

VLANs and Trunking

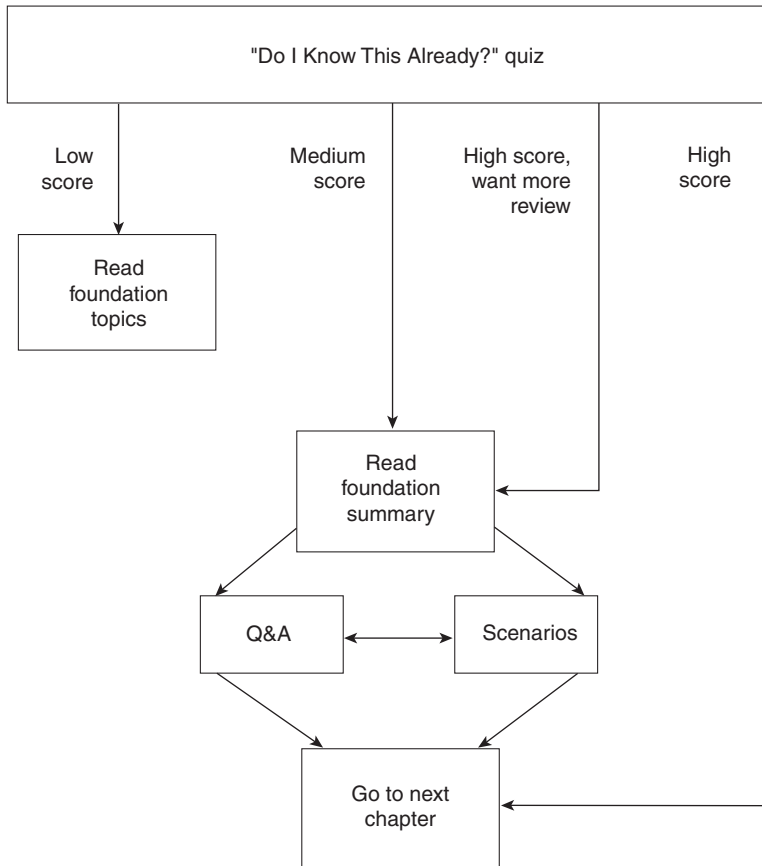
Switched campus networks can be broken up into distinct broadcast domains or virtual LANs (VLANs). A flat network topology, or a network with a single broadcast domain, can be simple to implement and manage. However, flat network topology is not scalable. Instead, the campus can be divided into segments using VLANs, while Layer 3 routing protocols manage interVLAN communication.

This chapter details the process of defining common workgroups within a group of switches. Switch configuration for VLANs is covered, along with the method of identifying and transporting VLANs on various types of links. VLAN administration and management is presented through the configuration of the VLAN Trunking Protocol (VTP).

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and answers for all your work with this book in one place, for easy reference.
- Take the “Do I Know This Already?” quiz and write down facts and concepts (even if you never look at the information again).
- Use the diagram in Figure 4-1 to guide you to the next step.

Figure 4-1 *How to Use This Chapter*

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The quiz helps you make good choices of how to spend your limited study time. The quiz is sectioned into five smaller “quizlets” that correspond to the five major headings in the “Foundation Topics” section of the chapter. Although your answer may differ somewhat from the answers given, finding out if you have the basic understanding that is presented in this chapter is most important. You will find that these questions are open-ended rather than multiple choice as found on the exams. This is done to focus more on understanding the subject matter than on memorizing details.

Use the scoresheet in Table 4-1 to record your score.

Table 4-1 *Scoresheet for Quiz and Quizlets*

Quizlet Number	Foundation Topics Sections Covering These Questions	Questions	Score
1	Virtual LANs	1–4	
2	VLAN Trunks VLAN Trunk Configuration	5–7	
3	VTP VTP Configuration	8–10	
4	VTP Pruning	11–12	
All questions		1–12	

1 What is a VLAN? When is it used?

2 What are two types of VLANs, in terms of spanning areas of the campus network?

3 Generally speaking, what must be configured (both switch and end user device) for a port-based VLAN?

4 What are the components of a Token Ring VLAN?

5 What is a trunk link?

6 What methods of VLAN frame identification can be used on a Catalyst switch?

7 What is the purpose of Dynamic Trunking Protocol (DTP)?

8 What VTP modes can a Catalyst switch be configured for? Can VLANs be created in each of the modes?

9 How many VTP management domains can a Catalyst switch participate in? How many VTP servers can a management domain have?

