Upon completion of this chapter, you will be able to perform the following tasks:

- Describe the process in which data is transferred from an application across a network.
- Given a network topology, identify the roles and functions of each network device and determine where each device best fits into the network.
- Given a network that combines switching, routing, and remote access, select the appropriate Cisco equipment.

CHAPTER 1

Internetworking Concepts Overview

The purpose of this chapter is to review basic internetworking concepts. These concepts are used throughout this book and are fundamental in understanding the functions of Cisco network devices.

Defining Network Components

The purpose of a data network is to help an organization increase productivity by linking all the computers and computer networks so that people have access to the information regardless of differences in time, location, or type of computer equipment.

Data networks have changed how we view our companies and employees. It is no longer necessary to have everyone in the same location in order to access the information needed to do the job. Because of this, many companies have changed their business strategy to utilize these networks in the way they do business. It is now typical for a company to organize the corporate internetwork in a way that allows it to optimize its resources. Figure 1-1 shows that the network is defined based on grouping employees (users) in the following ways:

- The main office is where everyone is connected to a LAN and where the majority of the corporate information is located. A main office could have hundreds or thousands of users who depend on the network to do their jobs. The main office might be a building with many local-area networks (LANs) or might be a campus of such buildings. Because everyone needs access to central resources and information, it is common to see a high-speed backbone LAN as well as a centralized data center with mainframe computers and application servers.
- The other connections are a variety of remote access locations that need to connect to the resources at the main offices and/or each other, including the following:
  - **Branch offices**—These are remote locations where smaller groups of people work. These users connect to each other via a LAN. In order to access the main office, these users access wide-area network (WAN) services. Although some information might be stored at the branch office, it

is likely that users will have to access much of the data from the main office. How often the main office network is accessed determines whether the WAN connection will be a permanent or dialup connection.

— **Telecommuters**—These are employees who work out of their homes. These users typically require an on-demand connection to the main office and/or the branch office in order to access network resources.

— **Mobile users**—These individuals work from various locations and rely on different services to connect to the network. While at the main or branch offices, these users connect to the LAN. When they are out of the office, these users usually rely on dialup services to connect to the corporate network.

Figure 1-1 *Corporate Networking Strategy*

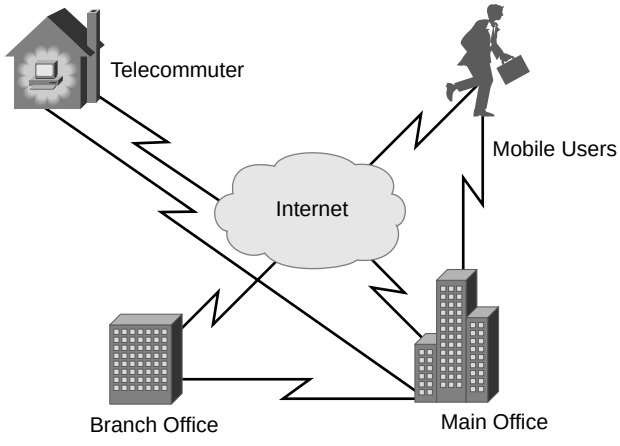

In order to understand what types of equipment and services to deploy in your network and when, it is important to understand business and user needs. You can then subdivide the network into a hierarchical model that spans from the end user's machine to the core (backbone) of the network. Figure 1-2 shows how the different employee groups interconnect.

To subdivide an internetwork into smaller components, Cisco uses a three-layer hierarchical model, as described in the following section.

Figure 1-2 *Group Interconnection*

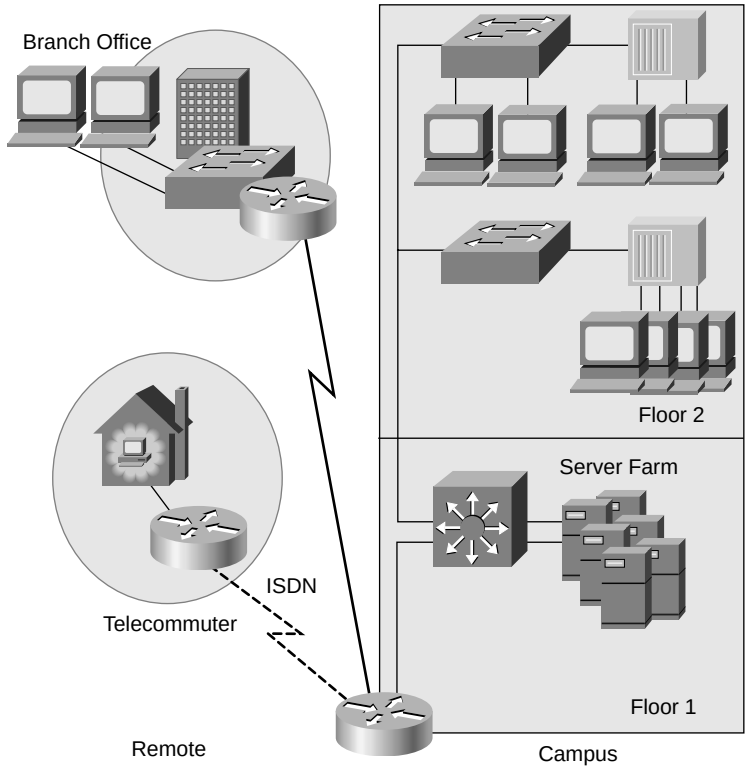

Mapping Business Needs to a Hierarchical Model

To simplify network designs, implementation, and management, Cisco uses a hierarchical model to describe the network. Although using this model is typically associated with designing a network, it is important to understand the model in order to know what equipment and features are needed in your network.

Campus networks have traditionally placed basic network-level intelligence and services at the center of the network and shared bandwidth at the user level. As businesses continue to place more emphasis on the network as a productivity tool, distributed network services and switching will continue to migrate to the desktop level.

User demands and network applications have forced networking professionals to use the traffic patterns in the network as the criteria for building an internetwork. Networks cannot be divided into subnetworks based only on the number of users. The emergence of servers that run global applications also has a direct impact on the load across the network. A higher

traffic load across the entire network results in the need for more efficient routing and switching techniques.

Traffic patterns now dictate the type of services needed by end users in networks. To properly build an internetwork that can effectively address a user's needs, a three-layer hierarchical model is used to organize traffic flow (see Figure 1-3).

Figure 1-3 *Three-Layer Hierarchical Network Model*

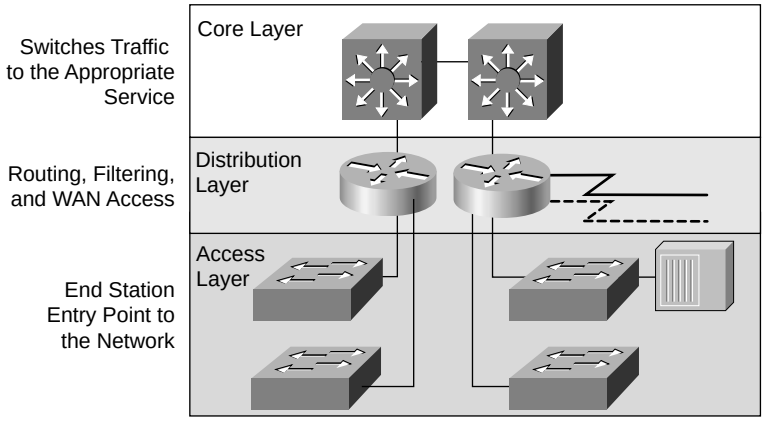

The model consists of three layers:

- Access
- Distribution
- Core

Each of these layers serves a function in delivering network services, as described in the following sections.

Access Layer

The access layer of the network is the point at which end users are connected to the network. This is why the access layer is sometimes referred to as the desktop layer. Users, and the resources they need to access most, are locally available. Traffic to and from local resources is confined between the resources, switches, and end users. Multiple groups of users and their resources exist at the access layer.

In many networks, it is not possible to provide users with local access to all services, such as database files, centralized storage, or dial-out access to the web. In these cases, user traffic for these services is directed to the next layer in the model, the distribution layer.

Distribution Layer

The distribution layer of the network (also referred to as the workgroup layer) marks the point between the access layer and the core services of the network. It is the primary function of this layer to perform functions such as routing, filtering, and WAN access. In a campus environment, the distribution layer represents a multitude of functions, including the following:

- Serving as an aggregation point for access layer devices
- Routing traffic to provide departmental or workgroup access
- Segmenting the network into multiple broadcast/multicast domains
- Translating between different media types, such as Token Ring and Ethernet
- Providing security and filtering services

The distribution layer can be summarized as the layer that provides policy-based connectivity, because it determines if and how packets can access the core services of the network. The distribution layer determines the fastest way for a user request (such as file server access) to be forwarded to the server. After the distribution layer chooses the path, it forwards the request to the core layer. The core layer then quickly transports the request to the appropriate service.

