DNS server finds the number that corresponds to that name and returns it to your browser, which connects to that numeric address.

You (or your network manager) will assign a name to each computer when you set up your network; your Internet service provider should set up a domain name for your connection to the Internet. Within a LAN, you can use simple descriptive names for each computer, such as "Sam" or "Kate."

On the other hand, the system for naming computers and networks connected to the Internet (rather than to your own LAN) follows some very specific rules called the *Domain Name System (DNS)*. In the Domain Name System, every name starts with a top-level domain name at the extreme right that can be either a generic description (such as *com*, *net*, or *edu*) or a two-letter country code (such as *uk* for the United Kingdom or *ca* for Canada). As you move to the left, the next word (or group of letters and numbers) is a name (called a *subdomain*) that has been reserved by a specific owner—an individual, a business, a government agency, or some other formal or informal organization. Large organizations might have one or more additional subdomain names to the left of the first one. Each part of the name is divided from the next one by a period (read as *dot*).

For example, the University of Washington's domain name is *washington.edu*. Within the university, the Department of Genome Science's address is *gs.washington.edu*. And within that department, the addresses of the research group studying evolutionary genetics is *evolution.gs.washington.edu*.

At the extreme left of a domain name, you will sometimes see a subdomain that identifies the type of server. This address might be the familiar *www* or some other Internet service such as *ftp (file transfer protocol)*.

Many addresses also include the *type* of Internet service (the *protocol*) that the web resource at that address uses as a leading part of the address, followed by a colon and two forward slashes (//), such as *http://host.sample.com/*. The *http* part stands for *HyperText Transfer Protocol*—the protocol that defines most websites. If you want to reach a different service at the same destination such as a file transfer server, a telnet host, or an Internet Relay Chat server, you might instead use *ftp://host.sample.com/*, *telnet://host.sample.com/*, or *irc://host.sample.com/*, respectively. When an address does not include the protocol type and the two forward slashes, your web browser will assume it's an http address. Some top-level domains that use country codes have other structures that differ from one country to another. Domain names that have a *us* (for United States) top-level domain sometimes use subdomains (also called *second-level domains*) that identify the state and city where the owner is located, such as *example.sf.ca.us*, which would be in San Francisco, California. In Canada and other countries, the domain name comes right before the country code (such as the Canadian Broadcasting Corporation's *cbc.ca*), whereas other countries use generic identifiers along with the geographic domain, such as *bbc.co.uk* for the British Broadcasting Corporation; the *co* stands for *commercial* and the *uk* for the *United Kingdom*.

NOTE *Just because a domain name address has a country code, the owner of that address is not necessarily located in that country. For example, many American FM radio stations have obtained addresses in the* .fm *domain, which belongs to the Federated States of Micronesia, and some television stations use the* .tv *domain assigned to the Pacific island nation of Tuvalu.*

Table 4-1 lists the most common generic top-level domains.

Table 4-1: Generic Top-Level Domains

Top-Level Domain	Used By
.com	Originally commercial, but now a generic domain
.net	Originally reserved for domains related to networks, but now a generic domain
.edu	Reserved for US colleges and universities
.org	Originally reserved for nonprofit organizations, but now a generic domain
.gov	Originally reserved for the US government, but now also used by state and local governments
.mil	Reserved for branches of the US military
.info	A generic domain with no restrictions
.biz	A generic domain restricted to businesses
.name	A generic domain reserved for individuals

Some other top-level domains such as *.asia*, *.coop*, *.museum*, and *.travel* are restricted to certain categories of users. Still others, such as

.测试

.испытание

.δοκιμή

.آزمایشی

are for addresses that don't use the Roman alphabet.

Name Servers

DNS name servers are an essential part of the Internet's internal plumbing, but most people don't know that they exist. If your computer can't find a DNS server, your email program, web browser, and other Internet programs won't work unless you use a numeric IP address to identify a destination.