10 What conditions must exist for two Catalyst switches to be in the same VTP management domains?

11 What is the purpose of VTP pruning?

12 Which VLAN numbers are never eligible for VTP pruning? Why?

The answers to the quiz are found in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections,” on page 477. The suggested choices for your next step are as follows:

- **6 or less overall score**—Read the entire chapter. This reading includes the “Foundation Topics” and “Foundation Summary” sections, the Q&A section, and the scenarios at the end of the chapter.
- **7–9 overall score**—Begin with the “Foundation Summary” section and then follow with the Q&A section and the scenarios at the end of the chapter.
- **10 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” section and then go to the Q&A section and the scenarios at the end of the chapter. Otherwise, move on to the next chapter.

Foundation Topics

Virtual LANs

Consider a network design that consists of Layer 2 devices only. For example, this design could be a single Ethernet segment, an Ethernet switch with many ports, or a network with several interconnected Ethernet switches. A fully Layer 2 switched network is referred to as a *flat network topology*. A flat network is a single broadcast domain, such that every connected device sees every broadcast packet that is transmitted. As the number of stations on the network increases, so does the number of broadcasts.

Due to the Layer 2 foundation, flat networks cannot contain redundant paths for load balancing or fault tolerance. The reason for this is explained in Chapter 5, “Redundant Switch Links.” To gain any advantage from additional paths to a destination, Layer 3 routing functions must be introduced.

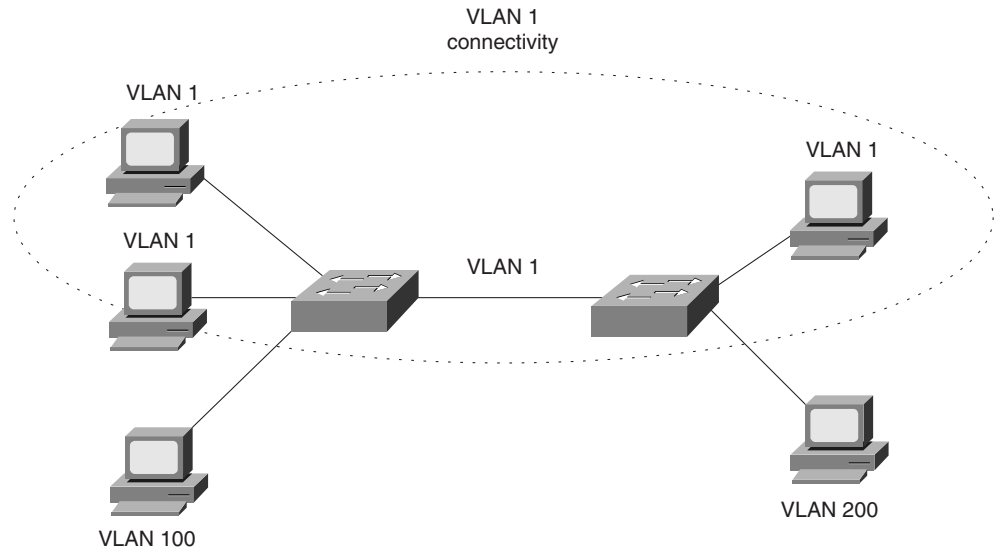
A switched environment offers the technology to overcome flat network limitations. Switched networks can be subdivided into virtual LANs (VLANs). By definition, a VLAN is a single broadcast domain. All devices connected to the VLAN receive broadcasts from other VLAN members. However, devices connected to a different VLAN will not receive those same broadcasts.

A VLAN is made up of defined members communicating as a logical network segment. In contrast, a physical segment consists of devices that must be connected to a physical cable segment. A VLAN can have connected members located anywhere in the campus network, as long as VLAN connectivity is provided between all members. Layer 2 switches are configured with a VLAN mapping and provide the logical connectivity between the VLAN members.