Core Layer

The core layer (also called the backbone layer) switches traffic as fast as possible to the appropriate service. Typically, the traffic being transported is to and from services common to all users. These services are referred to as global or enterprise services. Examples of these services are e-mail, Internet access, and videoconferencing.

When a user needs access to enterprise services, the request is processed at the distribution layer. The distribution layer device then forwards the user's request to the backbone. The backbone simply provides quick transport to the desired enterprise service. The distribution layer device provides controlled access to the core.

To properly build a network, you must first understand how your internetwork is used, your business needs, and your user needs. Those needs can then be mapped into a model that can be used to build your internetwork.

One of the best ways to understand how to build an internetwork is to first understand the way in which traffic is passed across the network. This is done through a conceptual network framework, the most popular of which is the OSI reference model. It is described in the following sections.

OSI Reference Model Overview

The OSI reference model serves several functions for the internetworking community:

- It provides a way to understand how an internetwork operates.
- It serves as a guideline or framework for creating and implementing network standards, devices, and internetworking schemes.

Here are some of the advantages of using a layered model:

- Breaks down the complex operation of networking into simple elements.
- Enables engineers to specialize design and development efforts on modular functions.
- Provides the capability to define standard interfaces for "plug-and-play" compatibility and multivendor integration.

As shown in Figure 1-4, the OSI reference model has seven layers. The four lower layers define ways for end stations to establish connections to each other in order to exchange data. The three upper layers define how the applications within the end stations will communicate with each other and with the users.

Figure 1-4 *OSI Reference Model*

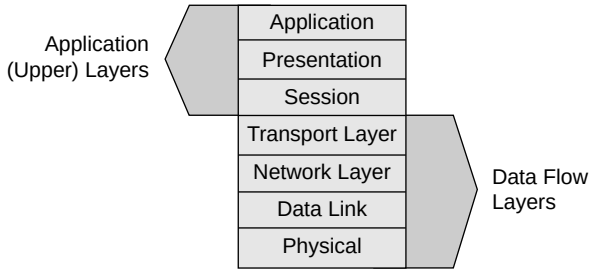

The following sections break down the layers and look at how they function to provide network connectivity.

Upper Layers

The three upper layers of the OSI reference model are often referred to as the *application* layers. These layers deal with the user interface, data formatting, and application access. Figure 1-5 shows the upper layers and provides information on their functionality with some examples.

Figure 1-5 *Upper Layers*

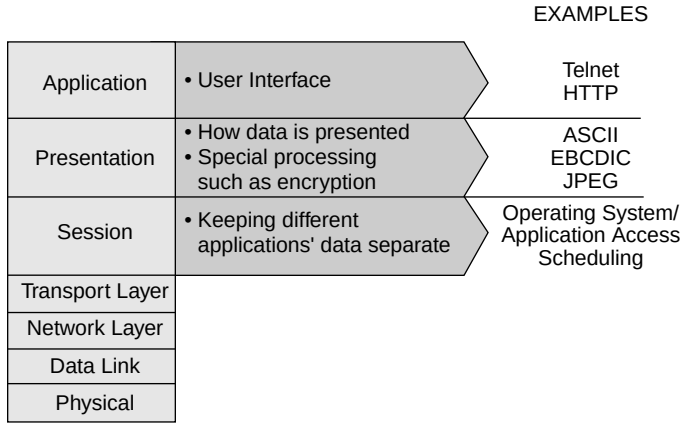

- **Application layer**—This is the highest layer of the model. It is the point where the user or application interfaces with the protocols to gain access to the network. For example, a word processor is serviced by file transfer services at this layer.
- **Presentation layer**—The presentation layer provides a variety of coding and conversion functions that are applied to application layer data. These functions ensure that data sent from the application layer of one system can be read by the application layer of another system. An example of coding functions is the encryption of data after it leaves an application. Another example is the jpeg and gif formats of images displayed on web pages. This formatting ensures that all web browsers, regardless of operating system, can display the images.
- **Session layer**—The session layer is responsible for establishing, managing, and terminating communications sessions between presentation layer entities. Communication at this layer consists of service requests and responses that occur between applications located in different devices. An example of this type of coordination would be between a database server and a database client.

Lower Layers

The four lower layers of the OSI reference model are responsible for defining how data is transferred across a physical wire, through internetwork devices, to the desired end station, and finally to the application on the other side. The focus of this book is Cisco's implementation of these layers. Figure 1-6 summarizes the basic functions of these four layers. We will discuss each layer in greater detail later in this chapter.

Figure 1-6 *Lower Layers*

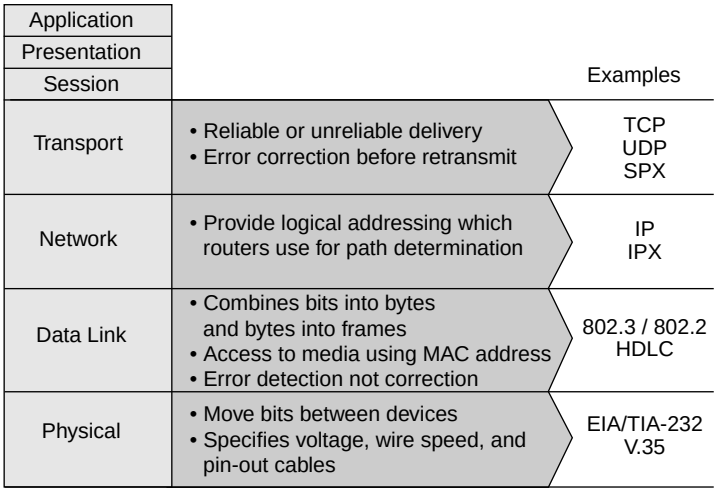

Communicating Between OSI Reference Model Layers

It is the responsibility of the protocol stack to provide communications between network devices. A protocol stack is the set of rules that define how information travels across the network. An example of this would be TCP/IP. The OSI reference model provides the basic framework common to most protocol stacks.

Each layer of the model allows data to pass across the network. These layers exchange information to provide communications between the network devices. The layers communicate with one another using protocol data units (PDUs). These PDUs control information that is added to the user data. The control information resides in fields called *headers* and *trailers*. In Figure 1-7, the Media Access Control (MAC) header and frame check sequence (FCS) at the data link layer represent a header and trailer.

Figure 1-7 *Data Encapsulation*

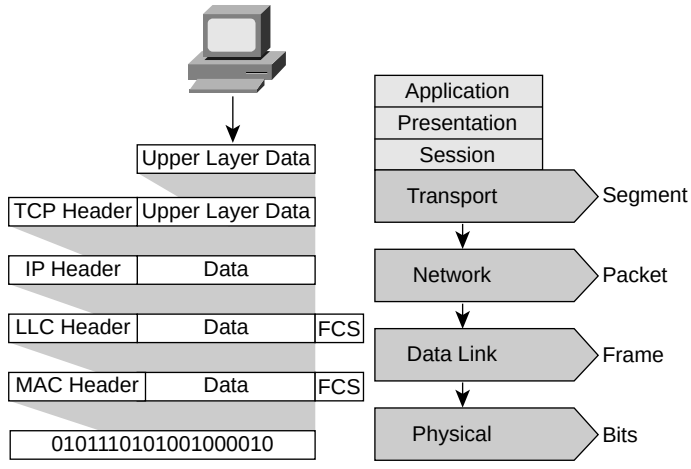

Because a PDU includes different information as it goes up or down the layers, it is given a name according to the information it is carrying. For example, in a TCP/IP stack (see Figure 1-7), after a transport layer TCP header has been added to the upper-layer data, that unit is called a *segment*. The segment is then passed down to the network layer, where an IP header is added, and it becomes a *packet*. The packet is packaged into a Layer 2 header, which becomes a *frame*. Finally, the frame is converted into bits, and the electrical signals are transmitted across the network media.

This method of passing data down the stack and adding headers and trailers is called *encapsulation*. After the data is encapsulated and passed across the network, the receiving device removes the information added, using the messages in the header as directions on how to pass the data up the stack to the appropriate application.

Data encapsulation is an important concept to networks. It is the function of like layers on each device, called *peer* layers, to communicate critical parameters such as addressing and control information.

