This chapter covers the following topics:

- An ounce of planning
- Configuration overview
- Initial configuration
- Connecting the switches
- Configuring the access layer
- Configuring SNMP

CHAPTER 7

Configuring Switches

Now that you have learned about the concepts behind Layer 2 and Layer 3 switching in some detail, you will focus on a start-to-finish configuration of a relatively simple campus switching design in this chapter.

An Ounce of Planning

Everyone has probably heard the old joke "ready, fire, aim." Unfortunately, this phrase can sometimes describe the implementation of some networks given what appears to be a lack of basic planning prior to configuration. The daily operation of a switched environment can be greatly simplified and future problems avoided by applying a few best practices and a little bit of planning. This begins with planning the method for remotely accessing the switch, followed by basic configuration of the switch, and then configuring connections between switches.

Management Interfaces

Believe it or not, one of the first things to think about when configuring a new network is management, primarily because network management typically is the last thing to be thought of when the network is implemented, and seemingly one of the most tedious things to change or improve after the network is operational. One item to consider is how to handle remote access to the switch. Catalyst switches support both in-band and out-of-band management. In-band management interfaces are connected to the switching fabric and participate in all the functions of a switchport including spanning tree, Cisco Discovery Protocol (CDP), and VLAN assignment. Out-of-band management interfaces are not connected to the switching fabric and do not participate in any of these functions.

Out-of-band management is achieved initially through the serial console port on the Supervisor module. Each Catalyst switch ships with the appropriate console cable and connectors to connect to a host such as a Windows workstation or terminal server. Consult the Catalyst documentation at Cisco.com to determine the kind of connectors and cables appropriate for each platform. After a physical connection is made between the console port on a Catalyst switch and a serial port on a workstation or terminal server, the administrator has full access to the switch for configuration. At this point, the administrator can assign an IP address to either an out-of-band management (sl0) interface via the Serial Line Internet Protocol (SLIP), a predecessor to the Point-to-Point Protocol (PPP), or assign an IP address to an in-band management interface (sc0 or sc1). Supervisors for the Catalyst 4500 series switches offer an additional out-of-band management interface via a 10 Mbps or 10/100 Mbps Ethernet interface (me1) depending on the Supervisor model.

The choice between out-of-band and in-band management is often not an easy one because each has its pros and cons. An in-band management connection is the easiest to configure and the most cost effective because management traffic rides the same infrastructure as user data. Downsides to in-band management include a potential for switches to be isolated and unmanageable if connectivity to the site or individual device is lost, for example in a spanning-tree loop or if fiber connections are cut accidentally. In addition, if the management interface is assigned to a VLAN that has other ports as members, any broadcast or multicast traffic on that VLAN is seen by the management interface and must be processed by the supervisor.

As the speed of processors has improved with newer supervisors, the risk of overwhelming a supervisor with broadcast/multicast traffic has declined somewhat, but has not been eliminated completely. With these drawbacks to in-band management, why doesn't everyone just use out-of band management? The answer is simple: time and money. Out-of-band management requires a secondary infrastructure to be built out around the devices such as terminal servers, switches, and modems. The benefit of an out-of-band management solution is that it offers a completely separate method of connecting to the devices for management that does not rely upon a properly functioning data infrastructure to work.

Many administrators find themselves implementing both in-band and out-of-band management solutions depending on the reliability of the data infrastructure between the networks that contain the management stations and the devices being managed. For example, Catalyst switches in a typical headquarters location are likely to be on reliable power grids, potentially with backup power, and have redundant connections between devices. A Catalyst switch in a remote office connected to headquarters via a router and a single nonredundant Frame Relay connection might justify out-of-band management. The remote router and switch could be connected to a terminal server and an analog dial-up connection for configuration and remote management. In an ideal world, networking devices would all be accessible via an out-of-band connection, if possible. Sometimes it takes only a wake-up call at 3:00 a.m. or an unplanned road trip to a remote location to compel an organization to install an out-of-band management solution.

sc0 and VLAN 1

All switchports must be members of a VLAN, and, by default, it is VLAN 1. Because VLAN 1 was selected as the default VLAN for all switchports, it was also chosen to handle special traffic such as VLAN Trunking Protocol (VTP) advertisements, CDP, Port Aggregation Protocol (PAgP), or Link Aggregation Control Protocol messages (LACP). By default, in-band management interfaces such as sc0 are members of VLAN 1.

Over the years, a common scenario involving VLAN 1 and the management interface developed. In this scenario, administrators assigned an IP address to sc0, left it in VLAN 1, and created other VLANs for all user traffic. All ports not changed or enabled remain in VLAN 1. Trunked ports between switches are created to connect VLANs, and, by default, all VLANs (1-1005 or 1-4096 depending on trunk type and switch software version) are allowed across a trunk. Because each switch will have a management interface, likely sc0, this can result in VLAN 1 spanning the entire switched network. Remember that IEEE spanning tree only allows seven switch hops between end stations, and many times large networks that allow all VLANs to be trunked can approach or exceed the limit, especially for VLAN 1. When a spanning tree exceeds seven switch hops, the spanning-tree topology can become unpredictable during a topology change and reconvergence can be slow if the spanning tree reconverges at all. A few different options should be considered to alleviate this problem. The first option is to use a different VLAN other than VLAN 1 for the management interfaces in the network. As of Catalyst OS version 5.4(1) and later, VLAN 1 can be cleared from both Inter-Switch Link (ISL) Protocol and 802.1q trunks, thus removing VLAN 1 from the spanning-tree topology on those trunks. Simply substituting a different VLAN number does not alleviate the problem of new VLAN spanning the switched network and potentially exceeding the allowed number of hops. To avoid the problem, either multiple VLANs must be dedicated to network management or the management interfaces must be placed in multiple VLANs along with user traffic. Either way, the management interfaces must be reachable by the network management stations. In the configuration examples later in this chapter, the sc0 interface is placed in a user VLAN along with other ports.

Figure 7-1 shows a simple network diagram of a small remote office with multiple switches. In this figure, VLAN 501 is used as the management VLAN at the remote office.

In a configuration like this, the VLAN numbers in the remote office are only locally significant. This is true because a Layer 3 routed connection separates the remote office from the headquarters location, and VLAN 501 is not carried across the WAN. As a result, the remote office could use any VLAN number for management including VLAN 1.

The example could get trickier if the routers and WAN connections are replaced by switches and a high-speed Gigabit connection between buildings in a campus environment. In this situation, as long as the links between buildings can still be Layer 3 connections and VLAN 501 is cleared from the trunks, it can yield the same result, as in Figure 7-1. Unfortunately, many times with existing implementations, because of legacy Layer 2-only implementations or application design considerations, the links between locations are Layer 2 trunks carrying all VLANs. As a result, VLAN 501 gets carried to the home office switches, and potential spanning-tree problems can result.

Figure 7-1 *Remote Office Using VLAN 501 for Management*

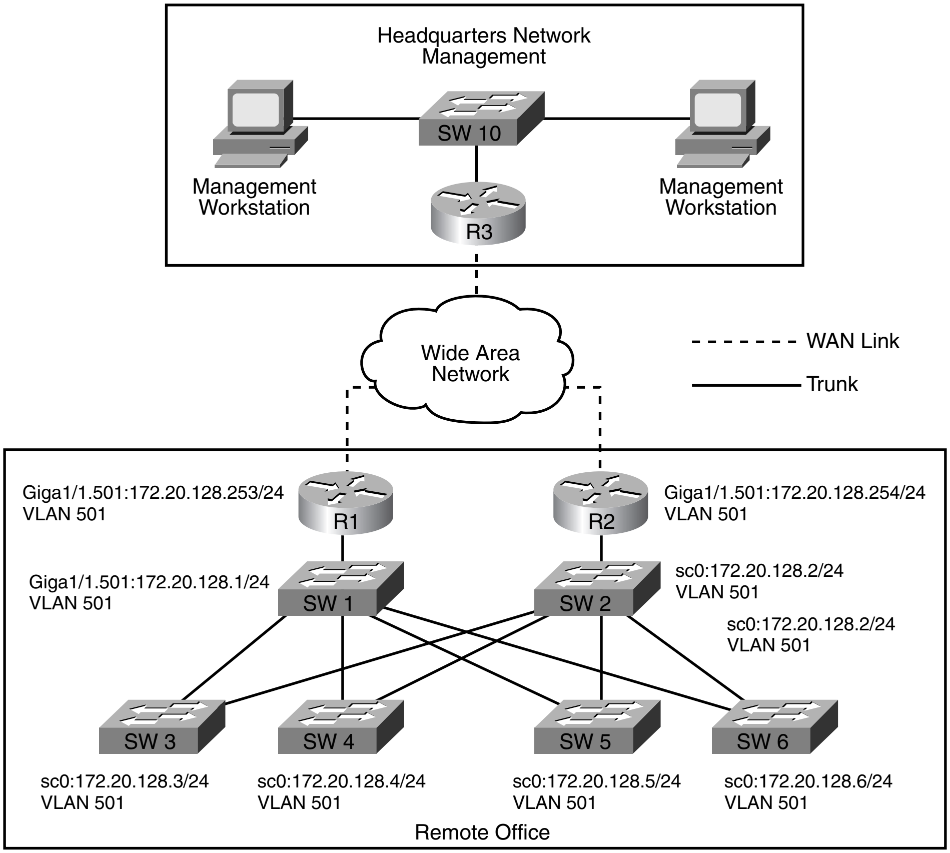

It is important to remember that even when VLAN 1 is cleared from a trunk, the previously mentioned special traffic, such as CDP, PAgP, and VTP, is still forwarded across the trunk with a VLAN 1 tag, but no user data is sent using VLAN 1. All trunks default to a native VLAN of 1 unless changed. In the case of an 802.1q trunk, where the native VLAN is untagged, 802.1q IEEE Bridge Protocol Data Units (BPDUs) are forwarded untagged on the common spanning-tree VLAN 1 for interoperability with other vendors, unless VLAN 1 has been cleared from the trunk. Cisco Per-VLAN Spanning Tree (PVST+) BPDUs are sent and tagged for all other VLANs. Refer to the sections on ISL and 802.1q trunking in Chapter 4, "Layer 2 Fundamentals," of this book for more information on trunking and native VLAN operation.