Figure 4-2 shows how a VLAN can provide logical connectivity between switch ports.

Two workstations on the left Catalyst switch are assigned to VLAN 1, while a third workstation is assigned to VLAN 100. In this example, there can be no communication between VLAN 1 and VLAN 100. Both ends of the link between the Catalysts are assigned to VLAN 1. One workstation on the right Catalyst is also assigned to VLAN 1. Because there is end-to-end connectivity of VLAN 1, any of the workstations on VLAN 1 can communicate as if they were connected to a physical network segment.

Figure 4-2 VLAN Functionality



VLAN Membership

When a VLAN is provided at an access layer switch, an end user must have some means to gain membership to it. Two membership methods exist on Cisco Catalyst switches: static VLANs and dynamic VLANs.

Static VLANs

Static VLANs offer *port-based* membership, where switch ports are assigned to specific VLANs. End user devices become members in a VLAN based on which physical switch port they are connected to. No handshaking or unique VLAN membership protocol is needed for the end devices; they automatically assume VLAN connectivity when they connect to a port. Normally, the end device is not even aware that the VLAN exists. The switch port and its VLAN are simply viewed and used as any other network segment, with other “locally attached” members on the wire.

Switch ports are assigned to VLANs by the manual intervention of the network administrator, hence the static nature. The ports on a single switch can be assigned and grouped into many VLANs. Even though two devices are connected to the same switch, traffic will not pass between them if they are connected to ports on different VLANs. To perform this function, either a Layer 3 device could be used to route packets or an external Layer 2 device could be used to bridge packets between the two VLANs.

The static port-to-VLAN membership is normally handled in hardware with application-specific integrated circuits (ASICs) in the switch. This membership provides good performance because all port mappings are done at the hardware level with no complex table lookups needed.

Configuring Static VLANs

This section describes the switch commands needed to configure static VLANs. By default, all switch ports are assigned to VLAN 1, are set to be a VLAN type of Ethernet, have a maximum transmission unit (MTU) size of 1500 bytes, and have a Security Association Identifier (SAID) of 100,000 plus the VLAN number.

First, the VLAN must be created on the switch, if it doesn't already exist. Then the VLAN must be assigned to specific switch ports.

NOTE

To create a new VLAN, several prerequisites relating to VTP must be met. The switch must be assigned to a VTP domain and be configured for either *server* or *transparent* VTP mode. VTP is covered in the “VLAN Trunking Protocol” section of this chapter.

To configure static VLANs on an IOS-based switch, you would enter the following commands in enable mode:

```
Switch# vlan database
Switch(vlan)# vlan vlan-num name vlan-name
Switch(vlan)# exit
Switch# configure terminal
Switch(config)# interface interface module/number
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan vlan-num
Switch(config-if)# end
```

The VLAN is created and stored in a database, along with its number and name. To assign a switch port to the VLAN, you would use the **switchport access vlan** interface configuration command. The **switchport mode access** command configures the port for static VLAN membership.

To configure static VLANs on a CLI-based switch, you would enter the following commands in enable mode:

```
Switch(enable) set vlan vlan-num [name name]
Switch(enable) set vlan vlan-num mod-num/port-list
```

The first command creates the VLAN numbered *vlan-num* on the switch and assigns a descriptive name to it. Note that a VLAN and its number are significant only on the local switch, unless some form of VLAN trunking is used to communicate with other switches. If the *name* field is not specified, the switch will create a name based on the VLAN number, in the form of

VLAN0002 for VLAN 2 for example. The second command assigns VLAN *vlan-num* to one or more switch ports, identified with the switch module number and the list of port numbers. For example, the command **set vlan 101 3/1,3-7** would assign ports 3/1, 3/3, 3/4, 3/5, 3/6, and 3/7 to VLAN 101.

To verify VLAN configuration, using the **show vlan** command will output a list of all VLANs defined in the switch, in addition to the ports assigned to each VLAN.

Dynamic VLANs

Dynamic VLANs are used to provide membership based on the MAC address of an end user device. When a device is connected to a switch port, the switch must query a database to establish VLAN membership. A network administrator must assign the user's MAC address to a VLAN in the database of a VLAN Membership Policy Server (VMPS).

