9

CHAPTER

Configuring and Managing the Campus Network

CHAPTER OUTLINE

OBJECTIVES

- Understand the purpose of the three layers of a campus network design
- Understand the auxiliary services needed to operate a network, such as DHCP and DNS
- Understand the process of requesting an IP address using DHCP
- Understand how to use SNMP tools for network management
- Investigate how to use network data packet statistics to monitor network performance

KEY TERMS

core
distribution layer
access layer
load balancing
per-destination load balancing
per-pack load balancing
BOOTP
lease time
unicast
MT Discover
MT Offer
MT Request
MT ACK
DNS
forward domain name service
reverse domain name service
TLD
country domain
root servers
NS record
reverse DNS
SNMP
Management Information Base (MIB)
Power over Ether (PoE)
PD
PSE
endpoint PSE
midspan (mid-point) PSE
Resistive Power Discovery
PoE Plus
VLAN (Virtual LAN)
port-based VLAN
tag-based VLAN
protocol-based VLAN
static VLAN
dynamic VLAN
configure terminal (conf t)
Switch(config)#
Switch(config-line)#
Spanning-Tree Protocol
Bridge Protocol Data Unit (BPDU
Configuration BPDU
Topology Change Notification (TCN)
Topology Change Notification Acknowledgement (TCA)

9-1 INTRODUCTION

The objective of this chapter is to examine the computer networking issues that arise when planning a campus network. The term *campus network* applies to any network that has multiple LANs interconnected. The LANs are typically in multiple buildings that are close to each other and are interconnected with switches and routers.

The previous chapters introduced the fundamental issues of computer networks. These included techniques for configuring the LAN, analyzing TCP/IP data traffic, router configuration, configuring the wide area network connection, and selecting and configuring the routing protocols. This chapter looks at the planning and design of a simple campus network, including network design, IP assignment and DHCP, domain name service (DNS), network management, switch and VLAN configuration, and analyzing a campus network's data traffic.

The basics of configuring the three layers of a campus LAN (core, distribution, and access) are first examined in section 9-2. This section also addresses the important issues of data flow and selecting the network media. IP allocation and configuring DHCP service are examined in section 9-3. This includes a step-by-step description of the process of how a computer obtains an IP address manually or via DHCP in a network. Section 9-4 examines the issues of configuring DNS service for a campus network. The concepts of the root servers, the top level domains, and the subdomains are examined. The next section (9-5) addresses network management. An overview of configuring a Cisco router for SNMP operation is first presented. This section includes an example of using SNMP management software to collect router information and data statistics. This section also provides an overview of Power over Ethernet (PoE). Section 9-6 examines configuring a VLAN for use in a campus network. Basic switch commands in addition to configuring a static VLAN are presented. This section also discusses the Spanning-Tree Protocol (STP). The chapter concludes with an example of using data collected with an SNMP management program to monitor network data traffic.

9-2 DESIGNING THE CAMPUS NETWORK

Most campus networks follow a design that has core, distribution, and access layers. These layers (Figure 9-1) can be spread out into more layers or compacted into fewer depending on the size of these networks. This 3-layer network structure is incorporated in campus networks to improve data handling and routing within the network. The issues of data flow and network media are examined here.

Core Layer

Core
The backbone of the network

The network core usually contains high-end layer 3 switches or routers. The **core** is the heart or the backbone of the network. The major portion of a network's data traffic passes through the core. The core must be able to quickly forward data to other parts of the network. Data congestion should be avoided at the core if possible. This means that unnecessary route policies should be avoided. An example of a route policy is *traffic filtering,* which limits what traffic can pass from one part of a network to another. Keep in mind that it takes time for a router to examine each data packet, and unnecessary route policies can slow down the network's data traffic.

It was mentioned that high-end routers and layer 3 switches are typically selected for use in the core. Of the two, the layer 3 switch is probably the best choice. A layer 3 switch is essentially a router that uses electronic hardware instead of software to make routing decisions. The advantage of the layer 3 switch is the speed at which it can establish a network connection.

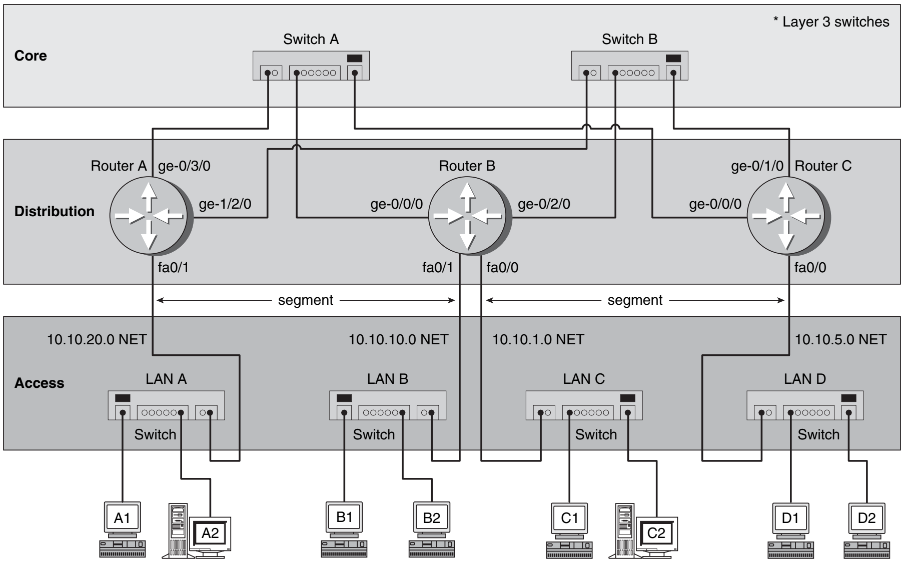

FIGURE 9-1 The core, distribution, and access layers of a campus network.

Another alternative for networking hardware in the core is a layer 2 switch. The layer 2 switch does not make any routing decisions and can quickly make network connection decisions based on the network hardware connected to its ports. The advantage to using the layer 2 switch in the core is cost. The disadvantage is that the layer 2 switch does not route data packets. High-speed layer 2 switches are more affordable than high-speed routers and layer 3 switches.

An important design issue in a campus network and the core is redundancy. *Redundancy* provides for a backup route or network connection in case of a link failure. The core hardware is typically interconnected, as shown in Figure 9-1, to all distribution network hardware. The objective is to ensure that data traffic continues for the whole network even if a core networking device or link fails.

Each layer beyond the core breaks the network into smaller networks with the final result being a group of networks that are capable of handling the amount of traffic generated. The design should thus incorporate some level of redundancy.

Distribution Layer

The **distribution layer** in the network is the point where the individual LANs connect to the campus network routers or layer 3 switches. Routing and filtering policies are more easily implemented at the distribution layer without having a negative impact on the performance of the network data traffic. Also, the speed of the network data connections at the distribution layer are typically slower than at the core. For

Distribution Layer
Point where the individual LANs connect together

example, connection speeds at the core should be the highest possible, such as 1 or 10 gigabits, where the data speed connections at the distribution layer could be 100 Mbps or 1 gigabit. Figure 9-1 shows the connections to the access and core layers via the router's Ethernet interfaces.

Access Layer

Access Layer
Where the networking devices in a LAN connect together

The **access layer** is where the networking devices in a LAN connect together. The network hardware used here is typically a layer 2 switch. Hubs can be used but are not recommended in any networks with significant amounts of data traffic. Remember, a switch is a better choice because it forwards data packets directly to destination hosts connected to its ports. Network data traffic is not forwarded to all hosts in the network. The exception to this is a broadcast that is sent to all hosts connected to the switch. Refer back to Chapter 4 for a review and a comparison of switch and hub operations.

Data Flow

An important networking issue is how data traffic flows in the core, distribution, and access layers of a campus LAN. In reference to Figure 9-1, if computer A1 in LAN A sends data to computer D1 in LAN D, the data is first sent through the switch in LAN A and then to RouterA in the distribution layer. RouterA then forwards the data to switch A or switch B. Switch A or switch B will then forward the data to RouterC. The data packet is then sent to the destination host in LAN D.

The following are some questions often asked when setting up a network that implements the core, distribution, and access layers.

- *In what layer are the campus network servers (Web, email, DHCP, DNS, etc.) located?* This varies for all campus networks and there is not a definitive answer. However, most campus network servers are located in the access layer.
- *Why not connect directly from RouterA to RouterC at the distribution layer?* There are network stability issues when routing large amounts of network data traffic if the networks are fully or even partially meshed together. This means that connecting routers together in the distribution layer should be avoided.
- *Where is the campus backbone located in the layers of a campus network?* The backbone of a campus network carries the bulk of the routed data traffic. Based on this, the backbone of the campus network connects the distribution and the core layer networking devices.

Selecting the Media

The choices for the media used to interconnect networks in a campus network are based on several criteria:

- Desired data speed
- Distance for connections
- Budget

The desired data speed for the network connection is probably the first consideration given when selecting the network media. Twisted-pair cable works well at 100 Mbps and 1 Gbps and is specified to support data speeds of 10-gigabit data traffic over twisted-pair cable. Fiber optic cable supports LAN data rates up to 10 Gbps or higher. Wireless networks support data rates up to 200+ Mbps.

The distance consideration limits the choice of media. CAT 6/5e or better have a distance limitation of 100 meters. Fiber optic cable can be run for many kilometers, depending on the electronics and optical devices used. Wireless LAN connections can also be used to interconnect networks a few kilometers apart.

The available budget is always the final deciding factor when planning the design for a campus LAN. If the budget allows, then fiber optic cable is probably the best overall choice especially in the high-speed backbone of the campus network. The cost of fiber is continually dropping, making it more competitive with lower-cost network media such as twisted-pair cable. Also fiber cable will always be able to carry a greater amount of data traffic and can easily grow with the bandwidth requirements of a network.

Twisted-pair cable is a popular choice for connecting computers in a wired LAN. The twisted-pair technologies support bandwidths suitable for most LANs, and the performance capabilities of twisted-pair cable are always improving.

Wireless LANs are being used to connect networking devices together in LANs where a wired connection is not feasible. For example, a wireless LAN could be used to connect two LANs in a building together. This is a cost-effective choice if there is not a cable duct to run the cable to interconnect the LANs or if the cost of running the cable is too high. Also wireless connections are playing an important role with mobile users within a LAN. The mobile user can make a network connection without having to use a physical connection or jack. For example, a wireless LAN could be used to enable network users to connect their mobile computers to the campus network. This topic is examined in greater detail in Chapter 11.