It is a good idea, if possible, to adopt some standards for VLAN numbering. Using consistent VLAN numbers for similar functions at multiple locations can many times help in the operation and troubleshooting of the networks later on. For example, many companies reserve certain VLAN ranges for specific functions. Table 7-1 is a sample of what a company might start with when implementing VLANs on existing and new networks.

Table 7-1 *Sample Plan for VLAN Numbering*

VLAN Numbers	Function
1	Not in use; clear from all trunks
2–99	Management VLANs (sc0)
100–399	Access layer devices
400–599	Data center devices
600–699	Internet and partner connections
700–899	Reserved for future use
900–999	Point-to-point links between switches (Layer 3)

Although this sample uses VLAN numbers in the 1–1005 range, newer versions of Cisco Catalyst OS and IOS support 4096 VLANs using 802.1q trunks. Again, because VLAN numbers are only locally significant when they are carried on trunks between switches, the sample numbering scheme provides great flexibility, and some companies may adopt a much more granular VLAN numbering system. For example, they may dictate that VLAN 50 be used as the management VLAN on all switches at all locations instead of allowing any VLAN in the range from 2–99 to be used. No hard and fast rules exist for VLAN numbering plans, and Table 7-1 represents one approach.

Configuration Overview

Figure 7-2 shows a diagram of four switches that are yet to have any connections configured. These four switches comprise the primary components of the network that will be configured and added onto throughout the remainder of this chapter.

Figure 7-2 *Four Primary Switches*

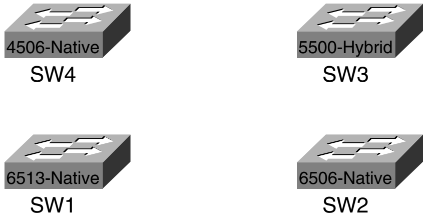

A mix of platforms, software levels, and interfaces was chosen to provide a variety of configuration examples. Table 7-2 lists the switches in use.

Table 7-2 *Switch Information*

Switch Name	Platform	Software Type	Software Version(s)
SW1	6513	Native	12.1(8b)E16
SW2	6506	Native	12.1(8b)E16
SW3	5500	Hybrid	4.5(5) Catalyst OS 12.2(10a) IOS
SW4	4506	Native	12.1(19)E

The software versions installed on these switches are not recommendations for these platforms, only versions that support the modules and features required for these exercises. Administrators should utilize the tools on the Cisco Software Center at Cisco.com, such as the IOS Upgrade Planner, Software Advisor, and Cisco Bug Navigator, to help select a satisfactory software level. After a software level has been selected, administrators should reference the release notes for a list of Open and Resolved Caveats in that version.

Initial Configuration

Configuration begins with naming each switch and assigning an IP address to a management interface on each switch shown in Figure 7-2. Refer to Chapter 5, "Using Catalyst Software," for examples of setting system and host names, along with setting an enable password. Private IP addresses described in Request for Comments (RFC) 1918 will be used in all the examples in this chapter.

NOTE RFC 1918 along with others can be viewed online at http://www.ietf.org/rfc. RFC 1918 defines private address ranges as

- 10.0.0.0–10.255.255.255 (10/8 prefix)
- 172.16.0.0–172.31.255.255 (172.16/12 prefix)
- 192.168.0.0–192.168.255.255 (192.168/16 prefix)

In this chapter, addresses from the 172.16.0.0–172.31.255.255 range are used.

Before implementing any IP equipment, take the time to develop an IP addressing standard. Going back and readdressing devices in production can be quite time consuming. Although development of an IP addressing standard is beyond the scope of this book, a few important items should be considered when developing a standard, including

- Planning the IP address space so it can be summarized, resulting in as few routes as possible being required to reach any network.
- Determining whether private, public, or a mix of private and public addressing will be used and how.
- Planning the IP address space to scale to the necessary number of devices. For example, assigning a network of 172.16.200.0/24 to a user VLAN on a switch provides 254 host addresses for user devices, but if 300 devices need to be supported, you must decide either to assign a second class C or /24 VLAN for the additional 46 devices or assign a larger network of 172.16.200/22.

In preparation for the configuration examples throughout the rest of this chapter, Table 7-3 provides a simple IP addressing scheme.

Table 7-3 *IP Address Ranges*

Function	IP Address Range
User VLANs	172.16.192–223.0 255.255.255.0
Loopback interfaces	172.16.224–239.0 255.255.255.255
Point-to-point links	172.16.240.4–252 255.255.255.252

Using the preceding ranges, IP addresses are plentiful because private addressing space is being used, but it is always a good practice to conserve addressing space whenever possible. The address ranges in Table 7-3 can all be summarized into a single 172.16.192.0/18 route advertisement.

Configuring VTP

Chapter 4 discussed the various modes and capabilities of VTP in detail. In this chapter, VTP transparent mode is used on all the example switches. A VTP domain name of Cisco is used. A VTP password is unnecessary in transparent mode but should be carefully chosen in client/server mode. Prior to Cisco IOS version 12.1(11b)E, VTP and VLANs could only be configured in VLAN database mode on IOS devices. In IOS version 12.1(11b)E and later, VTP and VLANs can be configured either in database mode or in global configuration mode. In either case, the VTP and VLAN configuration information is stored in a vlan.dat file and is not part of the running configuration. To properly back up a native IOS configuration, both the running-configuration and the vlan.dat file must be saved. CiscoWorks Resource Manager Essentials, starting with version 3.5, automatically saves the vlan.dat file. Using Examples 7-1 through 7-4, VTP is configured on each switch along with a VLAN that will be used for user devices later in the chapter.

Example 7-1 *Configuring VTP Using VLAN Database Mode on SW1*

```
SW1#vlan database
SW1(vlan)#vtp transparent
Setting device to VTP TRANSPARENT mode.
SW1(vlan)#vtp domain Cisco
Changing VTP domain name from NULL to Cisco
SW1(vlan)#vlan 110
VLAN 110 added:
    Name: VLAN0110
SW1(vlan)#exit
APPLY completed.
Exiting....
SW1#
```

Example 7-2 *Configuring VTP Using VLAN Database Mode on SW2*

```
SW2#vlan database
SW2(vlan)#vtp transparent
Setting device to VTP TRANSPARENT mode.
SW2(vlan)#vtp domain Cisco
Changing VTP domain name from NULL to Cisco
SW2(vlan)#vlan 120
VLAN 120 added:
    Name: VLAN0120
SW2(vlan)#exit
APPLY completed.
Exiting....
```

Example 7-3 *Configuring VTP on SW3 in Catalyst OS*

```
SW3 (enable) set vtp mode transparent
VTP domain  modified
SW3 (enable) set vtp domain Cisco
VTP domain Cisco modified
SW3 (enable) set vlan 130
Vlan 130 configuration successful
SW3 (enable)
```

Example 7-4 *Configuring VTP on SW4 in Global Configuration Mode*

```
SW4#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW4(config)#vtp mode transparent
Setting device to VTP TRANSPARENT mode.
SW4(config)#vtp domain Cisco
```

Example 7-4 *Configuring VTP on SW4 in Global Configuration Mode (Continued)*

```
Changing VTP domain name from NULL to Cisco
SW4(config)#vlan 140
SW4(config-vlan)#end
SW4#
```

When created, you can delete VLANs one at a time using either a **clear** command in Catalyst OS or the **no** form of the VLAN command in IOS. To delete all VTP and VLAN information in Catalyst OS, you can use the **clear config all** command. Although there is no vlan.dat file in Catalyst OS, the vlan.dat file is stored in NVRAM on Catalyst 6000/6500s in const_nvram: and 4000/4500s running native in cat4000_flash:. To delete all VTP and VLAN information in native IOS on the Catalyst 6000/6500, use the **erase const_nvram:** command. On the Catalyst 4000/4500 running native IOS, use the **erase cat4000_flash:**. As shown in Example 7-5, vlan.dat files can be copied to flash or to a TFTP server using the copy command.