With Cisco switches, dynamic VLANs are created and managed through the use of network management tools like CiscoWorks 2000 or CiscoWorks for Switched Internetworks (CWSI). Dynamic VLANs allow a great deal of flexibility and mobility for end users, but require more administrative overhead.

NOTE

Dynamic VLANs are not covered in this text. For more information, refer to the following Cisco resources:

- CLI-based switches: “Configuring Dynamic Port VLAN Membership with VMPS” at www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/re1_5_5/sw_cfg/vmps.htm
 - IOS-based switches: “How VMPS Works” at www.cisco.com/univercd/cc/td/doc/product/lan/c2900x1/29_35xu/scg/kivlan.htm#xtocid2442355
-

Extent of VLANs

To implement VLANs, you must give some consideration to the number of VLANs you'll need and how best to place them. As usual, the number of VLANs will be dependent on traffic patterns, application types, segmenting common workgroups, and network management requirements.

However, an important factor to consider is the relationship between VLANs and the IP addressing schemes used. Cisco recommends a one-to-one correspondence between VLANs and IP subnets. This recommendation means that if a Class C network address is used for a VLAN, then no more than 254 devices should be in the VLAN. As well, VLANs not extending beyond the Layer 2 domain of the distribution switch is recommended. In other words, the VLAN should not reach across the core of a network and into another switch block. The idea again is to keep broadcasts and unnecessary movement of traffic out of the core block.

VLANs can be scaled in the switch block by using two basic methods: end-to-end VLANs and local VLANs.

End-to-End VLANs

End-to-end VLANs, also called campus-wide VLANs, span the entire switch fabric of a network. They are positioned to support maximum flexibility and mobility of end devices. Users are assigned to VLANs regardless of physical location. As a user moves around the campus, that user's VLAN membership stays the same. This means that each VLAN must be made available at the access layer in every switch block.

End-to-end VLANs should group users according to common requirements. All users in a VLAN should have roughly the same traffic flow patterns, following the 80/20 rule. Recall that this rule estimates that 80 percent of user traffic stays within the local workgroup, while 20 percent is destined for a remote resource in the campus network. Although only 20 percent of the traffic in a VLAN is expected to cross the network core, end-to-end VLANs make it possible for *all* traffic within a single VLAN to cross the core.

Because all VLANs must be available at each access layer switch, VLAN trunking must be used to carry all VLANs between the access and distribution layer switches. (Trunking is discussed in later sections of this chapter.)

Local VLANs

Because most enterprise networks have moved toward the 20/80 rule (where server and intranet/Internet resources are centralized), end-to-end VLANs have become cumbersome and difficult to maintain. The 20/80 rule is reversed—only 20 percent of traffic is local, while 80 percent is destined to a remote resource across the core layer. End users require access to central resources outside their VLAN. Users must cross into the network core more frequently. In this type of network, VLANs are designed to contain user communities based on geographic boundaries, with little regard to the amount of traffic leaving the VLAN.

Local or geographic VLANs range in size from a single switch in a wiring closet to an entire building. Arranging VLANs in this fashion enables the Layer 3 function in the campus network to intelligently handle the inter-VLAN traffic loads. This scenario provides maximum availability by using multiple paths to destinations, maximum scalability by keeping the VLAN within a switch block, and maximum manageability.

VLAN Trunks

At the access layer, end user devices connect to switch ports that provide simple connectivity to a single VLAN each. The attached devices are unaware of any VLAN structure and simply attach to what appears to be a normal physical network segment. Remember, sending information from an access link on one VLAN to another VLAN is not possible without the intervention of an additional device—either a Layer 3 router or an external Layer 2 bridge.