Load Balancing

Load balancing is the concept of distributing the network data traffic over multiple interfaces so that one interface connection is not overloaded. It is also used as a tool to control the direction of the data flow when multiple routes are available. An example of load balancing is provided in Figure 9-1. In this case, data from LAN A going to LAN B will first pass through Router A at the distribution layer. The data will next travel to core Switches A or B. The data can all be sent to Switch A or Switch B or can be split over both switches.

Load Balancing
Concept of distributing the network data traffic over multiple interfaces so that one interface connection is not overloaded—also used as a tool to control the direction of the data flow

The selection of the data path leaving RouterA is determined by the "cost" of the data path. Remember, the "cost" of the route is typically based on the type of routing protocol and the "speed" of the connection. For example in Figure 9-1, RouterA is running OSPF for all interfaces, and the cost associated with RouterA's gigabit interface is lower than the cost of a RouterA's FastEthernet interface. The cost of a route can be programmed into the router by issuing the following command from the config-router prompt, ***ip ospf cost <number>*** where number is the cost of the network route ranging from 1 (lowest) to 65535 (highest). The lower the cost, the more preferred the route.

In this case, the cost of the route is being set by the network administrator who is using the cost metric to control the route of the data flow. For example, the command for setting the cost of the RouterA ge-0/3/0 interface to 10 is a follows. (*Note:* This example assumes that RouterA is running ospf off the ge-0/3/0 interface.)

```
RouterA<config-router># ip ospf cost 10
```

RouterA is also showing a link to Switch B off the ge-1/2/0 gigabit interface. If the primary data traffic at the core should travel through Switch A, then the path to Switch B should have a higher cost. For example, setting the cost off RouterA's ge-1/2/0 interface to 20 will make the data path to Switch A ge-1/2/0 interface have a higher cost and therefore establish the ge-0/3/0 interface (cost = 10) as the primary route.

Per-destination Load Balancing
Means that the data packets coming from the router are distributed based on the destination address

Per-packet Load Balancing
Load balance is guaranteed for all interfaces, but there is no guarantee that the packets will arrive at the destination in the proper order

What if the cost of both RouterA's gigabit interfaces is set to 10? In this case, the routers will load balance automatically, and the data traffic will be split equally over both interfaces. Load balancing off the router can be set to work per-destination or per-packet. **Per-destination load balancing** means that the data packets coming from the router are distributed based on the destination address. This means that all data packets with the same destination are sent over the same interface. This technique preserves the order of the packets but it does not guarantee equal load balancing. Another type of load balancing is per-packet. In **per-packet load balancing**, load balance is guaranteed for all interfaces, but there is no guarantee that the packets will arrive at the destination in the proper order considering the data packet can take different routes.

9-3 IP ASSIGNMENT AND DHCP

IP assignment is a process where subnets are created for each subgroup or department. The IP address assignment is typically tracked by the network operations center (NOC). The IP addresses are kept in a central log file so that NOC can troubleshoot network problems. For example, a machine could be causing network problems possibly due to hacked or corrupted software. NOC needs to be able to track down the network problem(s). The NOC database will have the MAC address, the IP address, and the name of the person who uses the computer.

IP addresses are assigned by NOC based on where the subnet for the computer is located. The subnet could be in a building, a floor of the building, a department, and so on. The subnets are created by the network administrators based on the expected number of users (hosts) in a subnet (refer back to Chapter 5). For example, the 192.168.12.0 network shown in Figure 9-2 has been partitioned into four subnets. The network addresses for each of the subnets are provided in Table 9-1. Any computer in subnet B is assigned one of the 62 IP addresses from the range 192.168.12.65 to 192.168.12.126. Remember, the first IP address in the subnet is reserved for the network address, and the last is reserved for the broadcast address.

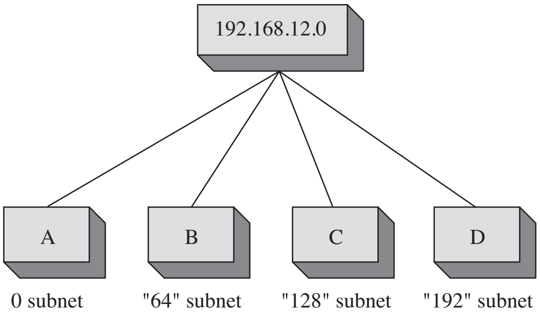

FIGURE 9-2 IP assignment of computers in a network's subnet.

TABLE 9-1 Subnet Addresses for the Subnets Shown in Figure 9-2

Subnet	Network Address	Broadcast Address
A	192.168.12.0	192.168.12.63
B	192.168.12.64	192.168.12.127
C	192.168.12.128	192.168.12.191
D	192.168.12.192	192.168.12.255

IP assignment is done either manually or dynamically. In the manual process for IP assignment within a campus network, IP addresses are entered manually for each computer in the network. This is a tedious process that requires that a file be edited with all of the necessary information about the computer, including the MAC address, the IP address, and the user (or owner) of the computer. The owner is typically the department or organizational unit within the campus LAN. This process can be automated to some extent using a program called **BOOTP** for IP assignment. BOOTP stands for *Bootstrap Protocol*, and it enables computers to discover their own IP addresses. When a client requests an IP address, it is assigned to the Ethernet address (MAC address) based on the BOOTP record. In this case, the IP and MAC addresses have a one-to-one relationship.

BOOTP
Bootstrap Protocol

DHCP (Dynamic Host Configuration Protocol) simplifies the steps for IP assignment. DHCP's function is to assign IP addresses from a pool to requesting clients. DHCP is a superset of BOOTP, and runs on the same port number. DHCP requests an IP address from the DHCP server. The DHCP server retrieves an available IP address from a pool dedicated to the subnet of the requesting client. The IP address is passed to the client and the server specifies a length of time that the client can hold the address. This is called the **lease time**. This feature keeps an unused computer from unnecessarily tying up an IP address.

Lease Time
The amount of time that a client can hold an IP address

The process of requesting an IP address with DHCP is as follows: The client boots up and sends out a DHCP request. This is a broadcast, meaning that the message is sent to all computers in the LAN. A DHCP server listening on the LAN will take the packet, retrieve an available IP address from the address pool, and send the address to the client. The client applies the IP address to the computer and is then ready to make network connections. An example is provided in Figure 9-3.

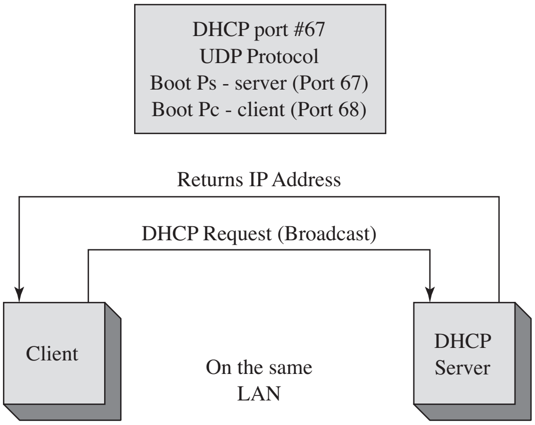

FIGURE 9-3 An example of a DHCP server and client in the same LAN.

Unicast
The packet has a fixed destination

What if a DHCP server is on the other side of the router (for example, not in the same LAN)? Remember, routers don't pass broadcast addresses, so the DHCP broadcast is not forwarded. This situation requires that a DHCP relay be used, as shown in Figure 9-4. The DHCP relay sits on the same LAN as the client. It listens for DHCP requests and then takes the broadcast packet and issues a **unicast** packet to the network DHCP server. *Unicast* means that the packet is issued a fixed destination and therefore is no longer a broadcast packet. The DHCP relay puts its LAN address in the DHCP field so the DHCP server knows the subnet the request is coming from and can properly assign an IP address. The DHCP server retrieves an available IP address for the subnet and sends the address to the DHCP relay, which forwards it to the client.

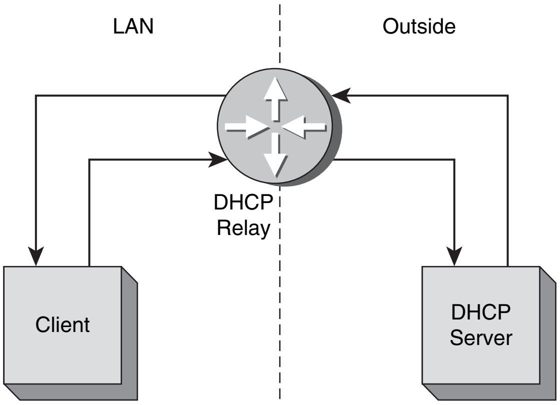

FIGURE 9-4 An example requiring the use of a DHCP relay.

Cisco routers have a DHCP relay built into their operating system. The router command to enable the DHCP relay is ***Router(config-if)# ip helper [ip address of the DHCP server]***. Notice that this command is issued from the interface that connects to the LAN. In fact, the IP address for the interface is typically the gateway address for the LAN.

DHCP is a UDP protocol and uses port number 68 for the BOOTP-client and port 67 for the BOOTP-server. (BOOTP and DHCP use the same port numbers.) The BOOTP-client is the user requesting the DHCP service. The BOOTP-server is the DHCP server. The following discussion describes how these services are used in a DHCP request. The DHCP proxy on the router listens for the packets that are going to DHCP or BOOTP port numbers.

The DHCP Data Packets

The following is a discussion on the TCP packets transferred during a DHCP request. The network setup is the same as shown in Figure 9-4. The data traffic shown in this example will contain only the data packets seen by the client computer. The Finisar Surveyor Demo software was used to capture the data packets. A portion of the captured data packets is provided in Figure 9-5. Packet 10 is a DHCP request with a message type discover (**MT Discover**). This is also called the DHCP Discover packet. The destination for the packet is a broadcast. The message source has a MAC address of Dell 09B956, and the IP address is 0.0.0.0. The IP address is shown in the middle panel, and the *0.0.0.0* indicates that an IP address has not been assigned to the computer. The source and destination ports are shown in the third panel in Figure 9-5. The source port is 68, which is for the Bootstrap Protocol Client (the computer requesting the IP address). The destination port is 67, the Bootstrap Protocol Server (the DHCP server).