DNS servers perform what seems like a simple task, but this task is more complicated than it first appears because millions of domain names are out there, and new ones are added all the time. Every DNS server in the world has to keep up with all the adds, moves, changes, and deletions. It accomplishes this through a system of *root servers* that are continuously

updated. If a local DNS server doesn't recognize a name, it consults the root server that keeps up with that name's top-level domain.

NOTE *There's actually a hierarchy of DNS servers, so a root server might end up consulting yet another server (and so on up the line) if it can't handle a name request itself.*

When you set up your computer for access to the Internet, you must specify the DNS servers that the computer will use to convert domain names to numeric IP addresses. In most cases, your Internet service provider or network manager will give you the numeric address of one or more nearby DNS servers. If your primary DNS server is not accessible, your computer will look for an alternate server if you have provided an alternate address.

It's generally best to use the DNS server address supplied by your ISP because the server with this address is probably closer to your own computer than any other server, and the system works best when total demand for DNS service is spread among as many servers as possible. But if you can't obtain reliable DNS service from your local service provider, a public DNS is often a useful alternative. You can find addresses for several public DNS servers though a Google or other web search for *Public DNS server*.

Some public DNS services can also provide some added features that your ISP might not offer. For example, OpenDNS (*http://www.opendns.com/*) can provide another layer of filtering against spyware, identity theft, adult sites, and other possible problems. It will also allow you to set up two three-letter shortcuts to frequently used addresses and will automatically correct common keystroke errors (such as typing *example.cmo* instead of *.com*). There's some controversy about some of these features, because they could lend themselves to returning names that are links to advertisements rather than the sites the original user requested.

Network Tools

You won't use them often if your LAN and your Internet connection are working properly, but you should know about a handful of troubleshooting tools that allow you to examine the innards of your network and its Internet connection.

All of these tools are simple text commands that you can use with just about any operating system. When you type a command, the system will display the results in the same window or screen. In Microsoft Windows, you can open a Command Prompt window after selecting Start ▸ Programs or by selecting Start ▸ Run and then typing **cmd**. In Mac OS X, select Applications ▸ Utilities and the Terminal program. If you're using Linux or Unix, use a command prompt or an XTerminal.

IPConfig

The IPConfig tool displays detailed information about your computer's current LAN and Internet connection, as shown in Listing 4-1.

```
C:\>IPConfig
Windows IP Configuration
Ethernet adapter Local Area Connection:
        Connection-specific DNS Suffix  . : domain.actdsltmp
        IP Address. . . . . . . . . . . . : 192.168.1.100
        Subnet Mask . . . . . . . . . . . : 255.255.255.0
        Default Gateway . . . . . . . . . : 192.168.1.1
```

Listing 4-1: The IPConfig tool displays the status of a computer's network configuration.

In this example, `Connection-specific DNS Suffix` is an address assigned by a DHCP host. This address is often an arbitrary name used internally within the network, but if your computer is connected directly to the Internet, it might be your computer's DNS address. If you try to connect to a domain name without a suffix (such as "example" rather than "example.net"), the network will assign this suffix to the address when it sends it to a DNS server.

The `IP Address` is the numeric address of this computer within the LAN or WAN. The `Subnet Mask` tells the network which parts of the numeric address identify individual computers, and the `Default Gateway` is the numeric address within the LAN of the gateway router that connects your LAN to the Internet.

For more details about your network connection, add `/all` to the command, as shown in Listing 4-2.

```
C:\>IPConfig /all

Windows IP Configuration

        Host Name . . . . . . . . . . . . : desktop
        Primary Dns Suffix  . . . . . . . :
        Node Type . . . . . . . . . . . . : Unknown
        IP Routing Enabled. . . . . . . . : No
        WINS Proxy Enabled. . . . . . . . : No
        DNS Suffix Search List. . . . . . : domain.actdsltmp

Ethernet adapter Local Area Connection:

        Connection-specific DNS Suffix  . : domain.actdsltmp
        Description . . . . . . . . . . . : Intel(R) PRO/100 VE Network
Connection
        Physical Address. . . . . . . . . : 00-0C-F1-AA-BF-BF
        Dhcp Enabled. . . . . . . . . . . : Yes
        Autoconfiguration Enabled . . . . : Yes
        IP Address. . . . . . . . . . . . : 192.168.1.100
        Subnet Mask . . . . . . . . . . . : 255.255.255.0
        Default Gateway . . . . . . . . . : 192.168.1.1
        DHCP Server . . . . . . . . . . . : 192.168.1.1
        DNS Servers . . . . . . . . . . . : 198.137.231.1
                                            206.63.63.1
        Lease Obtained. . . . . . . . . . : Wednesday, April 08, 2009 3:11:22 PM
        Lease Expires . . . . . . . . . . : Friday, April 10, 2009 3:11:22 PM
```

