

Allowing a computer to operate multiple clients is useful because services can be accessed simultaneously. For example, a user runs three applications at the same time: a web browser, an instant message application, and a video teleconference. Each application is a client that contacts one particular server independent of the other applications. In fact, the technology allows a user to have two copies of a single application open, each contacting a server (e.g., two web browser windows each contacting a different web site).

Allowing a given computer to run multiple server programs is useful for two reasons. First, using only one physical computer instead of many reduces the administrative overhead required to maintain the facility. Second, experience has shown that the demand for a service is usually sporadic — a given server often remains idle for long periods of time, and an idle server does not use the CPU. Thus, if the total demand for services is small enough, consolidating servers on a single computer can dramatically reduce cost without significantly reducing performance. To summarize:

*A single, powerful computer can offer multiple services at the same time; the computer runs one server program for each service.*

### 3.9 Server Identification And Demultiplexing

How does a client identify a server? The Internet protocols divide identification into two pieces:

- An identifier that specifies the computer on which a server runs
- An identifier that specifies a particular service on the computer

*Identifying A Computer.* Each computer in the Internet is assigned a unique identifier known as an *Internet Protocol address (IP address)*†. When it contacts a server, a client must specify the server's IP address. To make server identification easy for humans, each computer is also assigned a name, and the Domain Name System described in Chapter 4 is used to translate a name into an address. Thus, a user specifies a name such as *www.cisco.com* rather than an integer address.

*Identifying A Service.* Each service available in the Internet is assigned a unique 16-bit identifier known as a *protocol port number* (often abbreviated *port number*). For example, email is assigned port number 25, and the World Wide Web is assigned port number 80. When a server begins execution, it registers with its local system by specifying the port number for the service it offers. When a client contacts a remote server to request service, the request contains a port number. Thus, when a request arrives at a server, software on the server uses the port number in the request to determine which application on the server computer should handle the request.

Figure 3.4 summarizes the discussion by listing the basic steps a client and server take to communicate.

---

†Chapter 21 covers Internet addresses in detail and explains the addresses used with both IPv4 and IPv6.