Example 7-5 *Copying vlan.dat to Flash in Slot0:*

```
SW1#copy const_nvram:
SW1#copy const_nvram:vlan.dat slot0:
Destination filename [vlan.dat]?
660 bytes copied in 0.328 secs
```

Configuring sc0 and LO0

Because SW3 is running hybrid, it will be configured with both a sc0 management interface in Catalyst OS and a Loopback 0 (LO0) interface in IOS on the Route Switch Module (RSM). Figure 7-3 shows the management interfaces assigned to each of the four switches.

Figure 7-3 *IP Addresses Assigned to Management Interfaces*

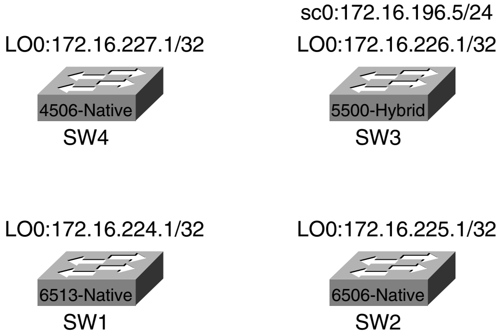

To prevent the use of a separate VLAN for switch management, the choice is made to place the sc0 interface in the user VLAN 130. Switches 1, 2, and 4 running native IOS are configured with only a Loopback interface (LO0), just like any other Cisco router. The primary benefit of a loopback interface is that it never goes down unless manually shut down. Example 7-6 shows the configuration of LO0 on SW1.

Example 7-6 *Configuring LO0 on SW1 (Native)*

```
SW1#config t
1w5d: %SYS-5-CONFIG_I: Configured from console by console
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#interface loopback0
SW1(config-if)#ip address 172.16.224.1 255.255.255.255
SW1(config-if)#end
1w5d: %SYS-5-CONFIG_I: Configured from console by console
SW1#show interface loopback0
Loopback0 is up, line protocol is up
  Hardware is Loopback
  Internet address is 172.16.224.1/32
  MTU 1514 bytes, BW 8000000 Kbit, DLY 5000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation LOOPBACK, loopback not set
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue :0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
  L2 Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes mcast
  L3 out Switched: ucast: 0 pkt, 0 bytes
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     0 packets output, 0 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
SW1#
```

In Example 7-6, an IP address of 172.16.224.1 is assigned using a 32-bit subnet mask. The output of the **show interface loopback0** command shows the interface in the UP/UP state. Because this is a loopback, the interface will show up even though no connectivity to the switch exists, and the loopback interface is, at the moment, unreachable. Example 7-7 shows the configuration of LO0 on SW2.

In Example 7-8, sc0 is assigned an IP address of 172.16.196.5/24 in VLAN 130. The default route added for sc0 will eventually point to the IP address of the VLAN 130 interface on the RSM.

Example 7-7 *Configuring LO0 on SW2 (Native)*

```
SW2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)#interface loopback0
SW2(config-if)#ip address 172.16.225.1 255.255.255.255
SW2(config-if)#end
```

Example 7-8 *Configuring sc0 on SW3 (Hybrid-Catalyst OS)*

```
SW3> (enable) set int sc0 130 172.16.196.5 255.255.255.0
Interface sc0 vlan set, IP address and netmask set.
SW3> (enable) set ip route default 172.16.196.1
Route added.
```

Example 7-9 shows the configuration of LO0 on SW3.

Example 7-9 *Configuring LO0 on SW3 (Hybrid-IOS)*

```
SW3 (enable) show module
Mod Module-Name         Ports Module-Type           Model    Serial-Num Status
--- ------------------- ----- --------------------- --------- --------- -------
1                       0     Supervisor III        WS-X5530  030061500 faulty
3                       1     Route Switch          WS-X5304  006578507 ok
4                       24    10/100BaseTX Ethernet WS-X5224  009607843 ok
6                       12    100BaseTX Ethernet    WS-X5113  002503515 ok
7                       24    10/100BaseTX Ethernet WS-X5234  019554483 ok
8                       24    10/100BaseTX Ethernet WS-X5225R 013458239 ok
13                            ASP/SRP

Mod MAC-Address(es)                        Hw     Fw         Sw
--- -------------------------------------- ------ ---------- ----------------
1   00-90-86-66-50-00 to 00-90-86-66-53-ff 3.5    5.1(2)     4.5(5)
3   00-e0-1e-91-b9-7c to 00-e0-1e-91-b9-7d 7.7    20.22      12.2(10a)
4   00-10-7b-78-57-00 to 00-10-7b-78-57-17 1.4    3.1(1)     4.5(5)
6   00-40-0b-b0-95-40 to 00-40-0b-b0-95-4b 1.2    1.2        4.5(5)
7   00-30-7b-b7-77-00 to 00-30-7b-b7-77-17 1.0    4.5(2)     4.5(5)
8   00-d0-06-9b-83-10 to 00-d0-06-9b-83-27 3.3    4.3(1)     4.5(5)

Mod Sub-Type Sub-Model Sub-Serial Sub-Hw
--- -------- --------- ---------- ------
1   NFFC II  WS-F5531A 0030060943 2.2
SW3 (enable) session 3
Trying Router-3...
Connected to Router-3.
Escape character is '^]'.

RSM1>en
RSM1#config t
```

continues

Example 7-9 *Configuring LO0 on SW3 (Hybrid-IOS) (Continued)*

```
Enter configuration commands, one per line.  End with CNTL/Z.
RSM1(config)#int loopback0
RSM1(config-if)#ip address 172.16.226.1 255.255.255.255
RSM1(config-if)#end
RSM1#sh interface loopback0
Loopback0 is up, line protocol is up
  Hardware is Loopback
  Internet address is 172.16.226.1/32
  MTU 1514 bytes, BW 8000000 Kbit, DLY 5000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation LOOPBACK, loopback not set
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue :0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     0 packets output, 0 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
RSM1#
```

In Examples 7-8 and 7-9, the switch is running hybrid Catalyst OS/IOS and the connection is to the console port on the supervisor. The first step is to determine in which slot the RSM is installed, and then session to the RSM. In this case, the RSM is installed in slot 3. After a session to the module in slot 3 is established, the loopback interface is configured the same way as in native. (See Example 7-10.)

Example 7-10 *Configuring LO0 on SW4 (Native)*

```
SW4(config)#interface loopback0
SW4(config-if)#ip address 172.16.227.1 255.255.255.255
SW4(config-if)#end
SW4#
```

In each of the loopback configuration examples, the loopback interface is administratively up and the line protocol is up even though no active ports are configured on the switch. This again is because loopback interfaces are special and cannot go down unless administratively shut down. This is not true for VLAN interfaces because of a feature called autostate. It is important to understand how autostate operates, as you learn in the next section.

Autostate

Hybrid and native switches have a feature called *autostate*. The feature is enabled by default and can only be disabled in hybrid. In hybrid, logical VLAN interfaces configured on the RSM/RSFC, MSFC, or Layer 3 module on the Catalyst 4000 rely on ports in Catalyst OS to be active in the same VLANs before communication is possible. For example, it is possible to configure a VLAN interface on an MSFC for VLAN 100 without any switchports in Catalyst OS belonging to VLAN 100, or VLAN 100 even being defined in Catalyst OS for that matter. Because this is possible, the Cisco IOS portion of the hybrid configuration attempts to prevent a routing "black hole" by placing the VLAN interface in a down/down state. After one or more active ports or a trunk is configured in the same VLAN as the interface in Cisco IOS, the VLAN interface changes to an up/up state. This checking mechanism is the result of the autostate feature. One exception to this feature is for the VLAN assigned to the management interface (sc0) on the switch. The sc0 interface can be shut down administratively.

To further prevent black holes, the autostate feature on the Catalyst 6000/6500 waits for the valid Layer 2 port(s) to transition into a forwarding state before allowing the Layer 3 VLAN interface to transition to an UP/UP state. The autostate on the Catalyst 6000/6500 feature began synchronizing with spanning tree in this way starting in 5.5(10) and 6.1(1) Catalyst OS software.

The commands in Example 7-11 disable autostate depending on the platform.

Example 7-11 *Disabling Autostate on Catalyst 6000/6500 Hybrid*

```
Switch (enable) set msfcautostate disable
Switch (enable) show msfcautostate
MSFC Auto port state: disabled
Switch (enable)
```

A Catalyst 6000/6500 with dual MSFCs would require autostate to be disabled to allow traffic to flow between the MSFCs on that VLAN if no active ports existed. In most situations, this is not necessary, and autostate should be enabled unless a specific need exists to disable it. Example 7-12 shows autostate being disabled on a Catalyst 5500 with an RSM.

Example 7-12 *Disabling Autostate on Catalyst 5000/5500 with RSM*

```
Switch (enable) set rsmautostate disable
RSM port auto state disabled.
Switch (enable) show rsmautostate
RSM Auto port state: disabled
Multi-RSM Option: enabled
Switch (enable)
```

If autostate is enabled and no active ports exist on a specific VLAN in the switch, the interface on the RSM remains up if there is more than one RSM. Essentially, the RSMs see each other's interfaces as valid. This allows traffic to flow between the two RSMs on that VLAN without disabling the autostate feature. The autostate feature is enhanced for multi-RSM configurations starting in 6.1(2) Catalyst OS software. Multi-RSM allows the interfaces on two RSMs to go down when the last active port on that VLAN in the switch goes down. Example 7-13 shows autostate being disabled on a Catalyst 4000 using hybrid software.