Although encapsulation seems like an abstract concept, it is actually quite simple. Imagine that you want to send a coffee mug to a friend in another city. How will the mug get there? Basically, it will be transported on the road or through the air. You can't go outside and set the mug on the road or throw it up in the air and expect it to get there. You need a service to pick it up and deliver it. So, you call your favorite parcel carrier and give them the mug. But, that's not all. You need to give the carrier some information as to where the mug is going. So you provide the parcel carrier with an address and send the mug on its way. But first, the mug needs to be packaged. Here's the complete process:

Step 1 Pack the mug in a box.

Step 2 Place an address label on the box.

Step 3 Give the box to a parcel carrier.

Step 4 The carrier drives it down the road.

This process is similar to the encapsulation method that protocol stacks use to send data across networks. After the package arrives, your friend has to reverse the process. He takes the package from the carrier, reads the label to see who it's from, and finally opens the box and removes the mug. The reverse of the encapsulation process is known as de-encapsulation. Figure 1-8 represents the de-encapsulation process up a protocol stack.

Figure 1-8 *De-Encapsulation*

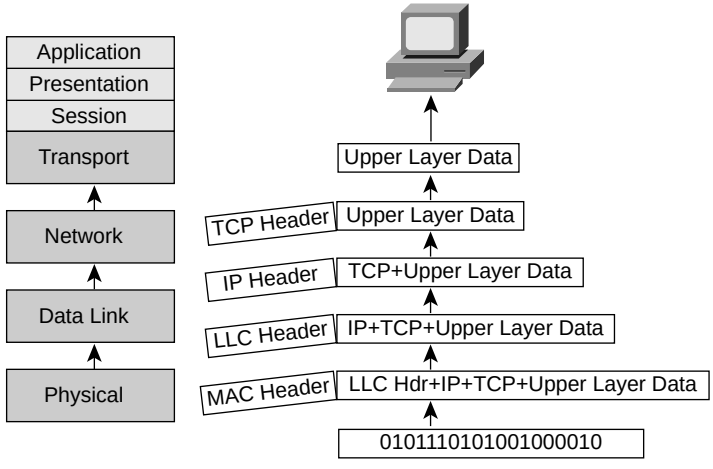

As networking professionals, it is our responsibility to implement networks that support the transport of user data. In order to implement and configure devices to do this, we must understand the processes of the lower layers of the OSI model. Understanding these processes makes configuring and troubleshooting network devices less troublesome.

Physical Layer Functions

To fully understand the network process, we must first closely examine each of the lower layers. Starting with the physical layer, shown in Figure 1-9, we will examine the function of each layer.

Figure 1-9 *Physical Layer*

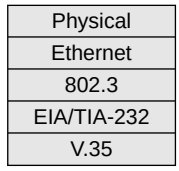

The physical layer defines the media type, connector type, and signaling type. It specifies the electrical, mechanical, procedural, and functional requirements for activating, maintaining, and deactivating the physical link between end systems. The physical layer also specifies characteristics such as voltage levels, data rates, maximum transmission distances, and physical connectors. In the analogy used earlier, the physical layer is the road on which the mug is carried. The roadway is a physical connection between different cities that allows us to go from one place to another. Different roads have different rules, such as speed limits or weight limits, just as different network media have different bandwidths or maximum transmission units (MTUs).

Physical Media and Connectors

The physical media and the connectors used to connect devices into the media are defined by standards at the physical layer. In this book, the primary focus is on the standards that are associated with Ethernet implementations.

The Ethernet and IEEE 802.3 (CSMA/CD) standards define a bus topology LAN that operates at a baseband signaling rate of 10 megabits per second (Mbps). Figure 1-10 shows three defined physical layer wiring standards, defined as follows:

- **10Base2**—Known as Thinnet. Allows network segments up to 185 meters on coaxial cable by interconnecting or chaining devices together.
- **10Base5**—Known as Thicknet. Allows network segments up to 500 meters on large coaxial cable with devices tapping into the cable to receive signals.
- **10BaseT**—Carries Ethernet signals up to 100 meters on inexpensive twisted-pair wiring back to a centralized concentrator called a *hub.*

Figure 1-10 *Defined Physical Layer 10Base Wiring Standards*

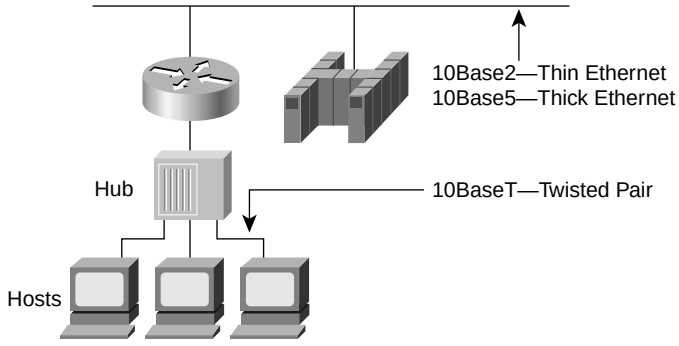

The 10Base5 and 10Base2 standards provide access for multiple stations on the same segment by physically connecting each device to a common Ethernet segment. 10Base5 cables attach to the bus using a cable and an attachment unit interface (AUI). 10Base2 networks chain devices together using coaxial cable and T connectors to connect the stations to the common bus.

Because the 10BaseT standard provides access for a single station at a time, each station must attach to a common bus structure to interconnect all the devices. The hub becomes the bus of the Ethernet devices and is analogous to the segment.

Collision/Broadcast Domains

Because all stations on an Ethernet segment are connected to the same physical media, signals sent out across that wire are received by all devices. This also means that if any two devices send out a signal at the same time, those signals will collide. The structure of Ethernet must therefore have rules that allow only one station to access the media at a time. There must also be a way to detect and correct errors known as *collisions* (when two or more stations try to transmit at the same time).

When discussing networks, it is critical to define two important concepts:

- **Collision domain**—A group of devices connected to the same physical media such that if two devices access the media at the same time, the result is a collision of the two signals
- **Broadcast domain**—A group of devices in the network that receive one another's broadcast messages

These terms help you understand the basic structure of traffic patterns and help define the needs for devices such as switches and routers.

Most Ethernet segments today are devices interconnected with hubs. Hubs allow the concentration of many Ethernet devices into a centralized device that connects all the devices to the same physical bus structure in the hub. This means that all the devices connected to the hub share the same media and, consequently, share the same collision domain, broadcast domain, and bandwidth. The resulting physical connection is that of a star topology as opposed to a linear topology. Figure 1-11 shows a common connection to the hub.

Figure 1-11 *Ethernet Hub*

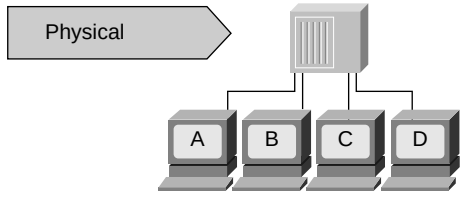

A hub does not manipulate or view the traffic that crosses that bus; it is used only to extend the physical media by repeating the signal it receives in one port out all the other ports. This means that a hub is a physical layer device. It is concerned only with propagation of the physical signaling, without any regard for upper-layer functions. This does not change the rules of Ethernet, however. Stations still share the bus of the hub, which means that contention still occurs.

Because all devices are connected to the same physical media, a hub is a single collision domain. If one station sends out a broadcast, the hub propagates it to all other stations, so it is also a single broadcast domain.

The Ethernet technology used in this instance is known as carrier sense multiple access collision detection (CSMA/CD). This means that multiple stations have access to the media, and before one station can access that media, it must first “listen” (carrier sense) to make sure that no other station is using the same media. If the media is in use, the station must wait before sending out any data. If two stations both listen and hear no other traffic, and then they both try to transmit at the same time, the result is a collision.

For example, in Figure 1-12, both cars try to occupy the same road at the same time, and they collide. In a network, as with cars, the resulting collision causes damage. In fact, the damaged frames become error frames, which the stations detect as a collision, forcing both stations to retransmit their respective frames. A backoff algorithm determines when the stations retransmit in order to minimize the chance of another collision. The more stations that exist on an Ethernet segment, the greater the chance that collisions will occur. These excessive collisions are the reason that networks are segmented (broken up) into smaller collision domains using switches and bridges.