Listing 4-2: The `IPConfig /all` command displays additional information about your connection.

Obviously, this command produces a lot more information. The `Host Name` is the name that this computer uses on the LAN. The `Description` identifies the type of network interface adapter that connects this computer to the network. The `Physical Address` is the *MAC address*—the unique hardware identifier—of the network adapter. The `DHCP Server` is the address of the device that assigns IP addresses to other devices on the LAN (in this case, this device is the same as the `Default Gateway`), and the `DNS Servers` are the computers that this network consults to convert DNS addresses into numeric IP addresses. The `Lease Obtained` and `Lease Expires` lines show the date and time that this computer obtained its IP address from the DHCP server and the time the computer will give up that address; the host automatically renews the lease long before it expires, so you don't have to worry about the expiry time.

ifconfig

The `ifconfig` command is available in Macinstosh OS X and in Unix and Linux. This command displays information about the current network interface, including the connection type and the connection's current status. The format of the information display, however, varies in different operating systems. Therefore, the best place to find a detailed explanation of the `ifconfig` display produced by your own system is the man page for the `ifconfig` command.

ping

The `ping` command is an echo request. When you type `ping` *`target address`*, your computer sends a series of "please answer" messages to the target address, and that computer sends you a reply, as shown in Listing 4-3. Your computer measures the amount of time for each roundtrip and displays the duration in milliseconds.

```
C:\>ping nostarch.com

Pinging nostarch.com [72.32.92.4] with 32 bytes of data:

Reply from 72.32.92.4: bytes=32 time=140ms TTL=48
Reply from 72.32.92.4: bytes=32 time=99ms TTL=48
Reply from 72.32.92.4: bytes=32 time=99ms TTL=48
Reply from 72.32.92.4: bytes=32 time=97ms TTL=48

Ping statistics for 72.32.92.4:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 97ms, Maximum = 140ms, Average = 108ms
```

Listing 4-3: The `ping` command sends a series of echo requests to a designated address.

Many books and people will tell you that ping is an acronym for *Packet InterNet Groper*, but Mike Muuss, who wrote the original program, always insisted that he chose the name to imitate the sound of a sonar system

aboard a submarine; the sonar system makes an audible "ping" when an echo pulse returns from a target.

`ping` has several uses. It can confirm that the distant computer is alive, and that your computer's connection is working properly. It can also provide a rough idea of the network's performance (less time means higher speed). `ping` is also useful for finding a DNS problem; if you get a successful ping echo when you enter the target's numeric IP address, but not when you enter the domain name, the glitch is almost certainly someplace in the DNS system.

In Listing 4-3, it took about one-tenth of a second (100 ms) for each test to go from Seattle to San Francisco and back. That's a perfectly reasonable amount of time. But if one or more of the attempts had taken around 500 milliseconds or more, that would indicate some kind of problem.

Ping has also become a verb in computer jargon. You'll hear a technician at a help desk ask you to "ping me" at a specific address, meaning that you should send a ping request to that address. Some people have extended that usage beyond computer networks: They'll talk about "pinging" somebody when they intend to get that person's attention, either by email, telephone, or even poking their head into the recipient's office.

Many large commercial Internet sites, such as *yahoo.com* and *microsoft.com*, have chosen to block ping requests from outside their own network. If you get a `no reply` response to a ping request, try another address before you assume the problem is with your own Internet connection.