MT Discover
Message type discover, a DHCP Discover packet

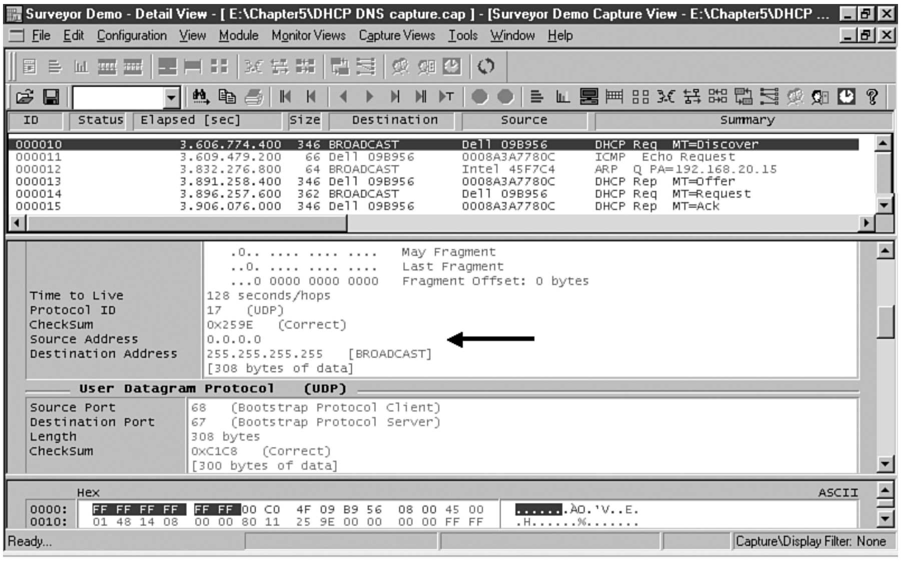
Surveyor Demo - Detail View - [E:\Chapter5\DHCP DNS capture.cap] - [Surveyor Demo Capture View - E:\Chapter5\DHCP ...
File Edit Configuration View Module Monitor Views Capture Views Tools Window Help

```
ID      Status  Elapsed [sec]     Size  Destination      Source           Summary
000010          3.606.774.400     346   BROADCAST        Dell 09B956      DHCP Req  MT=Discover
000011          3.609.479.200      66   Dell 09B956      0008A3A7780C     ICMP   Echo Request
000012          3.832.276.800      64   BROADCAST        Intel 45F7C4     ARP  Q PA=192.168.20.15
000013          3.891.258.400     346   Dell 09B956      0008A3A7780C     DHCP Rep  MT=Offer
000014          3.896.257.600     362   BROADCAST        Dell 09B956      DHCP Req  MT=Request
000015          3.906.076.000     346   Dell 09B956      0008A3A7780C     DHCP Rep  MT=Ack

                         .0.. .... .... ....   May Fragment
                         ..0. .... .... ....   Last Fragment
                         ...0 0000 0000 0000   Fragment Offset: 0 bytes
Time to Live           128 seconds/hops
Protocol ID            17    (UDP)
CheckSum               0x259E   (Correct)
Source Address         0.0.0.0
Destination Address    255.255.255.255   [BROADCAST]
                       [308 bytes of data]
        User Datagram Protocol    (UDP)
Source Port            68    (Bootstrap Protocol Client)
Destination Port       67    (Bootstrap Protocol Server)
Length                 308 bytes
CheckSum               0xC1C8   (Correct)
                       [300 bytes of data]

        Hex                                                     ASCII
0000:   FF FF FF FF  FF FF 00 C0  4F 09 B9 56  08 00 45 00     .......ÀO.'V..E.
0010:   01 48 14 08  00 00 80 11  25 9E 00 00  00 00 FF FF     .H......%.......
Ready...                                      Capture\Display Filter: None
```

FIGURE 9-5 The captured DHCP packets.

Packet 13 is a reply from the DHCP server, an offer of the IP address to the client. This is called the DHCP Offer packet (**MT Offer**). This packet contains the domain name, the domain name server, the default gateway, other network information the client may need to connect to the network. Packet 14 has a message type of **MT Request**. This packet is sent from the client back to the server that has been selected to provide the DHCP service. (*Note*: It is possible for a campus LAN to have more than one DHCP server answering the DHCP request.) The packet is sent through the DHCP relay to the DHCP server. This means that the client is accepting the IP address offer. Packet 15 is a message type of ACK (**MT ACK**). The DHCP server is acknowledging the client's acceptance of the IP address from the DHCP server. The client computer now has an IP address assigned to it.

MT Offer
Message type offer, a DHCP Offer packet

MT Request
Message type request, a DHCP Request packet

MT ACK
Message type acknowledgement, a DHCP ACK packet

9-4 NETWORK SERVICES—DNS

This section examines the DNS services that are typically available in a campus network. **DNS** is the domain name service. DNS translates a human readable name to an IP address or an IP address to a domain name. The translation of a name to an IP address is called **forward domain name service**, and translation of an IP address to a domain name is called **reverse domain name service**.

DNS
Domain name service

Forward Domain Name Service
Translation of a name to an IP address

Reverse Domain Name Service
Translation of an IP address to a name

TLD
Top level domain

Country Domain
Usually two letters, such as United States (.us) or Canada (.ca), that define the location of the domain server for that country

The domain name service is a tree hierarchy. It starts with the top level domains and then extends to subdomains. Examples of top level domains (**TLD**) are as follows:

.com .net .org .edu .mil .gov .us .ca .info .biz .tv

Country domains are usually defined by two letters, such as .us (United States) and .ca (Canada). The primary domain server for that domain has to exist in the same country; for example, the .us primary domain server is located in the United States. Figure 9-6 shows the top level domains and their relationship to the subdomains and root servers.

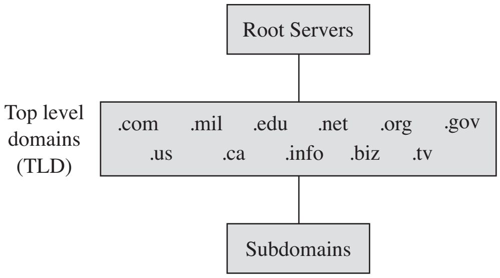

FIGURE 9-6 The domain name service tree hierarchy.

Root Servers
A group of servers that exist using well-known IP addresses that have been programmed into DNS servers

The **root servers** use well-known IP addresses that have been programmed into DNS servers. When the DNS is installed on a server, the root server's IP addresses are automatically configured in the DNS. The campus DNS will query the root servers to try to find name servers of known domains.

For example, if network-A wants to know the IP address for the www server at network-B.edu, DNS queries one of the root servers and it returns the IP address for the .edu domain. Then the network-A DNS queries the .edu domain for the IP address of the network-B DNS. The network-A DNS then queries the network-B DNS to obtain the IP address for www.network-B.edu. There are many steps for obtaining the IP address via DNS. However, the DNS server keeps a cache of recent queries so this multiple-step process of obtaining an IP address does not have to be repeated unnecessarily. The *www* entry, called the *A record,* is the name for an IP address. The A record is used by a DNS server at the parent company for network-B to convert the name www.network-B.edu to an IP address.

NS Record
A record that points to a name server

The top level domain for network-B.edu is .edu, and the subdomain is network-B.edu. At the top level of the .edu domain, and all domains, are the root servers, as shown in Figure 9-6. The .edu domain has an **NS record**, basically a record that points to a name server. They will have an NS record for network-A that points to the IP address of the network-A domain name server and the secondary DNS server's IP address and its DNS names.

The root servers have information only about the next level in the tree. Root servers know only about the top level domains (for example, .com, .gov, .mil, etc.). They will not know anything about www.network-B.edu. They only know the .edu domain server's IP address.

Campus DNS

The first step for providing DNS for a campus network is to obtain a domain name. This requires that the user seeking the domain name go to www.internic.net. Internic has a list of name registrars where a domain can be purchased. Select a company that registers domain names. When you get on the registrar's website you will be able to input a domain name. The registrar will check to see if the domain name is available. If the domain name is available, you will be prompted to complete the application for the domain name and put in the DNS servers that are to be used to host the domain. The DNS servers will be assigned an IP address and names. When the network's DNS servers are placed online, the root servers will point to the network's DNS servers.

Administering the Local DNS Server—A Campus Network Example The primary records are the A records of a campus network. These contain the host name and IP addresses for the computers. For example, network-B.edu has an assigned IP address of 172.16.12.1. When a host pings www.network-B.edu, the host computer first checks its DNS cache; assuming the DNS cache is empty, the host then sends a DNS request to the campus DNS server. Typically the host will know the IP addresses of the primary and secondary DNS server through either static input or dynamic assignment. The request is sent to the primary DNS server requesting the IP address for www.network-B.edu. The primary DNS server is the authority for network-B.edu and knows the IP address of the hosts in the network. The primary DNS server returns the IP address of www.network-B.edu, and then the ICMP process associated with a ping is started.

One might ask, "How does a PC in the campus network become part of the campus domain?" Specifically, how is an A record entered into the campus domain? Recall that the A record provides a host to IP address translation. Adding the PC to the campus domain is done either manually or dynamically.

The Steps for Manually Adding a Client to the Campus Network The steps for manually updating the DNS A records are graphically shown in Figure 9-7 and are listed as follows: A client PC updates the A record when an IP address is requested for a computer. The user obtains the PC name and the PC's MAC address. This information is sent to the network operation center (NOC). NOC issues an IP address to the client, updates the NOC database of clients on the network, and enters a new A record into the primary DNS. The entry is only made on the primary DNS. The entry will be later duplicated on the secondary DNS.

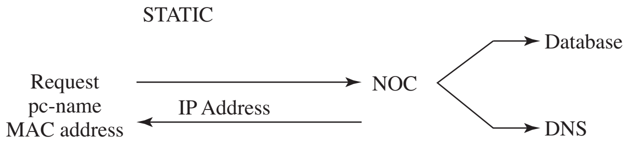

FIGURE 9-7 Manually updating the A record.

The Steps for Dynamically Adding a Client to the Campus Network A new A record can be entered dynamically when the client computer obtains an IP address through DHCP registration. This is graphically depicted in Figure 9-8. The DHCP server will issue an IP address to the client and at the same time send an updated A record to the network's primary DNS. Once again, the client name and the IP and MAC addresses are stored in the A record.

Why obtain the MAC address when entering the information into DNS? This record is used to keep track of all the machines operating on the network. The MAC address is a unique identifier for each machine. The MAC address is also used by BOOTP, which is a predecessor to DHCP. This is where a MAC address is specifically assigned to one IP address in the network.

Reverse DNS
Returns a hostname for an IP address

Reverse DNS returns a hostname for an IP address. This is used for security purposes to verify that your domain is allowed to connect to a service. For example, pc-salsa1-1 (10.10.20.1) connects to an FTP server that only allows machines in the salsa domain to make the connection. When the connection is made, the FTP server knows only the IP address of the machine making the connection (10.10.20.1). The server will use the IP address to request the name assigned to that IP. A connection is made to the salsa domain server, and the salsa DNS server returns pc-salsa1-1 as the machine assigned to 10.10.20.1. The FTP server recognizes this is a salsa domain machine and authorizes the connection.