Figure 1-12 *Ethernet Collisions*

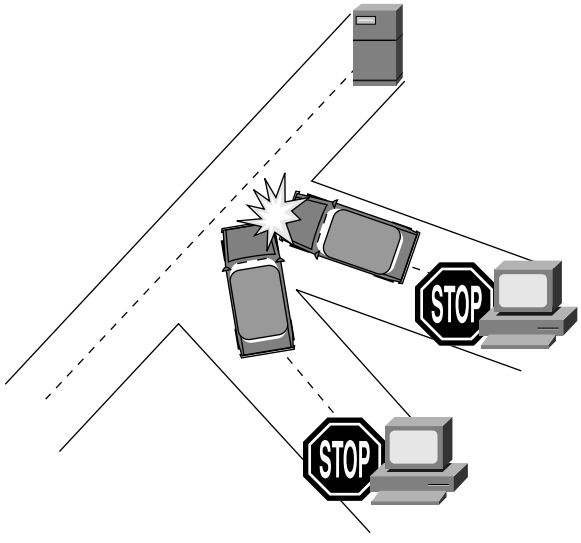

Data Link Layer Functions

Before traffic can be placed on the network, it must be given some details about where to go and what to do when it gets there. The data link layer provides this function. The data link layer is Layer 2 of the OSI reference model, and it differs depending on the topology. Figure 1-13 shows the various physical topologies and some corresponding data link encapsulation methods.

Figure 1-13 *Data Link Layer*

Physical	Data Link	
Ethernet		
802.3		802.2
EIA/TIA-232	HDLC	
V.35	Frame Relay	

The purpose of this layer is to provide the communications between workstations at the first logical layer above the bits on the wire. Because of this, many functions are provided by the data link layer. The physical addressing of the end stations is done at the data link layer to help the network devices determine whether they should pass a message up the protocol stack. Fields also exist in this layer to tell the device which upper-layer stack to pass the data to (such as IP, IPX, AppleTalk, and so on). The data link layer provides support for connection-oriented and connectionless services and provides for sequencing and flow control.

To provide these functions, the IEEE data link layer is defined by two sublayers:

- **Media Access Control (MAC) Sublayer (802.3)**—The Media Access Control (MAC) sublayer is responsible for how the data is transported over the physical wire. This is the part of the data link layer that communicates downward to the physical layer. It defines such functions as physical addressing, network topology, line discipline, error notification, orderly delivery of frames, and optional flow control.
- **Logical Link Control (LLC) Sublayer (802.2)**—The Logical Link Control sublayer is responsible for logically identifying different protocol types and then encapsulating them in order to be transmitted across the network. A type code or service access point (SAP) identifier does the logical identification. The type of LLC frame used by an end station depends on what identifier the upper-layer protocol expects. Additional LLC options include support for connections between applications running on the LAN, flow control to the upper layer, and sequence control bits. For some protocols, LLC defines reliable or unreliable services for data transfer, instead of the transport layer. (Reliable and unreliable services are discussed further in the section, "Transport Layer Functions.")

MAC Sublayer Frames

Figure 1-14 illustrates the frame structure for the MAC sublayer IEEE 802.3 frames.

Figure 1-14 *MAC Sublayer Frame*

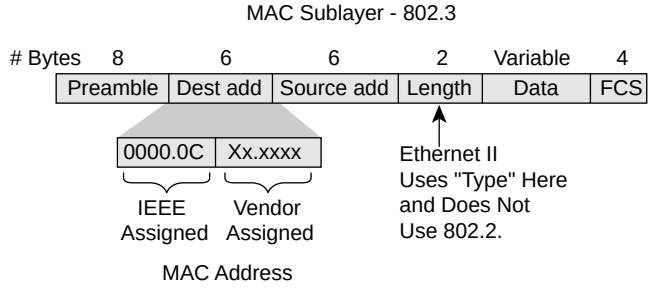

Figure 1-14 shows the standard frame structure to provide an example of how control information is used to transmit information at this layer. The definitions of the MAC sublayer fields are as follows:

- The IEEE 802.3 frame begins with an alternating pattern of 1s and 0s called a *preamble*. The preamble tells receiving stations that a frame is coming.
- Immediately following the preamble are the *destination* and *source physical address* fields. These addresses are referred to as *MAC layer addresses*. They are unique to each device in the internetwork. On most LAN interface cards, the MAC address is

burned into ROM, thus explaining the term burned-in-address (BIA). When the network interface card initializes, this address is copied into RAM to identify the device on the network.

The MAC address is a 48-bit address expressed as 12 hexadecimal digits. The first 24 bits or 6 hexadecimal digits of the MAC address contain a manufacturer identification or vendor code. Another name for this part of the address is the Organizationally Unique Identifier (OUI). To ensure vendor uniqueness, the IEEE administers OUIs. The last 24 bits or 6 hexadecimal digits are administered by each vendor and often represent the interface serial number.

The source address is always a unicast (single node) address, and the destination address might be unicast, multicast (group of nodes), or broadcast (all nodes).

- In IEEE 802.3 frames, the two-byte field following the source address is a *length* field, which indicates the number of bytes of data that follow this field and precede the frame check sequence (FCS) field.
- Following the length field is the *data* field, which includes the LLC control information, other upper-layer control information, and the user data.
- Finally, following the data field is a 4-byte *FCS* field containing a cyclic redundancy check (CRC) value. The CRC is created by the sending device and recalculated by the receiving device to check for damage that might have occurred to the frame in transit.

LLC Sublayer Frames

There are two LLC frame types: Service Access Point (SAP) and Subnetwork Access Protocol (SNAP). Which frame type your system uses depends on the applications that you have running on your system. Some applications define themselves by a SAP ID, and others define themselves using a type code. Figure 1-15 shows the format of the SAP and SNAP frame types.

Figure 1-15 *SAP and SNAP LLC Sublayer Frames*

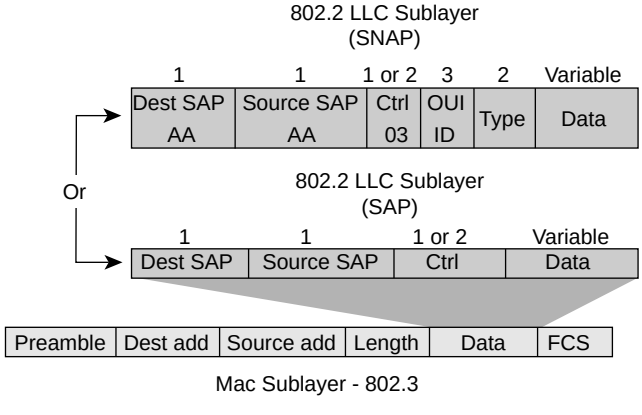

In the LLC header, the destination SAP (DSAP) and source SAP (SSAP) fields are 1 byte each and act as pointers to the upper-layer protocols in a station. For example, a frame with a SAP of 06 hex is destined for IP, and a frame with a SAP of E0 hex is destined for IPX. From the perspective of these lower MAC sublayers, the SAP process provides a convenient interface to the upper layers of the protocol stack. These SAP entries allow the physical and data link connections to provide services for many upper-layer protocols.

In order to specify that the frame uses SNAP, the SSAP and DSAP addresses are both set to AA hex, and the control field is set to 03 hex. In addition to the SAP fields, a SNAP header has a type code field that allows for the inclusion of the EtherType. The EtherType defines which upper-layer protocol receives the data.

In a SNAP frame, the first three bytes of the SNAP header after the control field are the OUI vendor code. Following the OUI vendor code is a two-byte field containing the EtherType for the frame. Here is where the backward compatibility with Ethernet Version II is implemented. As with the 802.3 frame, a 4-byte FCS field follows the data field and contains a CRC value.

Data Link Layer Devices

Bridges and Layer 2 switches are devices that function at the data link layer of the protocol stack. Figure 1-16 shows the devices typically encountered at Layer 2. Layer 2 switching is hardware-based bridging. In a switch, frame forwarding is handled by specialized hardware called application-specific integrated circuits (ASICs). ASIC technology allows a silicon chip to be programmed to perform a specific function as it is built. This technology allows functions to be performed at much higher rates of speed than that of a chip that is programmed by software. Because of ASIC technology, switches provide scalability to gigabit speeds with low latency.