TraceRoute

The TraceRoute tool measures and displays the amount of time it takes for your computer to receive an echo from each network device between your computer and the target. As a result, a TraceRoute display can show you the route between your computer and any other computer on the Internet and pinpoint the segment of that route where a problem is occurring. In Windows, the command is `tracert`; in OS X, Linux and Unix, it's `traceroute`. TraceRoute sends three requests to each intermediate node, and shows the timing for each request.

Listing 4-4 shows a TraceRoute from my office in Seattle to No Starch Press in San Francisco.

```
C:\>tracert nostarch.com

Tracing route to nostarch.com [72.32.92.4]
over a maximum of 30 hops:

❶  1    4 ms     3 ms     3 ms  192.168.0.1
❶  2    3 ms     3 ms     3 ms  192.168.0.1
❷  3   71 ms    63 ms    64 ms  --..blv.nwnexus.net [206.63..]
❸  4   57 ms    53 ms    48 ms  fe000.cr1.sea.nwnexus.net [206.63.74.1]
   5    *       44 ms    42 ms  fe000.br4.sea.nwnexus.net [206.63.74.20]
   6   45 ms    43 ms    41 ms  204.181.35.197
   7   42 ms    42 ms    43 ms  sl-bb20-sea-4-0-0.sprintlink.net [144.232.6.121]
```

```
❹  8    86 ms    85 ms    87 ms  sl-bb25-chi-5-0.sprintlink.net [144.232.20.84]
   9    96 ms    97 ms    97 ms  sl-bb20-kc-2-0.sprintlink.net [144.232.20.108]
  10   108 ms   109 ms   107 ms  sl-crs1-fw-0-4-0-1.sprintlink.net [144.232.20.56]
❺ 11   110 ms   110 ms   108 ms  sl-st20-dal-1-0.sprintlink.net [144.232.9.136]
  12    99 ms    98 ms    97 ms  sl-racks-5-0.sprintlink.net [144.223.244.138]
  13   101 ms    98 ms    97 ms  vlan903.core3.dfw1.rackspace.com [72.3.128.53]
  14   101 ms   100 ms    98 ms  aggr115a.dfw1.rackspace.net [72.3.129.109]
  15    97 ms   100 ms    99 ms  squid14.laughingsquid.net [72.32.92.4]

Trace complete.
```

Listing 4-4: TraceRoute shows the path to a distant computer through the Internet.

In this case, it took 15 hops to complete the connection:

❶ The first two lines show the very fast response from the router sitting on the same table as the computer through a 6-foot cable. Line 2 repeats line 1 because of a software problem in the router.

❷ Line 3, whose domain name and IP address I have hidden, is my Internet service provider's WAN, a couple of miles away in downtown Seattle. Completing that echo takes longer, but it's still pretty fast.

❸ Lines 4 to 7 show the packets moving through various routers in the same switching center in Seattle.

❹ Starting at line 8, the route apparently jumps through routers in Chicago, Kansas City, Fort Worth, and Dallas, which increases the response times.

❺ The path moves around a routing center in Dallas at lines 11 through 15 until it ends up at the Laughing Squid web host that houses the No Starch web server.

This connection goes from origin to destination with several thousand miles of detours. However, the whole thing takes only about a tenth of a second, so those detours don't really matter.

TraceRoute can help identify several possible problems:

- If a TraceRoute report ends with one or more lines of asterisks (***), that usually isolates the problem to either the router named in the preceding line or the connection from that router to the next one.
- If the report shows a very long path that includes router addresses that don't seem to be a on a reasonable route (such as a path from New York to Philadelphia by way of Singapore), one of the network routers is not configured correctly.
- If the route shows that a pair of routers are passing the signal back and forth until TraceRoute times out, that usually indicates that one of those routers has lost a connection and is returning the signal back to the previous router. *That* router still thinks the best path is through the other one, so it tries again.

- If the report shows a long delay that always begins at the same router, it could indicate a problem with that router or very high demand for service through that part of the Internet.

Unless you're a network manager, you probably won't have to analyze TraceRoute reports very often. But if you're having a connection problem, they can sometimes help you to understand why you're not getting through to a website or instant message recipient.