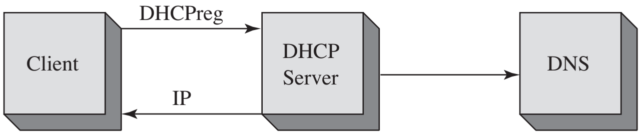

FIGURE 9-8 Dynamic updating of the A record using DHCP.

9-5 NETWORK MANAGEMENT

A campus network of moderate size has a tremendous number of data packets entering and leaving. The number of routers, switches, hubs, servers, and host computers can become staggering. Proper network management requires that all network resources be managed. This requires that proper management tools be in place.

SNMP
Simple Network Management Protocol

Management Information Base (MIB)
A collection of standard objects that are used to obtain configuration parameters and performance data on a networking device

A fundamental network management tool is **SNMP**, the Simple Network Management Protocol. SNMP, developed in 1988, is widely supported in most modern network hardware. SNMP is a connectionless protocol using the UDP (User Datagram Protocol) for the transmission of data to and from UDP port 161.

SNMP uses a **management information base (MIB)**, which is a collection of standard objects that are used to obtain configuration parameters and performance data on a networking device such as a router. For example, the MIB (ifDescr) returns a description of the router's interfaces. An example is shown in Figure 9-9. An SNMP software tool was used to collect the interface description information. The IP address of the router is 10.10.10.1, and a ***get request ifDescr*** was sent to port 161, the UDP port for SNMP. The descriptions of the interfaces were returned as shown.

Obtaining the SNMP data requires that SNMP be configured on the router. The following discussion demonstrates how to configure SNMP on a Cisco router.

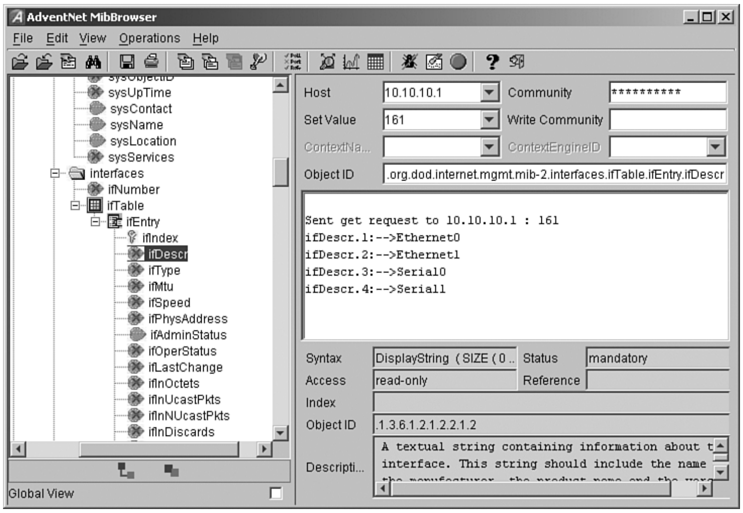

Host: 10.10.10.1
Community: **********
Set Value: 161
Write Community:
ContextNa...:
ContextEngineID:
Object ID: .org.dod.internet.mgmt.mib-2.interfaces.ifTable.ifEntry.ifDescr

```
Sent get request to 10.10.10.1 : 161
ifDescr.1:-->Ethernet0
ifDescr.2:-->Ethernet1
ifDescr.3:-->Serial0
ifDescr.4:-->Serial1
```

Syntax: DisplayString (SIZE (0 ..
Status: mandatory
Access: read-only
Reference:
Index:
Object ID: .1.3.6.1.2.1.2.2.1.2
Descripti...: A textual string containing information about t interface. This string should include the name

FIGURE 9-9 An example of using an SNMP software management tool to obtain descriptions of a router's interfaces using the MIB (ifDescr).

Configuring SNMP

The first step for configuring SNMP on a Cisco router is to enter the router's configuration mode using the ***conf t*** command:

```
 RouterB#conf t
Enter configuration commands, one per line. End with CNTL/Z.
```

From the router's (config)# prompt enter the command ***snmp community [community string] [permissions]***. The community string can be any word. The permissions field is used to establish if the user can read only (ro), write only (wo), or both (rw). The options for configuring SNMP on the router are shown here:

```
RouterB(config)#snmp community ?
WORD SNMP community string
```

The router was connected to the computer running the SNMP management software, as shown in Figure 9-10. The router's configuration mode was entered, and the ***snmp community public ro*** command was issued. The word *public* is used as the community string. The community string is the password used by the SNMP software to access SNMP (port 161) on the router. The ***ro*** sets the permission to read only.

```
RouterB(config)#snmp community public ro
```

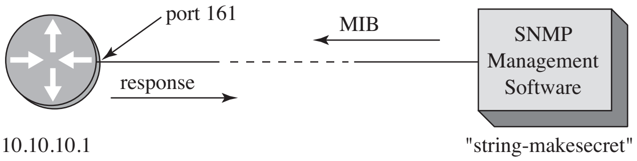

FIGURE 9-10 The setup for connecting the SNMP management software tool to the router.

In the next example, the community string password is set to *makesecret*, and the permission is set to read write (rw). Once again, the router's (config)# mode is entered and the command ***snmp community makesecret rw*** is entered:

```
RouterB(config)#snmp community makesecret rw
```

The configuration for SNMP can be verified using the ***show run*** command from the router's privileged mode prompt. A portion of the configuration file that lists the SNMP configuration for the router is shown here:

```
RouterB#sh run
.
.
snmp-server community makesecret RW
.
.
```

Figure 9-10 shows the setup of the configured router and the computer running the SNMP management software. The SNMP management software issues the MIB to the router at port 161, and the router returns the response. Figure 9-11 shows another example of using SNMP to obtain interface information about a router. The SNMP manager was configured with the host IP address of 10.10.10.1, a set value (port #) of 161 and the 10 character community string of *makesecret* shown as * * * * * * * * * *. The MIB (ifspeed) was sent to the router and a status for each of the interfaces was provided. The data displayed shows the speed settings for the router's interfaces.

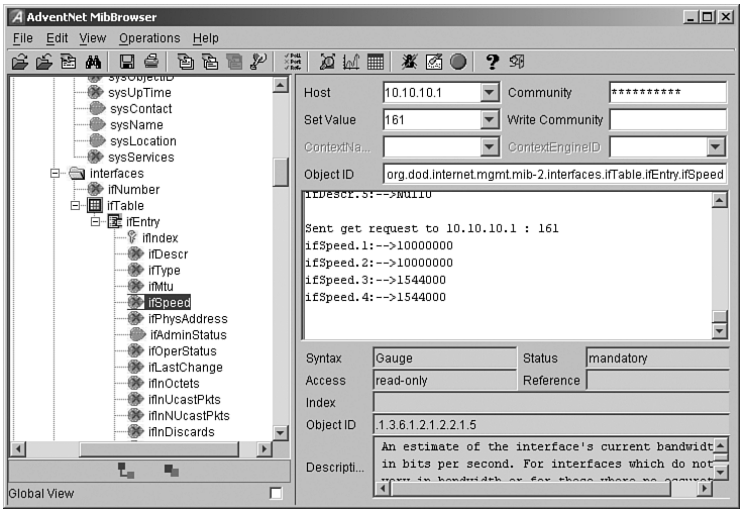

Host 10.10.10.1
Community **********
Set Value 161
Write Community
Object ID org.dod.internet.mgmt.mib-2.interfaces.ifTable.ifEntry.ifSpeed

```
ifDescr.5:-->Null0

Sent get request to 10.10.10.1 : 161
ifSpeed.1:-->10000000
ifSpeed.2:-->10000000
ifSpeed.3:-->1544000
ifSpeed.4:-->1544000
```

Syntax Gauge
Status mandatory
Access read-only
Reference
Index
Object ID .1.3.6.1.2.1.2.2.1.5
Descripti... An estimate of the interface's current bandwidt
in bits per second. For interfaces which do not

FIGURE 9-11 Using an SNMP software management tool to obtain interface speed settings.

Another important application of SNMP is for obtaining traffic data statistics. An example of this is shown in Figure 9-12. The SNMP management program issued the MIB (ifOutOctets), which returns the number of octets of data that have left the router. (The router has a counter that keeps track.) The first result shows ifOutOctets 7002270. The next result display shows that the ifOutOctets returns a value of 7002361.

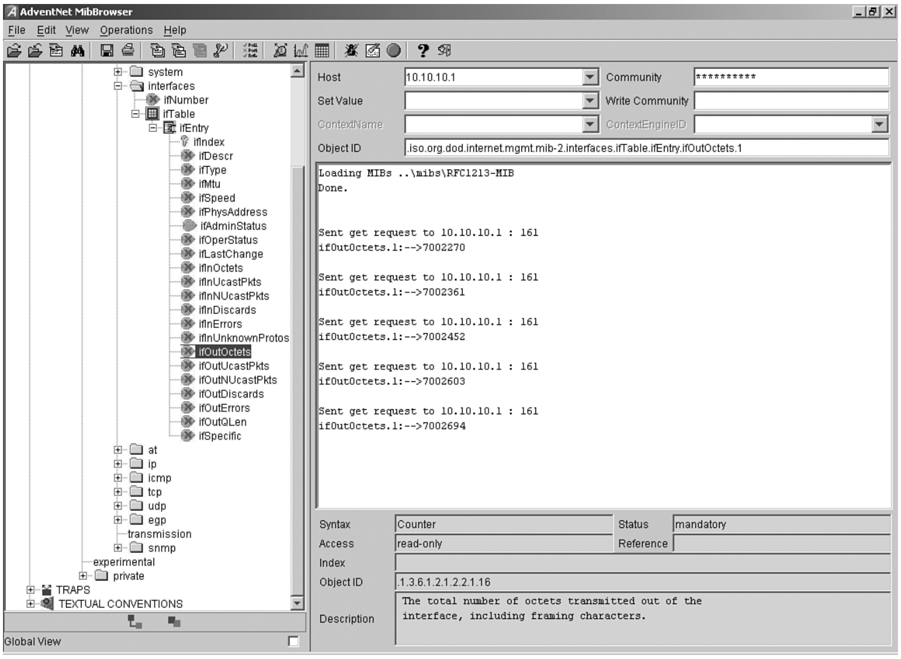

Host 10.10.10.1
Community **********
Set Value
Write Community
ContextName
ContextEngineID
Object ID .iso.org.dod.internet.mgmt.mib-2.interfaces.ifTable.ifEntry.ifOutOctets.1

```
Loading MIBs ..\mibs\RFC1213-MIB
Done.


Sent get request to 10.10.10.1 : 161
ifOutOctets.1:-->7002270

Sent get request to 10.10.10.1 : 161
ifOutOctets.1:-->7002361

Sent get request to 10.10.10.1 : 161
ifOutOctets.1:-->7002452

Sent get request to 10.10.10.1 : 161
ifOutOctets.1:-->7002603

Sent get request to 10.10.10.1 : 161
ifOutOctets.1:-->7002694
```

Syntax Counter
Status mandatory
Access read-only
Reference
Index
Object ID .1.3.6.1.2.1.2.2.1.16
Description The total number of octets transmitted out of the interface, including framing characters.