Figure 1-16 *Data Link Devices*

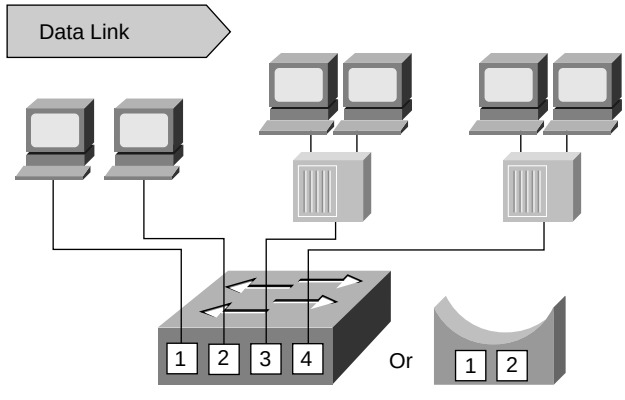

NOTE Although there are Layer 3 and Layer 4 switches that perform routing, this book uses the term *switch* to refer to a Layer 2 device.

When a bridge or switch receives a frame, it uses the data link information to process the frame. In a transparent bridge environment, the bridge processes the frame by determining whether it needs to be copied to other connected segments. A transparent bridge hears every frame that crosses a segment and views each frame and source address field to determine on what segment the source station resides. The transparent bridge stores this information in memory in what is known as a *forwarding table*. The forwarding table lists each end station (from which the bridge has heard a frame within a particular time period) and the segment on which it resides. When a bridge hears a frame on the network, it views the destination address and compares it to the forwarding table to determine whether to filter, flood, or copy the frame onto another segment.

This decision process occurs as follows:

- If the destination device is on the same segment as the frame, the bridge blocks the frame from going on to other segments. This process is known as *filtering*.
- If the destination device is on a different segment, the bridge forwards the frame to the appropriate segment.
- If the destination address is unknown to the bridge, the bridge forwards the frame to all segments except the one on which it was received. This process is known as *flooding*.

Because a bridge learns all the station destinations by listening to source addresses, it will never learn the broadcast address. Therefore, all broadcasts will always be flooded to all the segments on the bridge or switch. All segments in a bridged or switched environment are therefore considered to be in the same broadcast domain.

NOTE This book focuses on transparent bridging because this is the function performed by the Catalyst 1900 series of switches. This is also the most common form of bridging/switching in Ethernet environments. It should also be noted that there are other types of bridges, such as source-route bridging, in which the source determines the route to be taken through the network, and translational bridging, which allows the frame to move from a source route to a transparent environment between Ethernet and Token Ring.

A bridged/switched network provides excellent traffic management. The purpose of the Layer 2 device is to reduce collisions, which waste bandwidth and prevent packets from reaching their destinations. Part A of Figure 1-17 shows how a switch reduces collisions by giving each segment its own collision domain. Part B of Figure 1-17 shows that when two or more packets need to get onto a segment, they are stored in memory until the segment is available for use.

Figure 1-17 *Bridging Reduces Collisions*

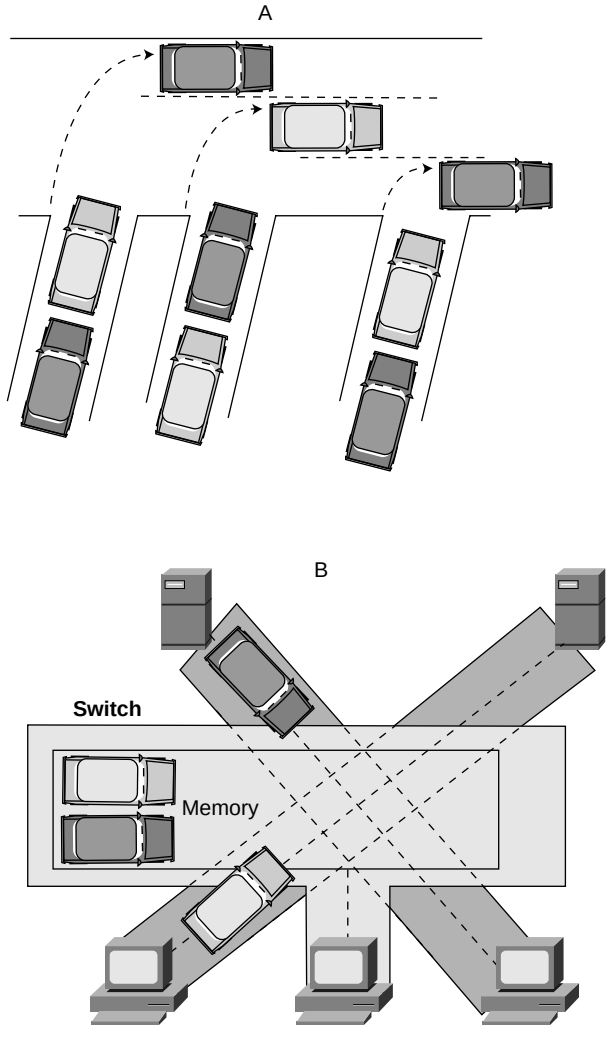

Bridged/switched networks have the following characteristics:

- Each segment is its own collision domain.
- All devices connected to the same bridge or switch are part of the same broadcast domain.
- All segments must use the same data link layer implementation, such as all Ethernet or all Token Ring. If an end station must communicate with another end station on different media, then some device, such as a router or translational bridge, must translate between the different media types.

- In a switched environment, there can be one device per segment, and each device can send frames at the same time, thus allowing the primary pathway to be shared.

Network Layer Functions

The network layer defines how to transport traffic between devices that are not locally attached in the same broadcast domain. Two pieces of information are required to achieve this:

- A logical address associated with the source and destination stations.
- A path through the network to reach the desired destination.

Figure 1-18 shows the location of the network layer in relation to the data link layer. The network layer is independent of the data link and can therefore be used to connect devices residing on different physical media. The logical addressing structure is used to provide this connectivity.

Figure 1-18 *Location of the Network Layer in the Protocol Model*

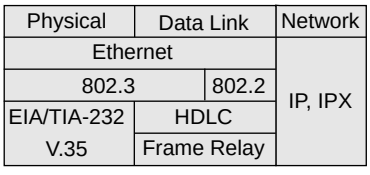

Logical addressing schemes are used to identify networks in an internetwork and the location of the devices within the context of those networks. These schemes vary based on the network layer protocol in use. This book discusses the network layer operation for the TCP/IP and IPX (Novell) protocol stacks.

Network Layer Addresses

Network layer addresses (also called *virtual* or *logical addresses*) exist at Layer 3 of the OSI reference model. Unlike data link layer addresses, which usually exist within a flat address space, network layer addresses are usually hierarchical in that they define networks first and then devices or nodes on each of those networks. In other words, network layer addresses are like postal addresses, which describe a person's location by providing a ZIP code and a street address. The ZIP code defines the city and state, and the street address is a particular location in that city. This is in contrast to the MAC layer address, which is flat in nature. A good example of a flat address space is the U.S. Social Security numbering system, in which each person has a single, unique Social Security number. Figure 1-19 shows a sample logical address as defined within a network layer packet.

Figure 1-19 *Network Layer Logical Addressing*

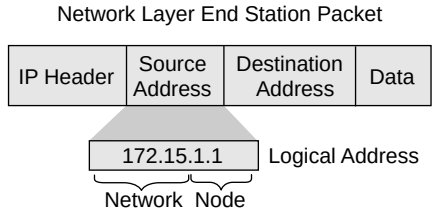

The logical network address consists of two portions. One part uniquely identifies each network within the internetwork, and the other part uniquely identifies the hosts on each of those networks. Combining both portions results in a unique network address for each device. This unique network address has two functions:

- The network portion identifies each network in the internetwork structure, allowing the routers to identify paths through the network cloud. The router uses this address to determine where to send network packets, in the same manner that the ZIP code on a letter determines the state and city that a package should be delivered to.
- The host portion identifies a particular device or a device's port on the network in the same manner that a street address on a letter identifies a location within that city.

There are many network layer protocols, and they all share the function of identifying networks and hosts throughout the internetwork structure. Most of these protocols have different schemes for accomplishing this task. TCP/IP is a common protocol that is used in routed networks. An IP address has the following components to identify networks and hosts:

- A 32-bit address, divided into four 8-bit sections called *octets*. This address identifies a specific network and a specific host on that network by subdividing the bits into network and host portions.
- A 32-bit subnet mask that is also divided into four 8-bit octets. The subnet mask is used to determine which bits represent the network and which represent the host. The bit pattern for a subnet mask is a string of recursive 1s followed by the remaining bits, which are 0. Figure 1-20 shows that the boundary between the 1s and the 0s marks the boundary for the network and host portions of the address, the two components necessary to define an IP address on an end device.