Example 7-13 *Disabling Autostate on Catalyst 4000 Hybrid with a Layer 3 Module*

```
Router#autostate disable
Disabling Autostate
Router#show autostate entries
Autostate Feature is currently disabled on the system.
```

System Logging

Cisco devices including Catalyst switches generate a variety of system messages for events such as changes in interface status, environmental conditions, parity memory errors, and security alerts. These messages are displayed on the system console by default. Console logging is a high-priority task in Cisco IOS, and, in some cases, enough console messages can effectively hang the router or switch and render the console unusable. Cisco recommends disabling console and monitor logging and configuring the switch or router to send console messages to an internal buffer that is adjustable in size. Disabling monitor logging prevents system messages from being displayed on terminal lines. Table 7-4 lists the levels of syslog messages supported on a Cisco device.

Table 7-4 *Syslog Severity Levels, Types, and Descriptions*

Severity Level	Severity Type	Description
0	Emergencies	System unusable
1	Alerts	Immediate action is required
2	Critical	Critical condition
3	Errors	Error conditions
4	Warnings	Warning conditions
5	Notifications	Normal, but significant condition
6	Informational	Informational messages
7	Debug	Debugging messages

Examples 7-14 and 7-15 show console logging being disabled and logging to a buffer being enabled on both native and hybrid software.

Example 7-14 *Disabling Console and Monitor Logging and Enabling Logging Buffered (Native)*

```
SW1(config)#no logging console
SW1(config)#no logging monitor
SW1(config)#logging buffered 16384
SW1(config)#end
SW1#
```

Example 7-15 *Disabling Console and Monitor Logging and Enabling Logging Buffered (Hybrid-Catalyst OS)*

```
SW3> (enable) set logging console disable
System logging messages will not be sent to the console.
SW3> (enable) set logging buffer 500
System logging buffer size set to <500>
SW3> (enable)
```

The logging buffers in Examples 7-14 and 7-15 are specified in bytes and are circular, meaning the oldest log messages will be overwritten by the newest messages after the buffer is full. The maximum logging buffer size in Catalyst OS is 500 bytes. To view the contents of the logging buffer, use the **show log** command. One problem with relying only on the logging buffer is that it is wiped clean during a reload. A more effective solution for logging system messages is the addition of a syslog server. A *syslog server* is simply a machine running a syslog daemon conforming to the Berkley Standard Distribution (BSD) standard. A syslog server stores the messages in the order received for later viewing. Many network management tools such as CiscoWorks and HP Openview can operate as a syslog server. In larger environments, it is generally recommended to set up a dedicated syslog server because of the number of messages that can be generated each day by dozens or hundreds of Cisco devices. Example 7-16 shows the configuration of logging to a syslog server using native software.

Example 7-16 *Completing the Logging Configuration (Native)*

```
SW1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#logging 10.10.10.1
SW1(config)#logging facility local7
SW1(config)#logging trap notifications
SW1(config)#logging source-interface lo0
SW1(config)#
```

In Example 7-16, the switch is pointed to a syslog server at IP address 10.10.10.1 and sets the default logging facility for logging. The syslog server specified should also be set for the same facility/level. The switch is configured to send notification level (5) messages and above to the syslog server and not send informational and debug level (6 and 7, respectively) messages because of the sheer number of level 6 and 7 messages generated during operation. Finally, the switch is configured to send log messages with a source address of loopback0. Example 7-17 shows the configuration of logging to a syslog server using hybrid software.

Example 7-17 *Completing the Logging Configuration (Hybrid-Catalyst OS)*

```
SW3> (enable) set logging server 10.10.10.1
10.10.10.1 added to System logging server table.
```

In Example 7-17, the switch is pointed to the same syslog server at 10.10.10.1. Catalyst OS does not support a loopback interface and log messages are sent with a source address of sc0.

By default, syslog messages are not time stamped. This can cause major issues when attempting to troubleshoot the switch because not knowing when the message occurred can sometimes render the messages almost useless. In Example 7-18, a switch running native software is configured for time stamping of syslog messages and system debug messages.

Example 7-18 *Configuring Debug and Log Message Time Stamps (Native)*

```
SW1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#service timestamps debug datetime localtime show-timezone msec
SW1(config)#service timestamps log datetime localtime show-timezone msec
SW1(config)#end
SW1#
```

In Example 7-19, a switch running hybrid software is configured for time stamping of syslog messages and system debug messages.

Example 7-19 *Configuring Log Message Time Stamps (Hybrid-Catalyst OS)*

```
SW3> (enable) set logging timestamp enable
System logging messages timestamp will be enabled.
```

NOTE A discussion of external time sources is beyond the scope of this book. You should reference documentation on the Network Time Protocol (NTP) on Cisco.com, along with publicly available information on the types of time sources that can be purchased for or accessed in networking environments.

Logging levels can be adjusted in both Catalyst OS and Cisco IOS for a wide variety of facilities or features. For example, spanning tree in Catalyst OS defaults to generating log messages for level 2 and higher, but is many times adjusted to level 6 so that more information is recorded during spanning-tree changes. Consult the Cisco web page at Cisco.com for a complete listing of facilities and their default levels for each platform and operating system.

Connecting the Switches

Now that a good portion of the "housekeeping" configuration items are complete, connections between the four switches in Figure 7-4 can be configured. All the physical connections between the switches are already in place.

Figure 7-4 *Physical Links Between Switches*

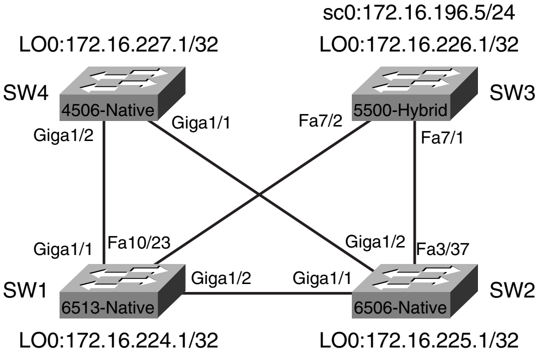

IOS Port/Interface Types

Because a combination of platforms is being used in the examples throughout this chapter, it is important to understand the different types of port/interface types that can be configured on a switch running IOS. Table 7-5 outlines the types of port/interfaces and their uses.

Configuring the Connections

The first connection to bring up is the Gigabit connection between SW1 and SW2. This connection is a single Gigabit link and will not be configured as a trunk, but as a routed physical interface. All interfaces on Catalyst 6000/6500s running native IOS default to routed physical interfaces.

Table 7-5 *Port/Interface Types in IOS*

Port/Interface Type	Function	Sample Configuration
Routed Physical Interface	Traditional Cisco IOS routed interface. Each interface represents a unique Layer 3 network.	**interface gigabitethernet 1/1** **no switchport** **ip address 172.16.100.1 255.255.255.0**
Routed Switch Virtual Interface (SVI)	Single routed interface for all the switchports assigned to a VLAN.	**interface vlan 901** **ip address 172.16.200.1 255.255.255.0**
Access Switch-Port Interface	To group a range of Layer 2 ports into a single VLAN.	**interface range fastethernet 2/1-48** **switchport mode access** **switchport access vlan 130**

In Example 7-20, the current configuration of the GigabitEthernet1/2 interface shows the interface is shut down and no IP address is assigned. In Example 7-20, a /30 IP address is assigned from the range for point-to-point links defined in Table 7-3, earlier in this chapter.

Example 7-20 *Configuring the GigabitEthernet Link on SW1*

```
SW1#show run interface gigabitethernet 1/2
Building configuration...

Current configuration : 61 bytes
!
interface GigabitEthernet1/2
 no ip address
 shutdown
end

SW1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#interface gigabitethernet 1/2
SW1(config-if)#ip address 172.16.240.5 255.255.255.252
SW1(config-if)#no shutdown
SW1(config-if)#end
SW1#
```

In Example 7-21, the GigabitEthernet interface on SW2 is configured.

Example 7-21 *Configuring the GigabitEthernet Link on SW2*

```
SW2#show run interface gig
SW2#show run interface gigabitethernet 1/1
Building configuration...

Current configuration : 61 bytes
```

Example 7-21 *Configuring the GigabitEthernet Link on SW2 (Continued)*

```
!
interface GigabitEthernet1/1
 no ip address
 shutdown
end

SW2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)#interface GigabitEthernet1/1
SW2(config-if)#ip address 172.16.240.6 255.255.255.252
SW2(config-if)#no shut
SW2(config-if)#end
SW2#
1w6d: %SYS-5-CONFIG_I: Configured from console by console
```

In Example 7-22, a **show interface gigabitethernet1/1** command is issued to determine if the interface is now UP/UP, and a **ping** command is issued to the IP address of the GigabitEthernet1/2 interface on SW1 to determine success.