FIGURE 9-12 An example of using SNMP to collect data traffic statistics.

The SNMP management program collecting the statistics keeps track of the time interval between measurements and the number of octets that have passed. This information can be used to calculate the average traffic flow by hour, day, week, or month, depending on the information needed. An example of collecting traffic route statistics is provided in section 9-7. A final note about the router's counter: the counter does not reset unless the router is rebooted.

Power over Ethernet (PoE)

One of the challenges the network administrator faces as the campus network grows is making sure that electrical power is available for the networking devices (for example, switches, wireless access points, and IP phones). It is not always practical or affordable to run electrical power every place a networking device is needed. This challenge is met with **Power over Ethernet (PoE)**. The Power over Ethernet standard (IEEE 802.3af) was approved in 2003 for networks running 10BASE-T, 100BASE-T, and 1000BASE-T technologies. This provided a standardized technology that can be used to supply power over existing CAT5 or better network cabling (CAT5 or better) to the networking devices.

Power over Ethernet (PoE)
Technology developed to supply power over the network cabling (CAT5 or better)

The benefits of PoE include the following:

- It is not necessary to run external power to all networking devices.
- You can run power and data over one cable.
- Monitoring of power management via SNMP.
- Networking devices can be moved easily.

The power provided by PoE as defined by IEEE 802.3af is as follows:

15.4 watts per port to Ethernet devices
48 volt system

PD
Powered Device

PSE
Power Sourcing Equipment

There are two pieces of networking hardware defined in PoE. These are the **Powered Device (PD)** and the **Power Sourcing Equipment (PSE)**. There are three main functions provided by the PSE:

- Capability of detecting a PD
- Supplying power to the PD
- Power supply monitoring

Endpoint PSE
Example is the source port on an Ethernet switch that connects to the PD

Midspan (mid-point) PSE
Used to provide power to a PD when a powered Ethernet port is not available

There are two types of power sourcing equipment (PSE). These are the **endpoint PSE** and the **midspan (mid-point) PSE**. An example of an endpoint PSE is the source port on an Ethernet switch that connects, via a cable, to a PD. The power to the PD can be delivered in two ways, over the active data pairs (for example, 1–2, 3–6) or via pairs 4–5, 7–8. This is shown on Figure 9-13 (a and b). Both types of power delivery can be used for 10BASE-T, 100 BASE-T, and 1000 BASE-T. The most common way of power delivery is over pairs 1–2/3–6 as shown in Figure 9-13 (a).

A midspan or mid-point PSE is used to provide power to a PD when a powered Ethernet port is not available. This setup requires the use of a power injector and typically uses pairs 4–5/7–8, and this does not support 1000 BASE-T connections.

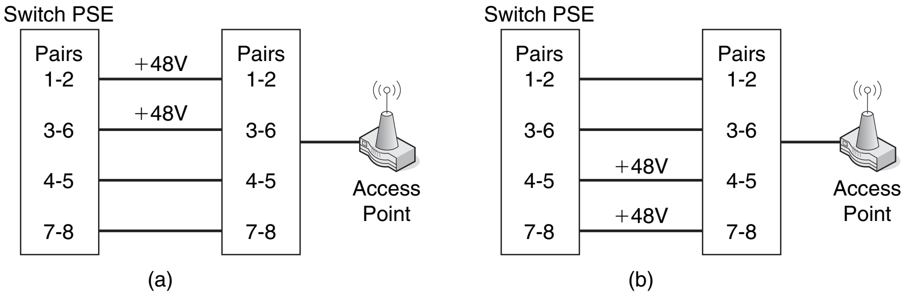

FIGURE 9-13 The two ways to deliver power to the PD: (a) pairs 1–2/3–6 and (b) 4–5/7–8

The PD is the actual device receiving power such as wireless access points and an IP phone. There are four classes of PD devices:

Class 0 (.44 to 12.95 watts)
Class 1 (.44 to 3.84 watts)
Class 2 (3.84 to 6.49 watts)
Class 3 (6.49 to 12.95 watts)

Resistive Power Discovery
Looking for devices that support PoE and have a 25kΩ resistor connected between the transmit and receive pairs

PD devices are "discovered" by the PSE by sending discovery signals on active and inactive Ethernet ports. The discovery process (called **Resistive Power Discovery**) is basically looking for devices that support PoE. Valid PDs will have a 25kΩ resistor connected between the transmit and receive pairs. Before full power is delivered to the PD, two low-voltage "discovery" signals are sent out to verify that a compatible PoE device is attached. The second of the two signals is a slightly higher voltage than the first but neither is large enough to damage an incompatible device. If the PSE detects a compatible PD, then the full 48 Volts is applied to all ports that have compatible PDs connected.

A new version of Power over Ethernet, called **PoE Plus**, based on the IEEE 802.3at standard is now available. PoE Plus provides the following features:

- supports both 802.3af (PoE) and 802.3at (PoE Plus) PDs
- support for midspan PSEs for 10000 BASE-T
- supports a minimum of 30 watts of power for the PD
- support for 10GBASE-T
- will operate with CAT5 and higher cabling

PoE Plus
A new version of PoE based on IEEE 802.3at

There are a limited number of problems with PoE as long as the devices support IEEE 802.3af. In cases where a vendor proprietary PoE equipment is being used, the PSEs and PDs must be compatible. Also the Network Administrator must be aware of how many PDs are connected to the PSE and that the total power requirements for the PDs does not exceed the PSEs limit. For example, a network could have 10 access point and 25 IP phone all requiring a PoE connection. The total number of devices requiring power can exceed the power output for a PSE and possibly damaging the device.

9-6 Switch/VLAN Configuration

The networking switch was introduced in Chapter 4. The basic functions and operation of a managed switch were introduced, as were the various modes of operations. This section examines the function of using a switch in a VLAN within the campus network. The terminology and steps for implementing VLANs will be first presented. The second part examines basic Cisco switch configuration and provides an introduction to the commands needed for configuring the VLAN. The third part of section 9-6 demonstrates the commands needed to set-up a static VLAN. The section concludes with a discussion on the Spanning-Tree Protocol (STP).

Virtual LAN (VLAN)

A switch can be configured as a **VLAN (Virtual LAN)** where a group of host computers and servers are configured as if they are in the same LAN even if they reside across routers in separate LANs. The advantage of using VLANs is the network administrator can group computers and servers in the same VLAN based on the organizational group (e.g. Sales, Engineering) even if they are not on the same physical segment or even the same building.

There are three types of VLANs; **port-based**, **tag-based**, and **protocol-based**. The port-based VLAN is one where the host computers connected to specific ports on a switch are assigned to a specific VLAN. For example, assume the computers connected to switch ports 2, 3, and 4 are assigned to the Sales VLAN 2 while the computers connected to switch ports 6, 7, and 8 are assigned to the Engineering VLAN 3 as shown in Figure 9-14. The switch will be configured as a port-based VLAN so that the groups of ports [2,3,4] are assigned to the sales VLAN while ports [6,7,8] belong to the Engineering VLAN. The devices assigned to the same VLAN will share broadcasts for that LAN however, computers that are connected to ports not assigned to the VLAN will not share the broadcasts. For example, the computers in VLAN 2 (Sales)

VLAN (Virtual LAN)
A group of host computers and servers that are configured as if they are in the same LAN even if they reside across routers in separate LANs

Port-based VLAN
Host computers connected to specific ports on a switch are assigned to a specific VLAN

Tagged-based VLAN
Used VLAN ID based on 802.1Q

Protocol-Based VLAN
Connection to ports is based on the protocol being used

share the same broadcast domain and computers in VLAN 3 (Engineering) share a different broadcast domain.

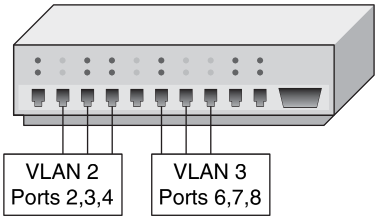

FIGURE 9-14 An example of the grouping for port-based VLANs.

In tag-based VLANs, a tag is added to the Ethernet frames. This tag contains the VLAN ID that is used to identify that a frame belongs to a specific VLAN. The addition of the VLAN ID is based on the 802.1Q specification. An advantage of an 802.1Q VLAN is it helps to contain broadcast and multicast data traffic that helps to minimize data congestion and improve throughput. This specification also provides guidelines for a switch port to belong to more than one VLAN. Additionally, the tag-based VLANs can help provide better security by logically isolating and grouping users

In a protocol-based VLANs, the data traffic is connected to specific ports based on the type of protocol being used. The packet is dropped when it enters the switch if the protocol doesn't match any of the VLANs. For example, an IP network could be set up for the Engineering VLAN on ports 6,7,8 and an IPX network for the Sales VLAN on ports 2,3, and 4. The advantage of this is the data traffic for the two networks is separated.

There are two approaches for assigning VLAN membership:

- **Static VLAN:** Basically a port-based VLAN. The assignments are created when ports are assigned to a specific VLAN.
- **Dynamic VLAN:** Ports are assigned to a VLAN based on either the computer's MAC address or the username of the client logged onto the computer. This means that the system has been previously configured with the VLAN assignments for the computer or the username. The advantage of this is the username and/or the computer can move to a different location, but VLAN membership will be retained.

Static VLAN
Basically a port-based VLAN

Dynamic VLAN
Ports are assigned to a VLAN based on either the compuuter's MAC addrress or the username of the client logged onto the computer

Switch Configuration

This section examines the basics of configuring a Cisco switch. The commands for switch configuration are similar to that of a router. Configuring a switch requires that the privileged mode be entered on the switch. The privileged EXEC mode (also called the enable mode) allows full access for configuring the switch ports and establishing a VLAN. This section focuses on general configuration steps for the switch, examining MAC address information and IP address configuration of the VLANs.