Figure 1-20 *IP Address Components*

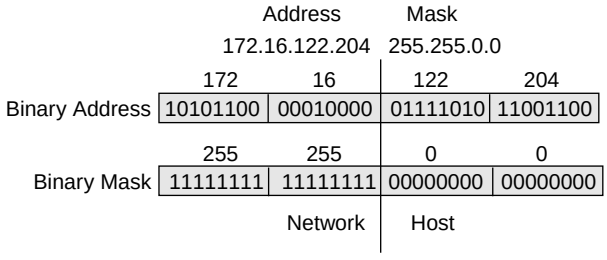

NOTE IP addresses are represented by taking the 8-bit octets and converting them to decimal and then separating the octets with dots or periods. This format is known as *dotted decimal* and is done to simplify addressing for those of us who count in Base10.

Router Operation at the Network Layer

Routers operate at the network layer by tracking and recording the different networks and choosing the best path to those networks. The routers place this information in a routing table, which includes the following items (see Figure 1-21):

- **Network addresses**—Represent known networks to the router. A network address is protocol-specific. If a router supports more than one protocol, it will have a unique table for each protocol.
- **Interface**—Refers to the interface used by the router to reach a given network. This is the interface that will be used to forward packets destined for the listed network.
- **Metric**—Refers to the cost or distance to the target network. This is a value that helps the router choose the best path to a given network. This metric changes depending on how the router chooses paths. Common metrics include the number of networks that must be crossed to get to a destination (also known as *hops*), the time it takes to cross all the interfaces to a given network (also known as *delay*), or a value associated with the speed of a link (also known as *bandwidth*).

Figure 1-21 *Routing Tables*

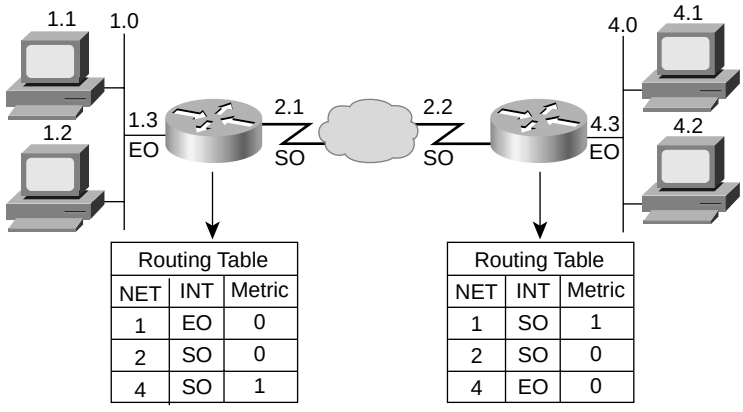

Because routers function at the network layer of the OSI model, they are used to separate segments into unique collision and broadcast domains. Each segment is referred to as a *network* and must be identified by a network address to be reached by end stations. In addition to identifying each segment as a network, each station on that network must also be uniquely identified by the logical address. This addressing structure allows for hierarchical network configuration (that is, a station is not known merely by a host identifier) but is defined by the network it is on as well as a host identifier. In order for routers to operate on a network, it is required that each interface be configured on the unique network it represents. The router must also have a host address on that network. The router uses the interface's configuration information to determine the network portion of the address to build a routing table.

In addition to identifying networks and providing connectivity, routers also provide other functions:

- Routers do not forward Layer 2 broadcast or multicast frames.
- Routers attempt to determine the optimal path through a routed network based on routing algorithms.
- Routers strip Layer 2 frames and forward packets based on Layer 3 destination addresses.
- Routers map a single Layer 3 logical address to a single network device; therefore, routers can limit or secure network traffic based on identifiable attributes within each packet. These options, controlled via access lists, can be applied to inbound or outbound packets.
- Routers can be configured to perform both bridging and routing functions.

- Routers provide connectivity between different virtual LANs (VLANs) in a switched environment.
- Routers can be used to deploy quality of service parameters for specified types of network traffic.

In addition to the benefits in the campus, routers can be used to connect remote locations to the main office using WAN services, as illustrated in Figure 1-22.

Figure 1-22 *Routers Connect Remote Locations to the Main Office*

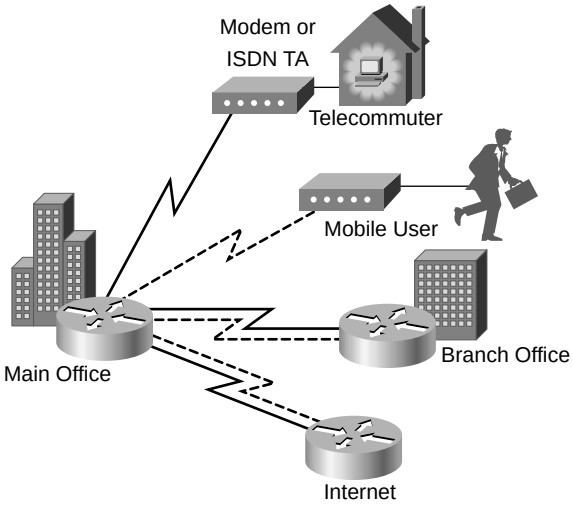

Routers support a variety of physical layer connectivity standards that allow you to build WANs. In addition, they can provide the security and access controls that are needed when interconnecting remote locations.

Transport Layer Functions

In order to connect two devices in the fabric of the network, a connection or session must be established. The transport layer defines the end-to-end station establishment guidelines between two end stations. A session constitutes a logical connection between the peer transport layers in source and destination end stations. Figure 1-23 shows the relationship of some transport layer protocols to their respective network layer protocols. Different transport layer functions are provided by these protocols.

Figure 1-23 *Transport Layer Protocols*

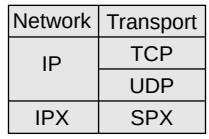

Specifically, the transport layer defines the following functions:

- Allows end stations to assemble and disassemble multiple upper-layer segments into the same transport layer data stream. This is accomplished by assigning upper-layer application identifiers. Within the TCP/IP protocol suite, these identifiers are known as *port numbers*. The OSI reference model refers to these identifiers as Service Access Points (SAPs). The transport layer uses these port numbers to identify application layer entities such as FTP and Telnet. An example of a port number is 23, which identifies the Telnet application. Data with a transport port number of 23 would be destined for the Telnet application.
- Allows applications to request reliable data transport between communicating end systems. Reliable transport uses a connection-oriented relationship between the communicating end systems to accomplish the following:
  - — Ensure that segments delivered will be acknowledged back to the sender.
  - — Provide for retransmission of any segments that are not acknowledged.
  - — Put segments back into their correct sequence order at the receiving station.
  - — Provide congestion avoidance and control.

At the transport layer, data can be transmitted reliably or unreliably. For IP, the TCP protocol is reliable or connection-oriented, and UDP is unreliable or connectionless. A good analogy to connection-oriented versus connectionless is a phone call versus a post card. With a phone call, you establish a dialogue that lets you know how well you are communicating. A post card offers no real-time feedback.

In order for a connection-oriented transport layer protocol to provide these functions reliably, a connection must be established between the end stations, data is transmitted, and then the session is disconnected.

Like a phone call, in order to communicate with a connection-oriented service, you must first establish the connection. To do this within the TCP/IP protocol suite, the sending and receiving stations perform an operation known as a three-way handshake (see Figure 1-24). A three-way handshake is accomplished by the sending and receiving of synchronization and acknowledgment packets. With a phone call, this would be like each party saying “hello” to indicate that they were ready to talk.

Figure 1-24 *The Three-Way Handshake*

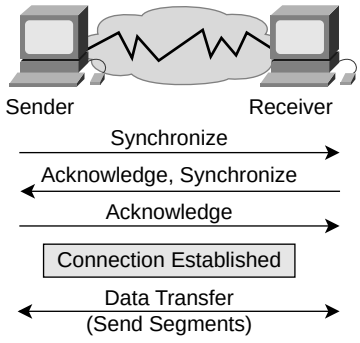

After the synchronization has occurred, the transfer of information begins. During the transfer, the two end stations continue to communicate with their network layer PDUs (headers) to verify that the data is received correctly. If the receiving station does not acknowledge a packet within a predefined amount of time, the sender retransmits the package. This ensures reliable delivery of all traffic. After the data transfer is complete, the session is disconnected, like saying "good-bye" during a telephone conversation.

OSI Lower Layer Review

Now that we have defined and discussed the lower four layers of the OSI model and defined the concepts of collision and broadcast domains, let's review what we have learned.