NOTE

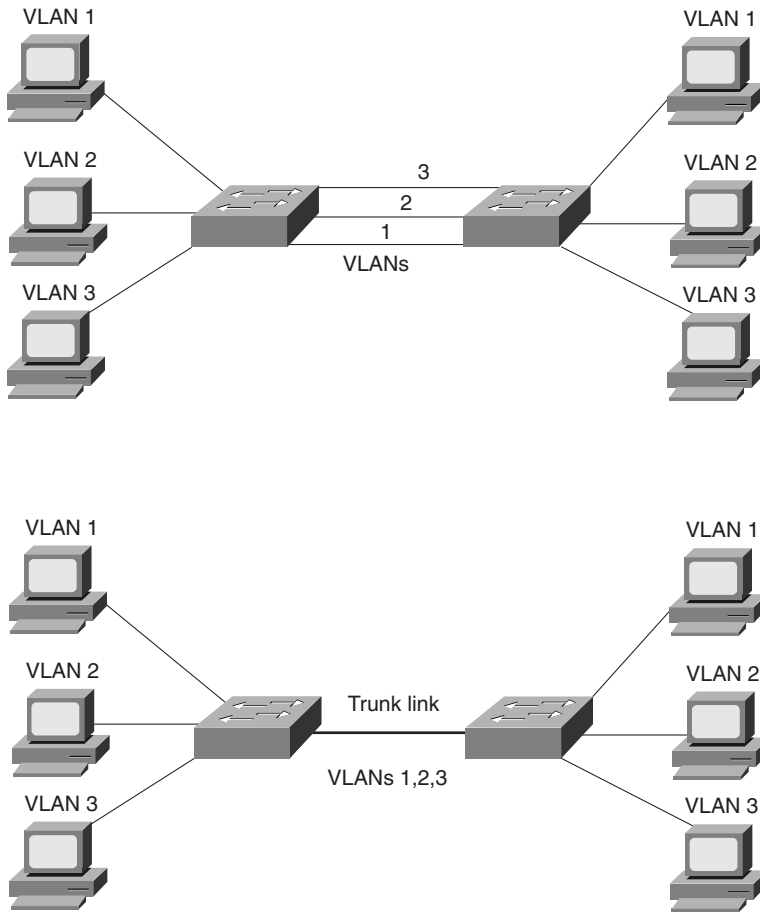
Note that a switch port can support more than one IP subnet for the devices attached to it. For example, consider a shared Ethernet hub that is connected to a single Ethernet switch port. One user device on the hub may be configured for 192.168.1.1 255.255.255.0, while another is assigned 192.168.17.1 255.255.255.0. Although these subnets are unique communicating on one switch port, they cannot be considered separate VLANs. The switch port supports one VLAN, but multiple subnets can exist on that single VLAN.

A *trunk link*, however, can transport more than one VLAN through a single switch port. Trunk links are most beneficial when switches are connected to other switches or switches are connected to routers.

A trunk link is not assigned to a specific VLAN. Instead, one, many, or all active VLANs can be transported between switches using a single physical trunk link. Connecting two switches with separate physical links for each VLAN is possible. Figure 4-3 shows how two switches might be connected in this fashion.

As VLANs are added to a network, the number of links can quickly grow. A more efficient use of physical interfaces and cabling involves the use of trunking. The right half of the figure shows how one trunk link can replace many individual VLAN links. A trunk link can be associated with a native VLAN, which is used if the trunk link fails for some reason.

Cisco supports trunking on both Fast Ethernet and Gigabit Ethernet switch links, as well as aggregated Fast and Gigabit EtherChannel links. To distinguish between traffic belonging to different VLANs on a trunk link, the switch must have a method of identifying each frame with the appropriate VLAN. Several identification methods are available and are discussed in the next section.

Figure 4-3 *Passing VLAN Traffic Using Single Links Versus Trunk Links*

VLAN Frame Identification

Because a trunk link can be used to transport many VLANs, a switch must identify frames with their VLANs as they are sent and received over a trunk link. *Frame identification*, or *tagging*, assigns a unique user-defined ID to each frame transported on a trunk link. This ID can be thought of as the VLAN number or VLAN “color,” as if each VLAN was drawn on a network diagram in a unique color.

VLAN frame identification was developed for switched networks. As each frame is transmitted over a trunk link, a unique identifier is placed in the frame header. As each switch along the way receives these frames, the identifier is examined to determine to which VLAN the frames belong.

If frames must be transported out another trunk link, the VLAN identifier is retained in the frame header. Otherwise if frames are destined out an access link, the switch removes the VLAN identifier before transmitting the frames to the end station. Therefore, all traces of VLAN association are hidden from the end station.