The privileged mode is entered using the command ***enable*** at the Switch> prompt as shown. The # sign after the switch name indicates you are in the privileged EXEC mode (Switch#).

```
Switch> enable
Password:
Switch#
```

Entry into the switch's privileged mode is typically password-protected. The exception to this is when a switch has not been configured and a password has not been assigned to it. In this case, pressing **Enter** on the keyboard from the Switch> prompt will promote the user to the privilege mode (Switch#) without requesting a password.

Use caution once you have entered the privileged mode in a switch. It is easy to make mistakes, and incorrectly entered switch configurations will adversely affect your network. This text comes with the Net-Challenge switch simulator on the companion CD-ROM to help you gain experience with switch configuration. In fact, most of the switch configuration commands presented in this section can be implemented in the switch simulator on the companion CD-ROM.

Hostname

The next commands examined require that the switch's terminal configuration mode be entered. To do this, enter the command ***configure terminal*** (abbreviated ***conf t***) at the switch# prompt to enter the switch's configuration mode.

Configure Terminal (*conf t*)
Command to enter the swtch's terminal configuration mode

```
switch#conf t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#
```

Or

```
switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)#
```

Note the change in the prompt to switch(config)#. This indicates that the switch is in terminal configuration mode.

The first switch configuration option examined enables the user to assign a hostname to the switch. The generic name or the name of an unconfigured Cisco switch is *switch*, and the ***hostname*** command enables the user to change the name to specifically identify the switch. For example, valid hostname structures include switchA, switch-A, or switch_A, while switch A is not valid. Valid switch hostnames may not have any spaces. The word *switch* does not have to be used in the hostname.

In the privileged mode (switch#) enter the command ***hostname [switch-name] <enter>***. This sets the hostname for the switch to *switch-name*. The following example demonstrates how the switch's hostname is changed to *SwitchA*. Notice the change with the switch's name from the first to the second line after the ***hostname*** command is entered.

```
switch(config)# hostname SwitchA
SwitchA#
```

Enable Secret

Switch(config)#
The prompt for the switch’s terminal configuration mode

Password protection for the privileged (enable) mode is configured by setting the enable secret. The steps are as follows: Enter the switch’s configure terminal mode by entering ***configure terminal*** or ***conf t*** at the **Switch#** prompt. Enter the command ***enable secret [your-password] <enter>***. An example is shown here:

```
SwitchA# conf t
SwitchA(config)#
SwitchA(config)# enable secret my-secret
```

This example sets the password for entering the switch’s privileged EXEC mode to *my-secret*. The password for entering the switch’s privileged mode must now be entered to gain access to the mode.

Setting the Line Console Passwords

Switch(config-line)#
The prompt indicating you are in the switch’s line configuration mode

The switch has two line connections through which a user may gain access to the switch. The line connections available on a switch can be displayed using the ***line ?*** command at the **Switch(config)#** prompt. The available line connections typically are as follows:

console	Primary terminal line (console port)
vty	Virtual terminal (for a telnet connection)

The *console* (primary terminal line) is the console port, and *vty* is the virtual terminal used for telnet connections. The following steps demonstrate how to configure password protection for the console port and the virtual terminal.

The console port configuration is as follows: Enter the command ***line console 0 <enter>***. Next enter the command ***login***, press **Enter**, and then input the command ***password [my-secret2]***, where *my-secret2* is the console port password. An example of using these commands follows.

```
SwitchA(config)# line console 0
SwitchA(config-line)# login
SwitchA(config-line)# password my-secret2
```

Note the change in the switch prompt to SwitchA(config-line)#, indicating you are in the switch’s line configuration mode.

Password protection for the virtual terminal (line vty) is set from the switch’s configuration mode. The virtual terminal is the method for entering the switch via a Telnet connection. The command ***line vty 0 15*** is first entered. This places the switch in the line configuration mode (config-line). The “0 15 “ indicates that sixteen virtual terminal connections can be simultaneously made: The sixteen virtual terminal connections are identified as 0, 1, 2, 3, 4 . . . . Next enter ***login***, press **Enter**, followed by entering the command ***password [my-secret3]***. The entry *my-secret3* is the password for the virtual terminal connection.

```
SwitchA(config)# line vty 0 4
SwitchA(config-line)# password my-secret3
SwitchA(config-line)# login
```

Layer 3 access to the switch is set by using the following command. Note that the IP address is being set for VLAN 1. The interface for the switch is also enabled at this same point using the ***no shutdown*** command as shown.

```
SwitchA(config)# interface VLAN 1
SwitchA(config-if)#ip address 172.16.32.2 255.255.255.0
SwitchA(config-if) no shutdown
```

The default gateway for the switch can be set using the following command. The default gateway instructs the switch where to forward the data packet if there isn't a specified route defined in the switch.

```
SwitchA(config)# ip default-gateway 172.16.35.1
```

The configuration settings entered on the VLAN 1 interface can be viewed by entering the following command:

```
SwitchA#show interface VLAN 1
```

The running configuration for the switch is viewed using the ***show running-config*** command as shown. The startup configuration is displayed using the ***show startup-config***, and the running-configuration file is copied to NVRAM using the ***copy running-config startup-config*** command as shown.

```
SwitchA#show running-configuration
SwitchA#copy running-configuration start-up configuration
```

These examples show that there is a lot of similarity between switch and router configuration at the command line interface. However, the major differences are apparent when configuring a VLAN. The next section demonstrates the steps for configuring a Static VLAN.

Static VLAN Configuration

This section demonstrates the steps for configuring a Static VLAN. In this example, the ports for VLAN 2 (Sales) and VLAN 3 (Engineering) will be defined. This requires that VLAN memberships be defined for the required ports. The steps and the commands will be demonstrated.

The first step for configuring the VLAN is to establish a terminal connection to the switch using the Cisco console cable. The console connection is used to perform the initial configurations that are needed to use the Cisco Network Assistant software.

1. Connect the Console cable to your workstation and switch.
2. Open the HyperTerminal software on your workstation.
3. Make a HyperTerminal connection to the switch. (Same steps as connecting to a router)

Now that you have made a HyperTerminal connection to the switch, the switch's initial prompt should appear as Switch>.

1. At the initial prompt type **enable** or **en**. Once this is done the prompt should change to Switch#. This allows you to enter the privileged mode of the switch.
2. Next type **configure terminal** or **conf t**. Once this is done the prompt should change to Switch(config)#. Doing this places you in the global configuration mode of the switch.
3. Now type **interface Vlan1** or **int Vlan1**. The prompt should now change to *Switch(config-if)#*. You are now able to make changes to Vlan1's interface.
4. Next enter **ip address 192.168.1.1 255.255.255.0**. This command changes the switches IP address to 192.168.1.1
5. Type **no shut**. This is needed for Vlan1 to stay up and active.

You have now set the IP address of the switch to 192.168.1.1. Notice that the subnet mask is set to 255.255.255.0 which places the switch in the 192.168.1.0 network. The IP address is set for Vlan1 because this is the default administrative VLAN for the switch, and it can never be removed. The workstation should be connected to port 1 on the switch, and the computers IP address's should be configured for 192.168.1.2. This places the computer in the same network as that defined for the VLAN1 interface. At this point, use the ***ping*** command to verify network connectivity from the computer to the switch.

The next step is to use the ***show vlan*** command to verify what ports have defined for the switch. By default, all ports are assigned to VLAN 1. An example using the ***show vlan*** command is provided next.

```
Switch# show vlan

VLAN Name                       Status          Ports
---- -------------------------- --------- -----------------------------
1    default                    active        Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                              Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                              Fa0/9, Fa0/10
```

This shows that all of the FastEthernet interfaces are currently assigned to VLAN 1. In the next step, two additional VLANs will be created for both the Sales and Engineering. This is accomplished modifying the VLAN database as shown in the next steps.

```
SwitchA#vlan database

Switch(vlan)#vlan 2 name Sales
VLAN 2 modified:
    Name: Sales
Switch(vlan)#vlan 3 name Engineering
VLAN 3 modified:
    Name: Engineering
```

The next step is used to verify that the new VLANs have been created.

```
Switch(vlan)# exit
Switch#show vlan

VLAN Name                       Status          Ports
---- -------------------------- --------- -----------------------------
1    default                    active        Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                              Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                              Fa0/9, Fa0/10
2    Sales                      active
3    Engineering                active
```

In the next steps, ports will be assigned to the newly created VLANs. This requires that the configuration mode be entered and each FastEthernet interface (port) must be assigned to the proper VLAN. An example is presented for FastEthernet interface 0/2 being assigned to VLAN 2.

```
Switch#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)#int fa 0/2
Switch(config-if)#switchport mode access
Switch(config-if)#switchport access vlan 2
Switch(config-if)#end
```

The next step is used to verify that FastEthernet 0/2 has been assigned to the Sales VLAN (VLAN2). This can be verified using the ***show vlan brief*** command as shown. This command only displays the interfaces assigned to each VLAN.

```
Switch#sh vlan

VLAN Name                               Status      Ports
---- ----------------------------- --------- -------------------------
1    default                           active     Fa0/1, Fa0/3, Fa0/4, Fa0/5
                                                  Fa0/6, Fa0/7, Fa0/8, Fa0/9
                                                  Fa0/10
2    Sales                             active     Fa0/2
```

The next steps are to assigned ports 3 and 4 to the Sales VLAN (VLAN 2) and ports 6,7,8 to Engineering (VLAN 3). Once this is completed, the port assignments can be verified using the show VLAN command as shown.

```
Switch#show vlan

VLAN Name                              Status    Ports
---- ---------------------------- --------- ---------------------------
1    default                           active    Fa0/1, Fa0/5, Fa0/9, Fa0/10

2    Sales                             active    Fa0/2, Fa0/3, Fa0/4

3    Engineering                       active    Fa0/6, Fa0/7, Fa0/8
```

You can look specifically at the assignments for only one of the VLANs by entering the command ***show vlan name <vlan-name>*** where vlan-name is the name assigned to the VLAN. Please note that the name is case-sensitive. You can also use the number of the VLAN instead using the command ***sh vlan id <vlan#>***. Examples of both are presented.

```
Switch#sh vlan name Engineering

VLAN Name                                 Status     Ports
---- ------------------------------- --------- ------------------------
3    Engineering                          active     Fa0/6, Fa0/7, Fa0/8

Switch#show vlan id 3

VLAN Name                                  Status    Ports
---- ------------------------------- --------- ------------------------
3    Engineering                           active    Fa0/6, Fa0/7, Fa0/8
```

The overall configuration of the switch can be viewed using the ***show running-config (sh run)*** command as shown. Only a part of the configuration is displayed.

```
Switch#sh run      -    -
Building configuration...

Current configuration : 1411 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Switch
!
ip subnet-zero
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface FastEthernet0/1
!-
```

```
 interface FastEthernet0/2
 switchport access vlan 2
 switchport mode access
    .       .
    .     .
    .       .
    .     .
interface FastEthernet0/5
!
interface FastEthernet0/6
 switchport access vlan 3
 switchport mode access
!
interface FastEthernet0/9
!
interface FastEthernet0/10
!
!
interface Vlan1
 ip address 192.168.1.1 255.255.255.0
 no ip route-cache
!
ip http server
!
line con 0
line vty 0 15
 login
end
```

The running-configuration for the switch shows that the FastEthernet interfaces have been assigned to the proper VLANs. Additionally this shows that an IP address has been assigned to the default interface VLAN1.