Each device shown in Figure 1-25 operates at a different layer of the OSI model:

- At Layer 1 (the physical layer) is the hub. The hub retransmits our packets and acts as a concentration device for our other network devices. The hub forms a single segment, providing one collision domain and one broadcast domain.
- The switch and the bridge are Layer 2 devices. These devices divide our network into separate segments, providing fewer users per segment. Each segment is a single collision domain, so in the figure, the bridge and switch each support four collision domains. Broadcast traffic, however, propagates across all segments, so only one broadcast domain is associated with each device.

- At Layer 3 (the network layer), the router provides paths to all the networks throughout the internetwork. The router segments the network into separate collision domains and broadcast domains. In Figure 1-25, we see that there are four collision domains and four broadcast domains.

Figure 1-25 *Network Device Functions*

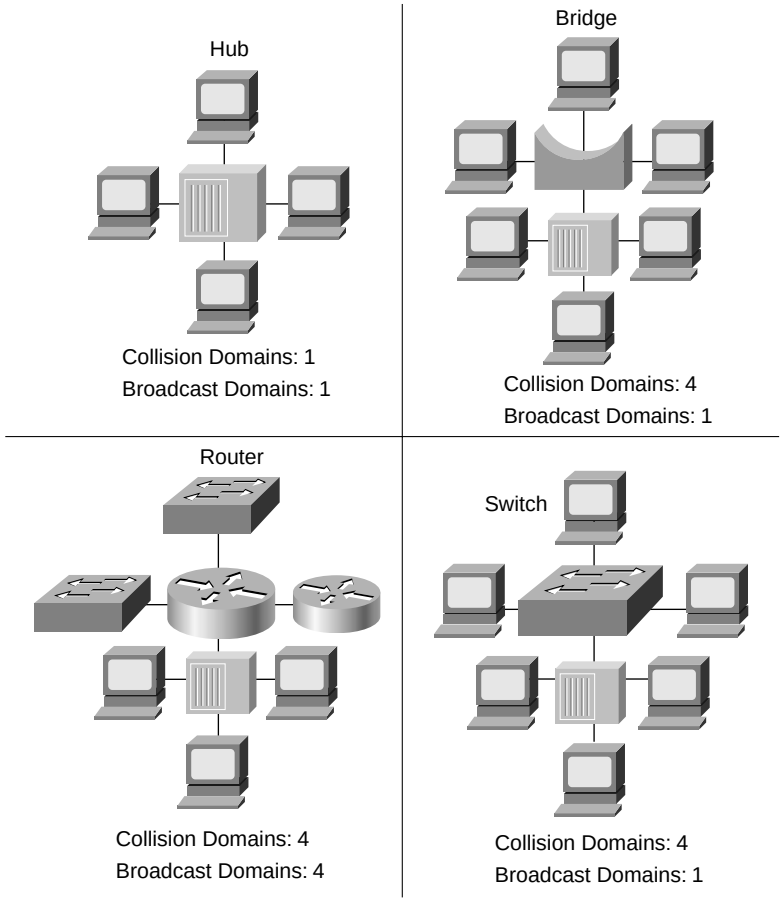

Selecting Cisco Products

Earlier in this chapter, we discussed the hierarchical model used to design and implement networks. Figure 1-26 reviews the structure of this model, shown earlier in Figure 1-3. Given a particular function of networking and what we have discussed about the service performed at each layer, you should be able to match Cisco products to your internetworking needs.

Figure 1-26 *The Three-Layer Hierarchical Network Model*

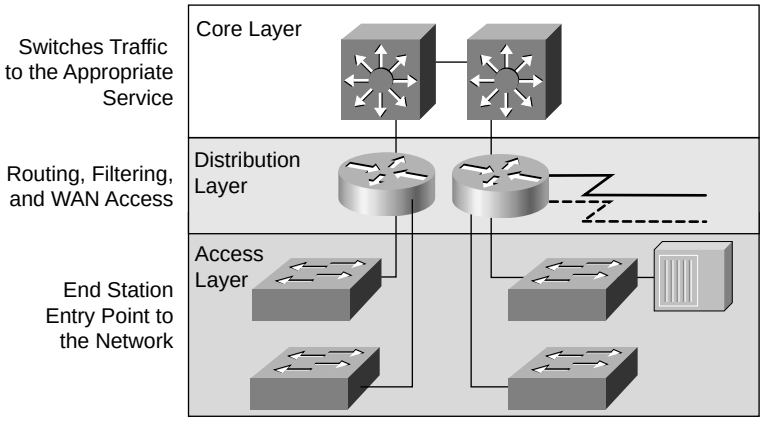

The following list summarizes the factors for selecting networking devices:

- Device provides desired functionality and features
- Device has required capacity and performance
- Device is easy to install and offers centralized management
- Device provides network resiliency
- Device provides investment protection in existing infrastructure
- Device provides migration path for change and growth

The most important task is to understand the needs and then identify the device functions and features that meet those needs. In order to accomplish this, obtain information about where in the internetworking hierarchy the device needs to operate, and then consider factors such as ease of installation, capacity requirements, and so forth.

Other factors, such as remote access, also play a role in product selection. When supporting remote access requirements, you must first determine the kind of WAN services that meet your needs. Then, you will be able to select the appropriate device.

The type and number of required WAN connections will significantly affect your choice of devices. The most important factor in choosing WAN services is the availability of the service. It is also important to know what your bandwidth requirements are and how much the service will cost. Figure 1-27 shows a graph relating cost to usage for some common WAN services. As you can see, depending on the usage, it might be more cost-effective to get a service that provides a fixed rate.

Figure 1-27 *WAN Cost Versus Usage*

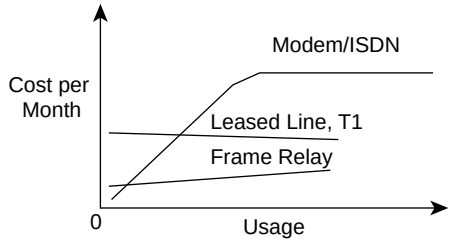

It is also important to choose a service that can be supported by your product.

When determining WAN service bandwidth requirements, you must look at the type of traffic that needs to cross the WAN service. Figure 1-28 gives you an idea of WAN technology as it maps to a given application.

Figure 1-28 *Application Bandwidth Requirements*

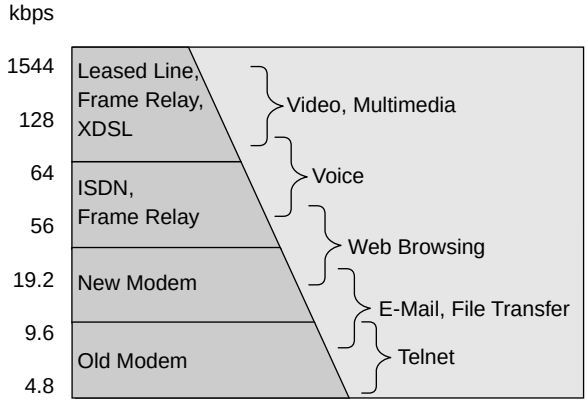

After you have chosen the type of network device you need, you can select a particular product. Cisco Systems offers a large variety of networking products, including hubs, switches, and routers.

Cisco Hub Products

Figure 1-29 shows the selection issues for hubs, along with a sampling of the Cisco hub product line. This figure represents the low-end to high-end line. The cost of these products also increases along this line.

Figure 1-29 *Cisco Hub Product Line*

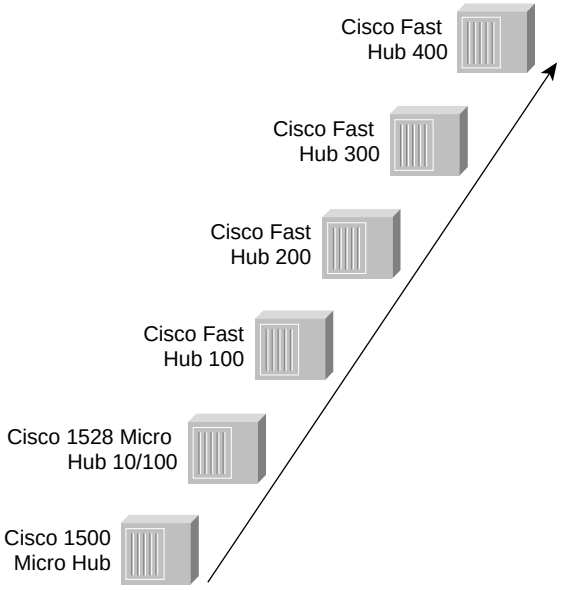

Criteria used in selecting hubs includes the media speed needed, the number of ports needed, ease of installation, and the need for remote management. The Micro Hub series represents the low-end hub with low-speed fixed-port densities. FastHub 100 and 200 represent the mid-level solution, offering higher-speed connectivity and some management. The FastHub 300 and 400 series offer the most flexibility with modular ports and manageability; however, they are 100 Mbps-only devices.