Example 7-22 *Testing the Connection Between SW1 and SW2*

```
SW2#show interface gigabitethernet 1/1
GigabitEthernet1/1 is up, line protocol is up
  Hardware is C6k 1000Mb 802.3, address is 0001.6471.d968 (bia 0001.6471.d968)
  Internet address is 172.16.240.6/30
!output truncated

SW2#ping 172.16.240.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.240.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
SW2#
```

Figure 7-5 shows the network as it is configured at this stage, and the IP addressing information assigned thus far.

Figure 7-5 *Link Operational Between SW1 and SW2*

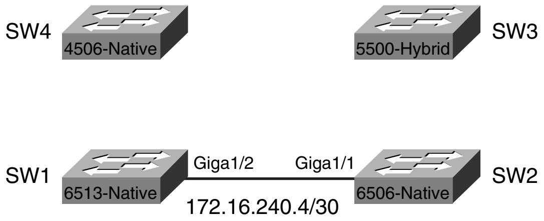

Next, the connection between SW2 and SW3 is configured, as shown in Example 7-23.

Example 7-23 *Configuring the Connection on SW2 to SW3*

```
SW2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)#interface fastEthernet 3/37
SW2(config-if)#ip address 172.16.240.9 255.255.255.252
SW2(config-if)#no shutdown
SW2(config-if)#end
SW2#
```

Because SW3 is running hybrid software, the Layer 2 connection in Catalyst OS is configured first, and then the switched virtual interface (SVI) on the RSM is configured. VLAN 901 is selected from the range of VLANs allocated to point-to-point Layer 2 links. Example 7-24 shows VLAN 901 being created and port 7/1 assigned to VLAN 901. The next step is to configure the RSM with a VLAN 901 interface.

Example 7-24 *Configuring the Connection on SW3 to SW2 (Catalyst OS)*

```
SW3> (enable) set vlan 901
Vlan 901 configuration successful
SW3> (enable) set vlan 901 7/1
VLAN 901 modified.
VLAN 1 modified.
VLAN  Mod/Ports
---- ----------------------
901   7/1

SW3> (enable)
SW3> (enable) show port 7/1
Port  Name               Status     Vlan       Level  Duplex Speed Type
----- ------------------ ---------- ---------- ------ ------ ----- ------------
 7/1                     connected  901        normal a-full a-100 10/100BaseTX
!Output truncated
SW3> (enable)
```

Example 7-25 shows VLAN 901 being configured on the RSM of SW3.

Example 7-25 *Configuring VLAN 901 on the RSM on SW3*

```
RSM1>en
RSM1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
RSM1(config)#interface vlan901
RSM1(config-if)#ip address 172.16.240.10 255.255.255.252
RSM1(config-if)#no shut
RSM1(config-if)#end
RSM1#
```

In Example 7-26, the **show interface vlan901** command is issued to determine if the interface is now UP/UP, and a **ping** is issued to the IP address of the FastEthernet3/37 interface on SW2 to determine success.

Example 7-26 *Testing the Connection Between SW3 and SW2*

```
RSM1#show interface vlan901
Vlan901 is up, line protocol is up
  Hardware is Cat5k Virtual Ethernet, address is 0010.f6b3.4800 (bia 0010.f6
800)
  Internet address is 172.16.240.10/30
!output truncated

RSM1#ping 172.16.240.9

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.240.9, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/9/40 ms
RSM1#
```

Figure 7-6 shows the network as it looks at this stage, and the IP addressing information assigned thus far.

Figure 7-6 *Link Operational Between SW3 and SW2*

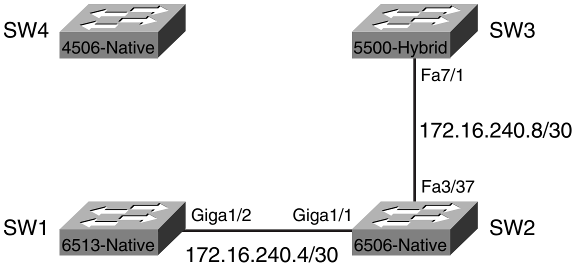

Next, the connection between SW3 and SW1 is configured (see Example 7-27). Again, the Layer 2 connection in Catalyst OS is configured first, followed by the SVI on the RSM. VLAN 902 is used for this link.

Example 7-27 *Configuring the Connection on SW3 to SW1 (Catalyst OS)*

```
SW3> (enable) set vlan 902
Vlan 902 configuration successful
SW3> (enable) set vlan 902 7/2
VLAN 902 modified.
```

continues

Example 7-27 *Configuring the Connection on SW3 to SW1 (Catalyst OS) (Continued)*

```
VLAN 1 modified.
VLAN  Mod/Ports
---- -----------------------
902   7/2

SW3> (enable) show port 7/2
Port  Name               Status     Vlan       Level  Duplex Speed Type
----- ------------------ ---------- ---------- ------ ------ ----- ------------
 7/2                     connected  902        normal a-full a-100 10/100BaseTX
!output truncated
```

The next step is to configure the RSM with a VLAN902 interface, as shown in Example 7-28.

Example 7-28 *Configuring VLAN902 on the RSM on SW3*

```
RSM1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
RSM1(config)#interface VLAN902
RSM1(config-if)#ip address 172.16.240.13 255.255.255.252
RSM1(config-if)#no shutdown
RSM1(config-if)#end
```

Now that the SW3 side of the connection is configured, the other side is configured on SW1 (see Example 7-29).

Example 7-29 *Configuring the Connection Between SW1 and SW3*

```
SW1#show run interface FastEthernet10/23
Building configuration...

Current configuration : 60 bytes
!
interface FastEthernet10/23
 no ip address
 shutdown
end

SW1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#interface FastEthernet10/23
SW1(config-if)#ip address 172.16.240.14 255.255.255.252
SW1(config-if)#no shutdown
SW1(config-if)#end
SW1#
```

In Example 7-30, a **show interface FastEthernet10/23** command is issued to determine if the interface is now UP/UP, and a **ping** is issued to the IP address of the VLAN902 interface on SW3 to determine success.

Example 7-30 *Testing the Connection Between SW1 and SW3*

```
SW1#show interface FastEthernet10/23
FastEthernet10/23 is up, line protocol is up
  Hardware is C6k 100Mb 802.3, address is 0005.7418.048a (bia 0005.7418.048a)
  Internet address is 172.16.240.14/30
!output truncated

SW1#ping 172.16.240.13

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.240.13, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
SW1#
```

Figure 7-7 shows the network as it looks at this stage, and the IP addressing information assigned thus far.

Figure 7-7 *Link Operational Between SW3 and SW1*

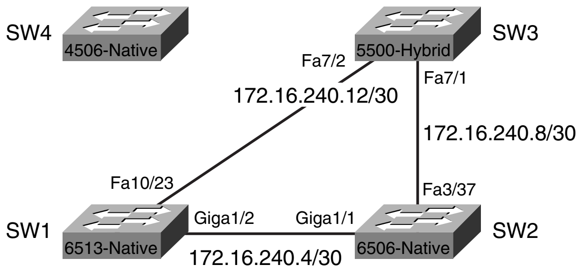

Next, the connection between SW1 and SW4 is configured, as shown in Example 7-31.

Example 7-31 *Configuring the Connection on SW1 to SW4*

```
SW1#show run interface gigabitethernet 1/1
Building configuration...

Current configuration : 61 bytes
```

continues

Example 7-31 *Configuring the Connection on SW1 to SW4 (Continued)*

```
!
interface GigabitEthernet1/1
 no ip address
 shutdown
end

SW1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#interface gigabitethernet 1/1
SW1(config-if)#ip address 172.16.240.17 255.255.255.252
SW1(config-if)#no shutdown
SW1(config-if)#end
SW1#
```

Next, the connection between SW1 and SW4 is configured, as shown in Example 7-32.

Example 7-32 *Configuring the Connection on SW4 to SW1*

```
SW4#show run interface gigabitethernet 1/2
Building configuration...

Current configuration : 36 bytes
!
interface GigabitEthernet1/2
end

SW4#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW4(config)#interface gigabitethernet 1/2
SW4(config-if)#no switchport
SW4(config-if)#ip address 172.16.240.18 255.255.255.252
SW4(config-if)#end
SW4#
```

It is important to understand that the Catalyst 4500 series switch defaults to all interfaces being configured as Access Port Switch Interfaces. To convert the gigabitethernet1/2 interface from a Layer 2 switchport to a Layer 3 routed physical interface, the **no switchport** command must be used prior to assigning the IP address in Example 7-32.

In Example 7-33, the **show interface gigabitethernet1/2** command is issued to determine if the interface is now up/up and a **ping** is issued to the IP address of the GigabitEthernet 1/1 interface on SW1 to determine success.

Figure 7-8 shows the network as it looks at this stage, and the IP addressing information assigned thus far.