This portion of the text has demonstrated the steps for creating a static VLAN. Both Sales and Engineering VLANs were created, and specific ports on the switch were assigned to the respective VLANs. Unassigned ports remained as part of the default VLAN 1.

Networking Challenge—Static VLAN Configuration

Use the simulator software included with the text's Companion CD-ROM to demonstrate that you can perform basic switch and static VLAN configuration. Place the CDROM in your computer's drive. Open the *Net-Challenge* folder, and click **NetChallenge.exe**. Once the software is running, click the **Select Challenge** button. This opens a Select Challenge drop-down menu. Select **Chapter 9—Static VLAN Configuration**. This opens a check box that can be used to verify that you have completed all the tasks.

1. Enter the privileged EXEC mode on the switch.
2. Enter the switch's configuration mode, **Router(config)**.
3. Set the hostname of the switch to switch-A.
4. Configure the IP address for VLAN 1 interface with the following:
   IP address: 10.10.20.250
   Subnet mask: 255.255.255.0
5. Enable the VLAN 1 interface.
6. Use the command to display the current VLAN settings for the switch.
7. Issue the command that lets you modify the VLAN database.

8. Create a VLAN called Sales
9. Verify that a new VLAN has been created.
10. Issue the command to enter the fa0/2 interface configuration mode.
11. Enter the sequence of commands that are used to assign interface fa0/2 and fa0/3 to the Sales VLAN.
12. Enter the command that enables you to display the interface assigned to each VLAN.
13. Enter the command that enables you to view specifically the assignments for the Sales VLAN.
14. Issue the command that allows you to view the switch's running-configuration.

Spanning-Tree Protocol

This last section on switches examines the **Spanning-Tree Protocol (STP)**. STP is a link management protocol that prevents looping and also controls data flow over possible redundant data paths. Looping is bad for Ethernet networks because duplicate packets can be sent over redundant paths. The switches should only send the packets over one path. The Spanning Tree Protocol is used to ensure only one data path is selected. The Spanning Tree Protocol also forces one of the redundant data paths into a stand-by mode by placing the path in a blocked state.

Switches that are participating in the Spanning-Tree Protocol exchange information with other switches in the form of **bridge protocol data units (BPDUs)**. Switches use the BPDUs for the following:

- Election of a root switch for the spanning-tree network topology.
- Removing redundant data paths.
- The shortest distance to a root switch is calculated.
- A port from each switch is selected as the best path to the root switch.
- Ports that are part of the Spanning-Tree Protocol are selected.

Switches assume they are the root switch until the BPDUs are exchanged and a root switch is elected. The root switch elected is the switch with the lowest MAC address. Part of the BPDU packet is shown. In this case a switch with the MAC address 0030194A6940 is issuing that data packet as start of the bidding process to see which switch will be elected as the "root" switch.

```
BPDU  Config  BID=0030194A6940 PID=0x801B
```

The "Config" indicates this is a **Configuration BPDU** and is used by the switches to elect the "root" switch. There are two other types of packets that can come from the switch. These are the **Topology Change Notification (TCN)**, which is used to indicate that a there has been a change in the switch network topology. The third is the **Topology Change Notification Acknowledgement (TCA)**. This is an acknowledgement from another switch that the TCN has been received.

An example of the contents of a BPDU is provided in Figure 9-15. This is showing that the Root ID—MAC address is 0030194A6940. The BPDUs are exchanged at regular intervals and are used to keep the switches notified of any changes in the network topology. The default notification interval is 2 seconds and is called the the "Hello Time," as shown in Figure 9-15.

Spanning-Tree Protocol
A link management protocol that prevents looping and also controls data flow over possible redundant data paths

Bridge Protocol Data Unit (BPDU)
Used by switches to share information with other switches that are participating in the Spanning-Tree Protocol.

Configuration BPDU
Used by switches to elect the "root" switch

Topology Change Notification (TCN)
Used to indicate that there has been a change in the switch

Topology Change Notification Acknowledgement (TCA)
An acknowledgement from another switch that the TCN has been received.

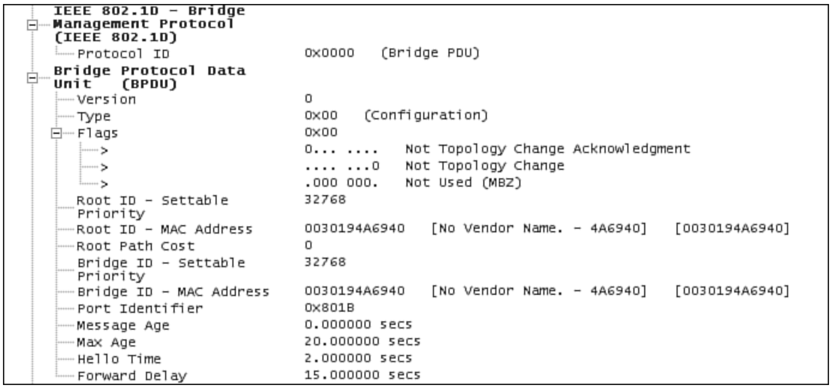

```
IEEE 802.1D - Bridge
Management Protocol
(IEEE 802.1D)
    Protocol ID                0x0000   (Bridge PDU)
Bridge Protocol Data
Unit   (BPDU)
    Version                    0
    Type                       0x00   (Configuration)
    Flags                      0x00
        >                      0... ....   Not Topology Change Acknowledgment
        >                      .... ...0   Not Topology Change
        >                      .000 000.   Not Used (MBZ)
    Root ID - Settable         32768
    Priority
    Root ID - MAC Address      0030194A6940   [No Vendor Name. - 4A6940]   [0030194A6940]
    Root Path Cost             0
    Bridge ID - Settable       32768
    Priority
    Bridge ID - MAC Address    0030194A6940   [No Vendor Name. - 4A6940]   [0030194A6940]
    Port Identifier            0x801B
    Message Age                0.000000 secs
    Max Age                    20.000000 secs
    Hello Time                 2.000000 secs
    Forward Delay              15.000000 secs
```

FIGURE 9-15 An example of a BPDU—Bridge Protocol Data Unit packet information

The switch will not begin to forward data packets when a networking device is connected to a port. Instead, during this delay, the switch will first begin to process the BPDUs to determine the topology of the switch network. This is called the forward delay, which is listed in Figure 9-15. This is showing that the forward delay is 15 seconds, which is the default value set by the root switch. During the delay period, the switch is going through the listening and learning states.

There are five Spanning-Tree Protocol states:

- **Blocking State:** In this state, the switch is not sending data out of the ports. However, the switch is receiving and monitoring the BPDUs. This state is used to prevent any possible switching loops.
- **Listening State:** BPDUs are being processed.
- **Learning State:** The switch is learning source MAC addresses from the received data packets and will add the addresses to the MAC address table.
- **Forwarding State:** The switch is now sending and receiving data packets. The BPDUs are still to be monitored for any possible change in the switch network.
- **Disabled:** This is a setting available for the network administrator to manually disable the port. This is not part of the Spanning-Tree Protocol but rather a function available on the switch.

9-7 ANALYZING CAMPUS NETWORK DATA TRAFFIC

The focus of this chapter has been on the issues of configuring and managing the campus network. A key issue in network management is network monitoring and the collection of utilization and error statistics.

Section 9-5 introduced the SNMP protocol for use in network management. An example was presented that shows how to obtain the number of octets leaving a router. This type of information can be used in a campus network to monitor the flow of data for many points in the network. Statistics can be obtained for hourly, daily, weekly, and monthly data traffic. This section discusses plots of network router utilization obtained via the router's SNMP port.

Figure 9-16 is a plot of a router's hourly data traffic. The plot shows the average number of bits coming into the router and the average number of bits out. The network administrator should become familiar with the typical hourly data traffic pattern for their network. Notice the decrease in data traffic in the early morning and the dramatic increase in data traffic around 12:00. The traffic clearly shows some type of disturbance around 12:00. The plot is showing that the bit rate significantly increases for a few minutes. This is not necessarily a problem, but it is something that a network administrator will want to watch.

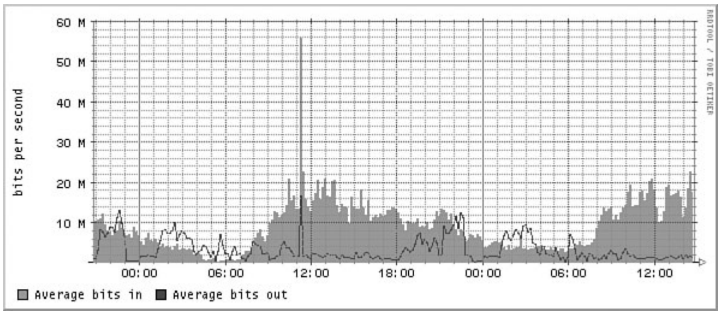

FIGURE 9-16 The hourly plot of a router's data traffic.

In this case, the network administrator looked at the daily log of network activity for the same router. This plot is shown in Figure 9-17. The cycle of the data traffic from morning to night is as expected, heavy data traffic about noon and very low data traffic in the mornings. An interesting note is the noon data traffic spikes on the first Wednesday and then repeats the following Wednesday. Whatever is causing the change in traffic appears to happen on Wednesdays. If this sudden change in data traffic turned out to be something of concern, a protocol analyzer could be set up to capture the data traffic on Wednesdays around noon so that the traffic pattern could be explained.