Before implementing hubs, assess which workstations need 10 Mbps and which higher-end stations need 100 Mbps. Lower-end hubs offer only 10 Mbps, whereas mid-range hubs offer both. The mid-range devices provide growth and migration potential.

The scope of consolidated connections refers to the issue of how many hub ports your users will require. Hubs allow for a variety of port densities, and you can stack hubs to get multiples of the hub densities.

Most hubs are simple to plug in and operate. For most hubs, there is no management or console port. If you want to be able to manage the hub, select from the higher-end hub series.

Catalyst Switch Products

Figure 1-30 shows a sampling of the Cisco switch product line. The figure represents the low-end to high-end selection of some of the switch products and shows where in the network these products can be used.

Figure 1-30 *Cisco Switch Product Line*

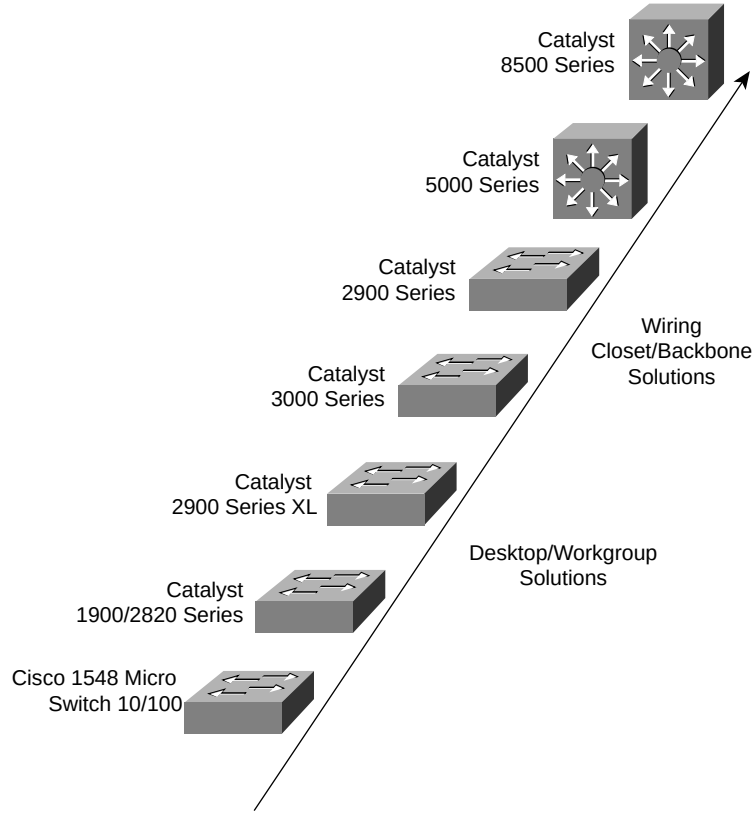

Here are the key selection issues when selecting switch products:

- Media speed requirements
- The need for interswitch communication (trunking)
- The need for broadcast segmentation (VLANs)
- Port density needs
- The need for configuration interface consistency

Because one of the major advantages of switches is the variety of link speeds that are offered, one of the key issues to consider is whether 10 or 100 Mbps access is required.

Other consideration factors for switches are the number of ports, the need for further segmentation using VLANs, and different media and topology connections and enterprise functionality, such as interswitched links for trunking. Many of these functions are discussed later in this book. Finally, you might want all the network devices to have a consistent user configuration interface. Cisco switch products have a variety of user interfaces, including command line, menu, and web. These interfaces could play a role in product selection.

Cisco Router Products

Figure 1-31 shows a sampling of the Cisco router product line. The figure represents the low-end to high-end selection of some of the router products and shows where in the network these products may be used.

Here are the key selection issues when selecting router products:

- Scale of the routing features needed
- Port density/variety requirements
- Capacity and performance
- Common user interface

Figure 1-31 *Cisco Router Product Line*

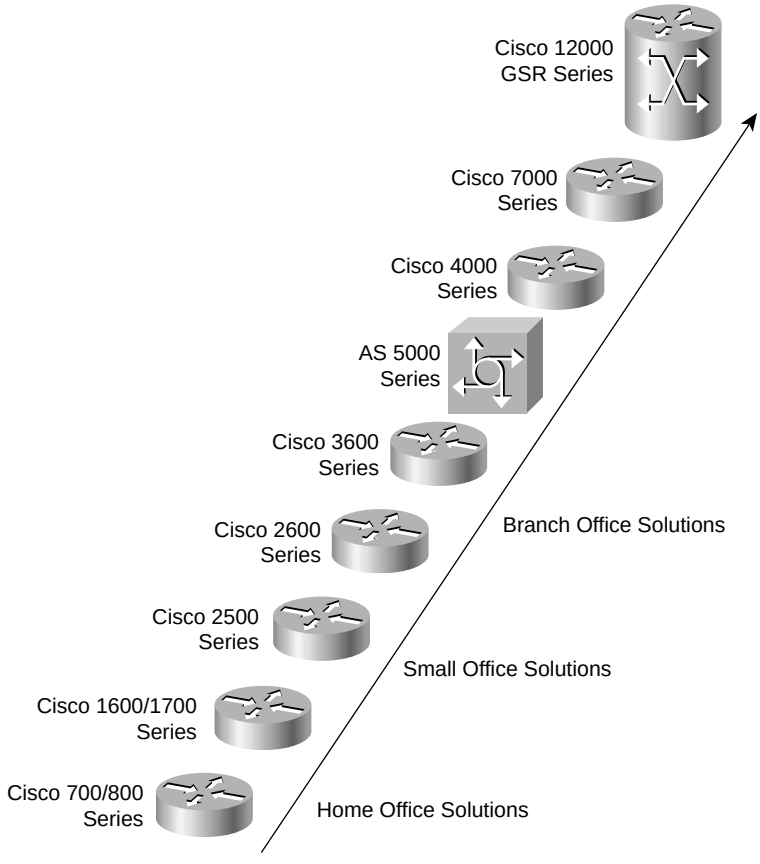

A key criterion in router selection is knowing what router service features are needed. Different routers in the Cisco product line incorporate different feature sets. You will learn about many advanced router features later in this book.

Port densities and interface speeds generally increase as you move to the upper end of the various Cisco router families. For example, the 12000 series is the first in a product class of gigabit switch routers (GSRs). The 12000 GSR initially supports an IP backbone link at OC-12 (622 Mbps) and can scale to handle links at OC-48 (2.4 Gbps). In contrast, the 800

series router is designed to handle 10 Mbps Ethernet connections to the SOHO network and 128 kbps ISDN services to the Internet or corporate office.

If your network requires WAN links, the router selection issues involve which router provides the necessary type and number of links in a cost-effective manner. A typical production network will have several LAN switches interconnected to the WAN by a router.

Please note that the products listed in these sections reflect a snapshot of Cisco's offerings. Cisco's product lines are continuously evolving in response to customer needs and other technology migration issues. For the current Cisco offerings, consult Cisco Connection Online (www.cisco.com) or your dealer/distributor.

NOTE Cisco offers a product selection tool at the web site http://www.cisco.com/pcgi-bin/front.x/corona/prodtool/select.pl. This tool is categorized into three groups—hubs, routers, and switches. It is an interactive JavaScript application used to help you select Cisco products.

Summary

In this chapter, we introduced some basic concepts of internetworking. These concepts include the ability to describe (using the OSI reference model) the process in which data is transferred from an application across the network. You learned the roles of each network device that will be discussed in this book and saw how each fits into the hierarchy of network design. You learned at which layer of the OSI model each of these devices functions. Finally, you learned how to use this information to select products based on the needs of your network.

Review Questions

1. Which three functions are defined by the Cisco hierarchical model?
2. What is one advantage of the OSI reference model?
3. Describe the data encapsulation process.
4. Define a collision domain, and give an example of a device that combines all devices in a single collision domain.
5. Define a broadcast domain, and give an example of a device that separates each segment into different broadcast domains and provides connectivity between the segments.
6. At which layer of the OSI model does a bridge or switch operate?

7. How many broadcast domains are associated with a bridge or switch (assuming no VLANs)?
8. Which OSI layer defines an address that consists of a network portion and a node portion?
9. Which OSI layer defines a flat address space?
10. Which process establishes a connection between two end stations using a reliable TCP/IP transport layer protocol?