Example 7-33 *Testing the Connection Between SW4 and SW1*

```
SW4#show interface gigabitethernet 1/2
GigabitEthernet1/2 is up, line protocol is up (connected)
  Hardware is Gigabit Ethernet Port, address is 000b.fdd5.62bf (bia 000b.fdd5
bf)
  Internet address is 172.16.240.18/30
!output truncated

SW4#ping 172.16.240.17

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.240.17, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
SW4#
```

Figure 7-8 *Link Operational Between SW1 and SW4*

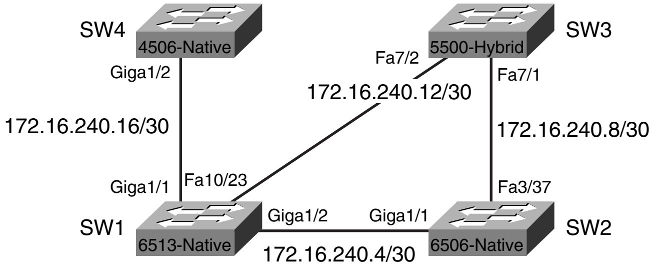

Finally, the connection between SW4 and SW2 is configured, starting with SW4. (See Example 7-34.)

Example 7-34 *Configuring the Connection from SW4 to SW2*

```
SW4#show run interface GigabitEthernet1/1
Building configuration...

Current configuration : 36 bytes
!
interface GigabitEthernet1/1
end

SW4#config t
Enter configuration commands, one per line.  End with CNTL/Z.
```

continues

Example 7-34 *Configuring the Connection from SW4 to SW2 (Continued)*

```
SW4(config)#interface GigabitEthernet1/1
SW4(config-if)#no switchport
SW4(config-if)#ip address 172.16.240.21 255.255.255.252
SW4(config-if)#end
SW4#
```

The connection is completed by configuring the GigabitEthernet1/2 interface on SW2, as shown in Example 7-35.

Example 7-35 *Configuring the Connection from SW2 to SW4*

```
SW2#show run interface GigabitEthernet1/2
Building configuration...

Current configuration : 61 bytes
!
interface GigabitEthernet1/2
 no ip address
 shutdown
end

SW2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)#interface GigabitEthernet1/2
SW2(config-if)#ip address 172.16.240.22 255.255.255.252
SW2(config-if)#no shutdown
SW2(config-if)#end
SW2#
```

In Example 7-36, the **show interface gigabitethernet 1/2** command is issued to determine if the interface is now UP/UP, and a **ping** is issued to the IP address of the GigabitEthernet 1/1 interface on SW4 to determine success.

Example 7-36 *Testing the Connection Between SW2 and SW4*

```
SW2#show interface gigabitethernet 1/2
GigabitEthernet1/2 is up, line protocol is up
  Hardware is C6k 1000Mb 802.3, address is 0001.6471.d969 (bia 0001.6471.d969)
  Internet address is 172.16.240.22/30
!output truncated

SW2#ping 172.16.240.21

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.240.25, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
SW2#
```

Figure 7-9 shows the completed switch connections, and all the network addresses assigned.

Figure 7-9 *Completed Connections Between All Four Switches*

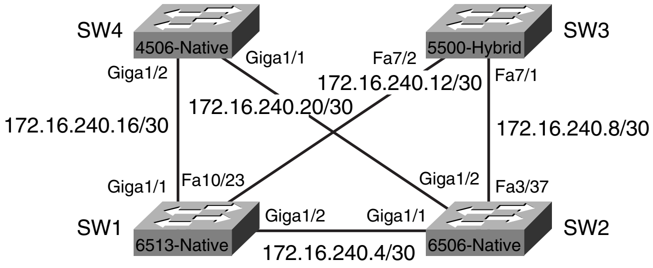

Now that the basic connections between switches have been established, ports to be used for access layer devices, such as workstations and servers, will be configured.

Configuring the Access Layer

Switchports on the Catalyst 5500 SW3 and interfaces on the Catalyst 4506 SW4 will be configured in VLANs to support access layer devices. Figure 7-10 shows the IP network numbers assigned to these VLANs.

Figure 7-10 *Addition of Access Layer VLANs*

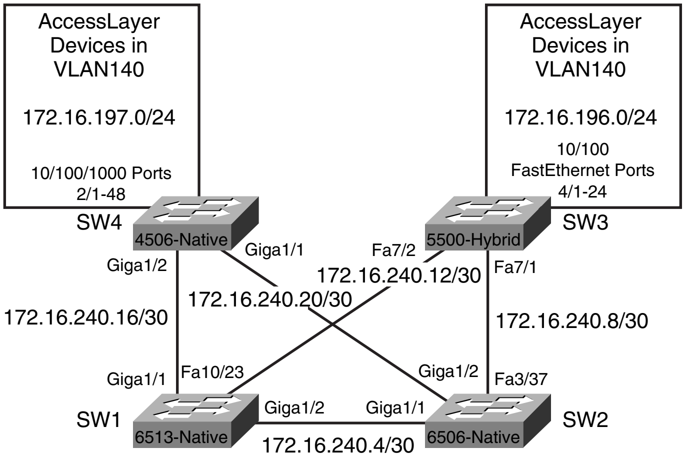

Configuring the access layer begins with configuring ports on SW3 to be in VLAN 130. VLAN 130 is one of the access layer VLANs in the VLAN addressing scheme outlined in Table 7-1 earlier in this chapter. Remember VLAN 130 was created on SW3 earlier in VTP configuration (refer to Example 7-3). In Example 7-37, module 4 on SW3 is a 24-port 10/100 Mb FastEthernet module, and will have all ports assigned to VLAN 130.

Example 7-37 *Configuring Ports on SW3 as Members of VLAN 130*

```
SW3> (enable) show mod 4
Mod Module-Name         Ports Module-Type           Model     Serial-Num Status
--- ------------------- ----- --------------------- --------- --------- -------
4                       24    10/100BaseTX Ethernet WS-X5224  009607843 ok

Mod MAC-Address(es)                       Hw     Fw         Sw
--- -------------------------------------- ------ ---------- -----------------
4   00-10-7b-78-57-00 to 00-10-7b-78-57-17 1.4    3.1(1)     4.5(5)
SW3> (enable)

SW3> (enable) set vlan 130 4/1-24
VLAN 130 modified.
VLAN 1 modified.
VLAN  Mod/Ports
---- -----------------------
130   4/1-24

SW3> (enable)
```

For these ports to be reachable from other networks, an SVI must be configured on the RSM for VLAN 130. The SVI for VLAN 130 is configured in Example 7-38. Remember sc0 on the switch is already assigned to VLAN 130 with an IP address of 172.16.196.5/24 in Example 7-8, earlier in the chapter.

Example 7-38 *Configuring a SVI for VLAN 130 on the RSM of SW3*

```
RSM1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
RSM1(config)#int vlan130
RSM1(config-if)#ip address 172.16.196.1 255.255.255.0
RSM1(config-if)#end
RSM1#
```

In Example 7-39, the **show interface vlan130** command is issued to confirm the SVI is UP/UP, and a **ping** from the SVI to the sc0 interface on the supervisor is issued.

In Example 7-40, the interfaces on module 2 of SW4 are configured for VLAN 140. Module 4 on the SW4 is a 48-port 10/100/1000BASE-TX module.

Example 7-39 *Verifying the Status of the VLAN130 Interface and sc0*

```
RSM1#show interface vlan130
Vlan130 is up, line protocol is up
  Hardware is Cat5k Virtual Ethernet, address is 0010.f6b3.4800 (bia 0010.f6b3.4
800)
  Internet address is 172.16.196.1/24
(output truncated)

RSM1#ping 172.16.196.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.196.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/23/112 ms
```

Example 7-40 *Configuring Ports 2/1-48 on SW4 for VLAN 140*

```
SW4#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW4(config)#interface range gigabitethernet 2/1 - 48
SW4(config-if-range)#switchport mode access
SW4(config-if-range)#switchport access vlan 140
SW4(config-if-range)#end
SW4#

SW4#show vlan

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active
140  VLAN0140                         active    Gi2/1, Gi2/2, Gi2/3, Gi2/4
                                                Gi2/5, Gi2/6, Gi2/7, Gi2/8
                                                Gi2/9, Gi2/10, Gi2/11, Gi2/12
                                                Gi2/13, Gi2/14, Gi2/15, Gi2/16
                                                Gi2/17, Gi2/18, Gi2/19, Gi2/20
                                                Gi2/21, Gi2/22, Gi2/23, Gi2/24
                                                Gi2/25, Gi2/26, Gi2/27, Gi2/28
                                                Gi2/29, Gi2/30, Gi2/31, Gi2/32
                                                Gi2/33, Gi2/34, Gi2/35, Gi2/36
                                                Gi2/37, Gi2/38, Gi2/39, Gi2/40
                                                Gi2/41, Gi2/42, Gi2/43, Gi2/44
                                                Gi2/45, Gi2/46, Gi2/47, Gi2/48
!output truncated
```

The **interface range** command must be entered exactly as shown in Example 7-40 with spaces to be accepted. The output of the **show vlan** command shows ports 2/1–48 assigned successfully to VLAN 140.

Now that these ports have been assigned, an SVI must be created on SW4 so that VLAN 140 can be reached from other networks. The SVI for VLAN 140 on SW4 is created in Example 7-41.