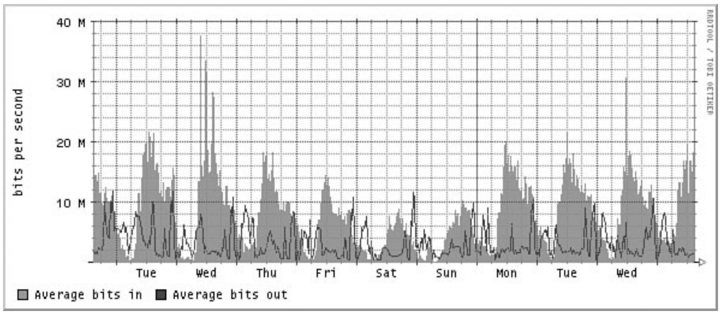

FIGURE 9-17 The daily plot of a router's data traffic.

Sometimes the graph of the network traffic over a longer period of time is needed. Figure 9-18 shows the data traffic through the router over a six-week period. The traffic shows some consistency except for a change from week 11 to week 12. Most likely this can be explained by examining the network trouble reports and maintenance logs to see if this router was briefly out of service.

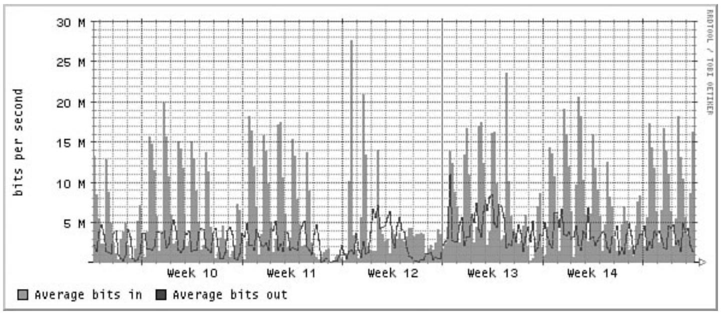

FIGURE 9-18 The weekly plot of a router's data traffic.

Justifying the expansion of a network's capability (for example, higher data rate or better core or distribution service) requires showing the manager data traffic statistics. Figure 9-19 is a plot of the router's monthly data traffic. The summer shows a significant decrease in data traffic. The plot also shows that the network was down once in the June–July period and again in January. The manager wants to know if there is justification to increase the data rate of the router to 1 gigabit (1 GB). (The router's current data rate is 100 Mbps.) Is there justification to upgrade the router to 1 GB? Probably not, at least not immediately. The maximum measured average data rate is about 16 Mbps. The router's 100 Mbps data rate does not seem to be causing any traffic congestion problems.

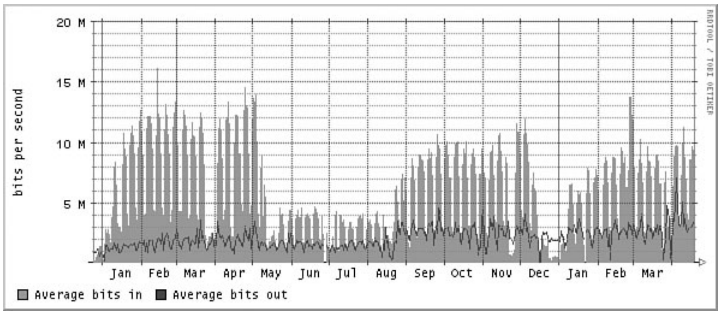

FIGURE 9-19 The monthly plot of a router's data traffic.

This section has shown how keeping logs of data traffic can be used to spot potential network problems and to help plan for possible future expansion of the network.

SUMMARY

The fundamentals of configuring and managing a campus network have been presented in this chapter. This has been a brief overview of the campus network, and you should understand that each of the topics presented in this chapter could easily be expanded to fill an entire textbook(s). What you should understand from this reading is that configuring and managing a campus network is a major task. It is critical for the networking group to have people who truly understand the many aspects of configuring and managing the campus network. Configuring and managing a campus DNS and DHCP service is a challenging task. Planning the integration of the equipment in the layers for the campus network is also a challenging task. Monitoring the network activity requires additional networking staff resources. You should appreciate the fact that configuring and managing a campus type network requires the expertise of many people with many different networking capabilities as well as understand the following:

- The importance and function of the three layers of a campus network
- How DHCP and DNS services work
- The importance of incorporating network services such as DHCP and DNS into a campus network
- How SNMP management tools can be used to monitor the data traffic and performance within a campus network

QUESTIONS AND PROBLEMS

Section 9-2

1. What networking equipment is usually found in the core of a campus network?
2. How are route policies applied in the core?
3. What is the advantage of using a layer 3 switch in the core of the campus network?
4. Can a layer 2 switch be used in the core of the campus network? Why or why not?
5. What is the function of the distribution layer in a campus network?
6. Can routing policies be implemented in the distribution layer? Why or why not?
7. What is the purpose of the access layer?
8. The campus network servers are typically located in what layer?
9. Why are routers typically not interconnected at the distribution layer?
10. What is the name for the part of the campus network that carries the bulk of the routed data traffic?
11. List three criteria for selecting the network media. Which is the final decision factor?
12. Which media is the best choice in a campus network?
13. Define load balancing in terms of data traffic flow in a computer network.
14. Define per-destination load balancing.
15. Define per-packet load balancing.

16. Referring to Figure 9-1 from the beginning of the chapter, discuss how data flows from a computer in LAN B to a computer in LAN D. Assume that the routing protocol is OSPF, the cost of RouterB's ge-0/0/0 interface has been set to 10, and the cost of RouterB's ge-0/2/0 interface has been set to 20.

Section 9-3

17. With regards to campus DHCP service, the IP address assignment is based on what?
18. What is the subnet mask for creating the four subnets illustrated in Figure 9-2 earlier in the chapter?
19. How are BOOTP and DHCP related?
20. Define *lease time.*
21. What networking function is required if the DHCP server is not on the same LAN? Why is this networking function required?
22. What command enables a DHCP relay on a Cisco router?
23. Why is packet 14 in the captured DHCP packets shown in Figure 9-5 (shown earlier in the chapter) a broadcast?
24. What are the port numbers for the DHCP protocol?

Section 9-4

25. List 11 top level domains.
26. What is the purpose of a root server in DNS?
27. A new network wants to obtain a domain name. The first step is what?
28. The hostname and IP address for a computer is stored in what for a campus DNS service?
29. How is it possible for the command ***ping www.networkB.edu*** to find the destination without an IP address?
30. What is the purpose of reverse DNS? Where is it used?

Section 9-5

31. What port number does SNMP use and what transport protocol?
32. The SNMP MIB get request ***ifDescr*** returns what information from a router?
33. What is the purpose of the MIB?
34. Write the Cisco router command for configuring SNMP on a Cisco router. Assume a community string of networking and set the permissions to read-only. Show the router prompt.
35. The command ***show run*** is entered on a Cisco router. Describe what the output "SNMP-server test RO" means.
36. What SNMP MIBs were most likely issued to the router discussed in section 9-7?

Use Figure 9-20 to answer questions 37 to 41.

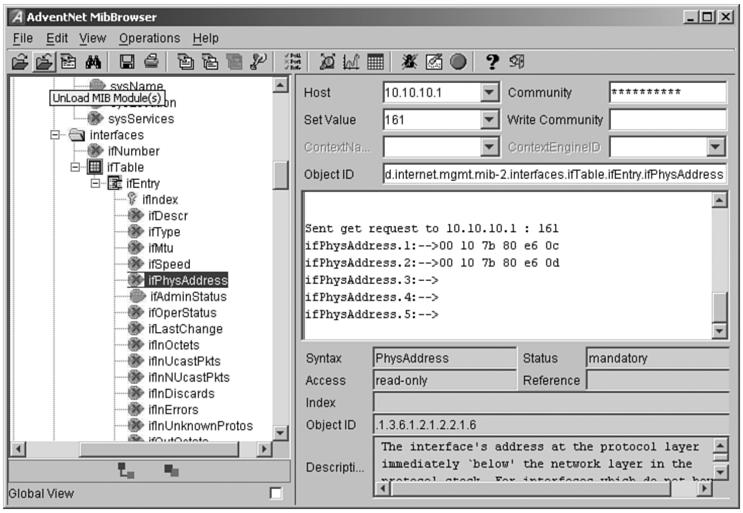

Host: 10.10.10.1
Community: **********
Set Value: 161
Write Community:
ContextNa...:
ContextEngineID:
Object ID: d.internet.mgmt.mib-2.interfaces.ifTable.ifEntry.ifPhysAddress

```
Sent get request to 10.10.10.1 : 161
ifPhysAddress.1:-->00 10 7b 80 e6 0c
ifPhysAddress.2:-->00 10 7b 80 e6 0d
ifPhysAddress.3:-->
ifPhysAddress.4:-->
ifPhysAddress.5:-->
```

Syntax: PhysAddress
Status: mandatory
Access: read-only
Reference:
Index:
Object ID: .1.3.6.1.2.1.2.2.1.6
Descripti...: The interface's address at the protocol layer immediately `below' the network layer in the

FIGURE 9-20 For problems 37–41.

37. What MIB was issued?
38. What information was returned?
39. What port number was used?
40. What protocol is being used? How do you know?
41. Who is the manufacturer of this networking device?
42. What are the two types of devices defined by PoE?
43. What should you check if you are installing a Power over Ethernet connection using computer equipment from two different manufactures?
44. Cite four benefits of Power over Ethernet.
45. What is resistive power discovery, and how does it work?
46. What wire pairs are used in PoE?
47. How much power can a class 0 PD PoE device source?
48. What are the benefits of PoE Plus?

Section 9-6

49. What is a VLAN?
50. List the three types of VLANs.
51. What type of VLAN is port-based?
52. What commands are used to assign the IP address 192.168.20.5 to VLAN1?
53. What switch command is used to display the interfaces assigned to a VLAN?
54. What is the purpose of the VLAN database?
55. List the commands used to create VLAN5 and name this VLAN Marketing group.
56. List the commands used to assign FA0/5 to the Marketing-group VLAN (VLAN5). Show the switch prompts.
57. What is the purpose of the Spanning-Tree Protocol?
58. What is a BPDU, and what its purpose?
59. Discuss how a root switch is elected.

60. What are the five STP protocol states?

61. A BPDU data packet shows that the “Hello Time” is 2.0 secs. What information does this provide?

62. A BPDU data packet lists the “Forward Delay” as 15 seconds. What information does this provide?

Critical Thinking

63. Your supervisor asks you if a layer 2 switch could be used in the core of the campus network. Prepare a response to your supervisor. Be sure to justify your recommendation.

64. A 1Gbps data link is to be set-up between building A and building B in a campus network. Does it matter if the link is fiber or microwave or some other media? Explain your answer.

This page intentionally left blank