Example 7-41 *Configuration of a SVI on SW4 for VLAN 140*

```
SW4#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW4(config)#interface VLAN140
SW4(config-if)#ip address 172.16.197.1 255.255.255.0
SW4(config-if)#no shutdown
SW4(config-if)#end
SW4#
```

In Example 7-42, the **show interface vlan140** command is issued to confirm the SVI is UP/UP.

Example 7-42 *Verifying the Status of the VLAN140 Interface*

```
SW4#show interface vlan140
Vlan140 is up, line protocol is up
  Hardware is Ethernet SVI, address is 000b.fdd5.62bf (bia 000b.fdd5.62bf)
  Internet address is 172.16.197.1/24
```

Dynamic Routing

Now that the Layer 3 connections between the four switches are configured, the access layer VLANs created, and access ports assigned, a dynamic routing protocol is configured to allow connectivity between VLANs. In these examples, EIGRP is used as the dynamic routing protocol. EIGRP will be enabled on all four switches using Autonomous System (AS) 100, starting with SW1. Refer to the documentation on Cisco.com for more information about EIGRP and other dynamic routing protocols. Example 7-43 shows EIGRP being configured on SW1.

Example 7-43 *EIGRP Configured on SW1*

```
SW1(config)#router eigrp 100
SW1(config-router)#network 172.16.192.0 0.0.63.255
SW1(config-router)#end
SW1#show ip eigrp interfaces
IP-EIGRP interfaces for process 100

                    Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface    Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Gi1/1          0        0/0         0       0/10           0           0
Gi1/2          0        0/0         0       0/10           0           0
Fa10/23        0        0/0         0       0/10           0           0
Lo0            0        0/0         0       0/10           0           0
SW1#
```

The output of the **show ip eigrp interfaces** command in Example 7-43 indicates the four interfaces that have been configured on SW1 with IP addresses in the previous exercises now part of EIGRP AS 100. The same commands are repeated on SW2 in Example 7-44, on SW3 in Example 7-45, and on SW4 in Example 7-46.

Example 7-44 *EIGRP Configured on SW2*

```
SW2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)#router eigrp 100
SW2(config-router)#network 172.16.192.0 0.0.63.255
SW2(config-router)#end
SW2#show ip eigrp interfaces
IP-EIGRP interfaces for process 100

                    Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface    Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Gi1/1          1        0/0      1044       0/10         5216           0
Gi1/2          0        0/0         0       0/10            0           0
Fa3/37         0        0/0         0       0/10            0           0
Lo0            0        0/0         0       0/10            0           0
SW2#
```

Example 7-45 *EIGRP Configured on the RSM of SW3*

```
RSM1(config)#router eigrp 100
RSM1(config-router)#network 172.16.192.0 0.0.63.255
RSM1(config-router)#end
RSM1#show ip eigrp interfaces
IP-EIGRP interfaces for process 100

                    Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface    Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Vl130          0        0/0         0       0/10            0           0
Vl901          1        0/0       726       0/10         3632           0
Vl902          1        0/0       752       0/10         3760           0
Lo0            0        0/0         0       0/10            0           0
RSM1#
```

Example 7-46 *EIGRP Configured on SW4*

```
SW4#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW4(config)#router eigrp 100
SW4(config-router)#network 172.16.192.0 0.0.63.255
SW4(config-router)#end
SW4#show ip eigrp interfaces
IP-EIGRP interfaces for process 100
```

continues

Example 7-46 *EIGRP Configured on SW4 (Continued)*

```
                    Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface    Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Vl140          0        0/0         0       0/10           0           0
Gi1/1          1        0/0         0       0/10           0           0
Gi1/2          1        0/0         0       0/10           0           0
Lo0            0        0/0         0       0/10           0           0
SW4#
```

Now that dynamic routing for network 172.16.192.0 and its subnets has been configured on all four switches, a look at the routing table of SW1 in Example 7-47 shows that the networks for the access layer VLANs (172.16.196.0 and 172.16.197.0) are now reachable via the uplinks to those switches.

Example 7-47 *Output of* **show ip route** *on SW1*

```
SW1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     172.16.0.0/16 is variably subnetted, 11 subnets, 3 masks
C       172.16.240.12/30 is directly connected, FastEthernet10/23
D       172.16.240.8/30
           [90/28416] via 172.16.240.6, 00:03:08, GigabitEthernet1/2
C       172.16.240.4/30 is directly connected, GigabitEthernet1/2
D       172.16.240.20/30
           [90/3072] via 172.16.240.6, 00:03:11, GigabitEthernet1/2
           [90/3072] via 172.16.240.18, 00:03:11, GigabitEthernet1/1
D       172.16.225.1/32
           [90/130816] via 172.16.240.6, 00:03:08, GigabitEthernet1/2
C       172.16.240.16/30 is directly connected, GigabitEthernet1/1
C       172.16.224.1/32 is directly connected, Loopback0
D       172.16.227.1/32
           [90/130816] via 172.16.240.18, 00:03:11, GigabitEthernet1/1
D       172.16.226.1/32
           [90/156160] via 172.16.240.13, 00:06:00, FastEthernet10/23
D       172.16.196.0/24
           [90/30720] via 172.16.240.13, 00:06:00, FastEthernet10/23
D       172.16.197.0/24
           [90/3072] via 172.16.240.18, 00:03:12, GigabitEthernet1/1
C    127.0.0.0/8 is directly connected, EOBC0/0
SW1#
```

Channeling and Trunking

All the configuration examples thus far have assumed it is possible to restrict a VLAN to a single switch. Although this is the cleanest and simplest configuration, it is not always possible. Many times, connections between access layer and distribution layer switches are Layer 2, and VLANs must span many switches because of application or administrative requirements. Chapter 11, "Design and Implementation Best Practices," discusses additional design options and considerations.

In Figure 7-11, a requirement for two additional VLANs with ports on both SW1 and SW2 is introduced. VLANs 401 and 402 are used for the exercises. VLANs 401 and 402 have been created on SW1 and SW2 using the same procedures as in Examples 7-1 through 7-4. While the Gigabit connection between SW1 and SW2 could be converted to a trunk to carry these additional VLANs, some unused FastEthernet ports will be configured in a channel to carry only these new VLANs and VLAN 1.

Figure 7-11 *Addition of VLANs 401 and 402*

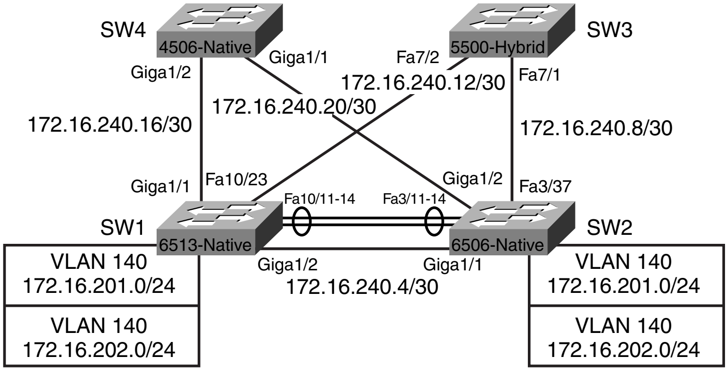

Configuration begins with creating the channel group on SW1, as shown in Example 7-48.

Example 7-48 *Creating the Channel Group on SW1*

```
SW1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#interface range fastEthernet 10/11 - 14
SW1(config-if-range)#no ip address
SW1(config-if-range)#switchport
SW1(config-if-range)#switchport trunk encapsulation dot1q
SW1(config-if-range)#switchport mode trunk
SW1(config-if-range)#switchport mode dynamic desirable
SW1(config-if-range)#switchport trunk allowed vlan remove 2-400,403-1005
```

continues

Example 7-48 *Creating the Channel Group on SW1 (Continued)*

```
SW1(config-if-range)#channel-group 1 mode desirable
SW1(config-if-range)#no shutdown
SW1(config-if-range)#end
SW1#
```

The channel is completed by configuring the other side on SW2, as shown in Example 7-49.

Example 7-49 *Creating the Channel Group on SW2*

```
SW2(config)#interface range fastEthernet 3/11 - 14
SW2(config-if-range)#no ip address
SW2(config-if-range)#switchport

SW2(config-if-range)#switchport trunk encapsulation dot1q
SW2(config-if-range)#switchport mode trunk
SW2(config-if-range)#switchport mode dynamic desirable
SW2(config-if-range)#switchport trunk allowed vlan remove 2-400,403-1005
SW2(config-if-range)#channel-group 1 mode desirable
Creating a port-channel interface Port-channel1
SW2(config-if-range)#no shutdown
SW2(config-if-range)#end
SW2#
```

Issuing a **show run interface fastEthernet 3/11** command displays the configuration of one of the ports in the channel (see Example 7-50).

Example 7-50 *Verifying the Configuration on SW2*

```
SW2#show run interface fastEthernet 3/11
Building configuration...

Current configuration : 182 bytes
!
interface FastEthernet3/11
 no ip address
switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 1,401,402
 channel-group 1 mode desirable
end
```

The operation of the new channel group can be verified by issuing a **show interfaces port-channel 1** command. The operation of the trunk can be verified by issuing the **show interfaces trunk** command, as shown in Example 7-51.

Example 7-51 *Output of the* **show interfaces port-channel** *and* **show interfaces trunk** *Commands on SW1*

```
SW1#show interfaces port-channel 1
Port-channel1 is up, line protocol is up
  Hardware is EtherChannel, address is 0009.1267.9ffa (bia 0009.1267.9ffa)
  MTU 1500 bytes, BW 400000 Kbit, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Full-duplex, 100Mb/s
  Members in this channel: Fa10/11 Fa10/12 Fa10/13 Fa10/14
!output truncated

SW1#show interfaces trunk

Port      Mode         Encapsulation  Status        Native vlan
Po1       desirable    802.1q         trunking      1

Port      Vlans allowed on trunk
Po1       1,401-402

Port      Vlans allowed and active in management domain
Po1       1,401-402

Port      Vlans in spanning tree forwarding state and not pruned
Po1       1,401-402
SW1#
```

The bandwidth reported on the channel is 400000 Kbit, and the members of the channel are listed in the output.

Configuring UniDirectional Link Detection

One best practice to follow when configuring a network like the one used in this chapter is the configuration of UniDirectional Link Detection (UDLD) in Aggressive mode. UDLD is designed to mitigate certain fault conditions on fiber and copper Ethernet interfaces. UDLD is designed to shutdown any miswired ports or unidirectional links by putting the port in an errDisabled state. UDLD is a Layer 2 protocol and, when run in combination with autonegotiation Layer 1 mechanisms, UDLD can validate the physical (Layer 1) and logical (Layer 2) integrity of a link. UDLD accomplishes this task by learning about neighbors and keeping neighbor status in a cache. Neighbors are learned by the sending of UDLD echo or hello messages.

The UDLD Aggressive feature provides additional protection against unidirectional link conditions in certain situations, and attempts to re-establish a connection with the neighbor when a failure is detected. UDLD Aggressive works by detecting when one side of a link remains up while the other side of the link has gone down, and after eight failed retries, transitions the port to an errDisabled state and generates a syslog message.

Cisco recommends configuring UDLD in Aggressive mode on point-to-point FastEthernet/GigabitEthernet links between Cisco switches, and setting the message interval to 15 seconds. UDLD is globally disabled by default and can be enabled globally or on a port–by-port basis. In the examples in this section, UDLD Aggressive should be configured on all the links between switches. An example of this configuration on a per-port basis is shown in Example 7-52 using SW1 and SW2.

Example 7-52 *Enabling Aggressive UDLD on SW1 and SW2*

```
SW1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW1(config)#interface range gigabitethernet 1/1 - 2
SW1(config-if-range)#udld enable
SW1(config-if-range)#udld aggressive
SW1(config-if-range)#end

SW2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW2(config)#interface range gigabitethernet 1/1 - 2
SW2(config-if-range)#udld enable
SW2(config-if-range)#udld aggressive
SW2(config-if-range)#end
SW2#
```

The output of the **show udld** command on SW1 shows the status of the UDLD configuration. In the output in Example 7-53, SW1 detects SW2 as a UDLD neighbor, because both SW1 and SW2 have been configured, but does not detect SW4 on GigabitEthernet1/1 because it has yet to be configured.

Example 7-53 *Output of* **show udld** *Command on SW1*

```
SW1#show udld

Interface Gi1/1
---
Port enable administrative configuration setting: Enabled / in aggressive mode
Port enable operational state: Enabled / in aggressive mode
Current bidirectional state: Unknown
Current operational state: Advertisement
Message interval: 7
Time out interval: 5
No neighbor cache information stored

Interface Gi1/2
---
Port enable administrative configuration setting: Enabled / in aggressive mode
Port enable operational state: Enabled / in aggressive mode
Current bidirectional state: Bidirectional
Current operational state: Advertisement - Single neighbor detected
Message interval: 60
```

Example 7-53 *Output of* **show udld** *Command on SW1 (Continued)*

```
Time out interval: 5

    Entry 1
    ---
    Expiration time: 168
    Device ID: 1
    Current neighbor state: Bidirectional
    Device name: SAD04281ARM
    Port ID: Gi1/1
    Neighbor echo 1 device: SAD050814BH
    Neighbor echo 1 port: Gi1/2

    Message interval: 5
    CDP Device name: SW2
```

Portfast and BPDU Guard

You can find a detailed discussion of portfast and BPDU Guard in Chapter 10, "Implementing and Tuning Spanning Tree," but the configuration of the access layer ports in this chapter's examples would not be complete without enabling portfast and BPDU Guard.

Portfast is a feature that bypasses the normal spanning-tree operation of listening and learning and places a port immediately into forwarding when a port is connected. Portfast should only be used on ports connecting to end-station devices such as workstations and servers. Portfast is disabled by default and is enabled on a port-by-port basis.

The addition of BPDU Guard as an additional protection allows the switch to place any port configured with portfast into an errDisabled state if a BPDU is received on that port. Because ports 2/1 through 2/48 on SW4 were configured for access layer devices in VLAN 140 in Example 7-40 earlier in the chapter, those ports will have portfast and BPDU Guard enabled as follows in Example 7-54.

Example 7-54 *Enabling Portfast and BPDU Guard on SW4*

```
SW4#config t
Enter configuration commands, one per line.  End with CNTL/Z.
SW4(config)#interface range gigabitethernet 2/1 - 48
SW4(config-if-range)#spanning-tree portfast
%Warning: portfast should only be enabled on ports connected to a single
 host. Connecting hubs, concentrators, switches, bridges, etc... to this
 interface  when portfast is enabled, can cause temporary bridging loops.
 Use with CAUTION
%Portfast will be configured in 48 interfaces due to the range command
 but will only have effect when the interfaces are in a non-trunking mode.

SW4(config-if-range)#spanning-tree bpduguard enable
SW4(config-if-range)#end
```

Configuring SNMP

Now that the sample network in this chapter has been configured and is operational, the switches should be configured so that management stations can gather information via the Simple Network Management Protocol (SNMP). SNMP is used to gather statistics, counters, and tables in the Management Information Base (MIB) of a device.

The SNMP framework consists of three parts:

- SNMP manager
- SNMP agent
- MIB

The SNMP manager is a host that monitors the activities of network devices using SNMP. The SNMP manager is typically referred to as the Network Management Station (NMS). The SNMP agent is software running on the device being monitored by the SNMP manager. A MIB is a virtual storage area for network management information consisting of collections of managed objects. MIBs are written in the SNMP MIB module language as defined in RFCs 2578, 2579, and 2580. SNMP agents can be configured to allow read-only or read-write access to the device. Management stations like CiscoWorks use the read-only functions of the agent to monitor the device, and can use the read-write functions of the agent to make changes to the device configuration. SNMP uses passwords called *community strings* to grant access to the SNMP agent. Access to the agents can be further limited via SNMP access lists.

A detailed discussion of SNMP and network management is beyond the scope of this book. SNMP MIB information for each Cisco device can be found on Cisco.com.

You should configure SNMP on each Cisco device to be monitored by NMS. A sample SNMP configuration is shown on SW1 in Example 7-55.

Example 7-55 *Sample SNMP Configuration on SW1*

```
SW1#config t
Enter configuration commands, one per line.  End with CNTL/Z.

SW1(config)#access-list 10 permit 10.10.10.2
SW1(config)#snmp enable
SW1(config)#snmp-server community alpha ro 10
SW1(config)#snmp-server community beta rw 10
SW1(config)#snmp-server contact John Smith (555)789-2653
SW1(config)#snmp-server location Denver Data Center
SW1(config)#snmp enable traps
SW1(config)#snmp trap-source lo0
```

After SNMP is configured on SW1, the statistics can be viewed using the **show snmp** command, as shown in Example 7-56.

Example 7-56 *Output of* **show snmp** *Command on SW1*

```
SW1#show snmp
Chassis: SAD050814BH
Contact: John Smith (555)789-2653
Location: Denver Data Center
0 SNMP packets input
    0 Bad SNMP version errors
    0 Unknown community name
    0 Illegal operation for community name supplied
    0 Encoding errors
    0 Number of requested variables
    0 Number of altered variables
    0 Get-request PDUs
    0 Get-next PDUs
    0 Set-request PDUs
0 SNMP packets output
    0 Too big errors (Maximum packet size 1500)
    0 No such name errors
    0 Bad values errors
    0 General errors
    0 Response PDUs
    0 Trap PDUs

SNMP logging: disabled
```

Summary

This chapter provided configuration examples for connecting switches running both native and hybrid software configurations, and offers examples of both Layer 3 and Layer 2 connectivity between switches. The material in this chapter can serve as a reference for initial configuration of either hybrid or native switches, and illustrates some options and best practices when configuring access layer ports, trunks, channels, and point-to-point Layer 3 links. In each of the examples, no logical Layer 2 loops were created, so spanning-tree issues are avoided. Avoiding spanning tree is not always possible, so you should reference Chapter 10 of this book when attempting to optimize and troubleshoot a network with logical loops.

A detailed look at other design options and implementation best practices can be found in Chapter 11.