VLAN identification can be performed using several methods. Each uses a different frame identifier mechanism, and some are suited for specific network media. These methods are described in the sections that follow.

Inter-Switch Link Protocol

The Inter-Switch Link (ISL) protocol is a Cisco proprietary method for preserving the source VLAN identification of frames passing over a trunk link. ISL performs frame identification in Layer 2 by encapsulating each frame between a header and trailer. Any Cisco switch or router device configured for ISL can process and understand the ISL VLAN information. ISL is primarily used for Ethernet media, although Cisco has included provisions to carry Token Ring, FDDI, and ATM frames over Ethernet ISL. (A frame-type field in the ISL header indicates the source frame type.)

When a frame is destined out a trunk link to another switch or router, ISL adds a 26-byte header and a 4-byte trailer to the frame. The source VLAN is identified with a 10-bit VLAN ID in the header. The trailer contains a cyclic redundancy check (CRC) to assure the data integrity of the new encapsulated frame. Figure 4-4 shows how Ethernet frames are encapsulated and forwarded out a trunk link. Because tagging information is added at the beginning and end of each frame, ISL is sometimes referred to as *double tagging*.

Figure 4-4 ISL Frame Identification



If a frame is destined for an access link, the ISL encapsulation (both header and trailer) is removed before transmission. This removal preserves ISL information only for trunk links and devices that can understand the protocol.

IEEE 802.1Q Protocol


The IEEE 802.1Q protocol can also be used to preserve VLAN associations over trunk links. However, this frame identification method is standardized, allowing VLAN trunks to exist and operate between equipment from multiple vendors.

In particular, the IEEE 802.1Q standard defines an architecture for VLAN use, services provided with VLANs, and protocols and algorithms used to provide VLAN services. Further information about the 802.1Q standard can be found at grouper.ieee.org/groups/802/1/pages/802.1Q.html

Like Cisco ISL, IEEE 802.1Q can be used for VLAN identification with Ethernet trunks. Instead of encapsulating each frame with a VLAN ID header and trailer, 802.1Q embeds its tagging information within the Layer 2 frame. This method is referred to as *single-tagging* or *internal tagging*. 802.1Q also introduces the concept of a *native VLAN* on a trunk. Frames belonging to this VLAN are not encapsulated with tagging information. In the event that an end station is connected to an 802.1Q trunk link, the end station will be able to receive and understand only the native VLAN frames.

In an Ethernet frame, 802.1Q adds a four-byte tag just after the source address field, as shown in Figure 4-5.

Figure 4-5 IEEE 802.1Q Frame Tagging Standard



The first two bytes are used as a Tag Protocol Identifier (TPID). The first two bytes also always have a value of 0x8100 to signify an 802.1Q tag. The remaining two bytes are used as a Tag Control Information (TCI) field. The TCI information contains a 3-bit priority field, which is used to implement class of service functions in the accompanying 802.1Q/802.1p prioritization standard. One bit of the TCI is a Canonical Format Indicator (CFI), flagging whether the MAC addresses are in canonical format. The last 12 bits are used as a VLAN Identifier (VID) to indicate the source VLAN for the frame. The VID can have values from 0 to 4095, but VLAN 0, 1, and 4095 are reserved.

NOTE Note that both ISL and 802.1Q tagging methods have one implication: they add to the length of an Ethernet frame. ISL adds a total of 30 bytes to each frame, while 802.1Q adds 4 bytes. Because Ethernet frames cannot exceed 1518 bytes, the additional VLAN tagging information can cause the frame to be too large. Frames that barely exceed the MTU size are called *baby giant frames*. Switches will usually report these frames as Ethernet errors or oversized frames.

LAN Emulation (LANE)

Trunking VLANs between switches over an Asynchronous Transfer Mode (ATM) link is possible. Here, VLANs are transported using the IEEE LAN Emulation (LANE) standard. LANE is discussed in greater detail in Chapter 6, “Trunking with ATM LANE.”

IEEE 802.10

Cisco offers a proprietary method for transporting VLAN information inside the standard IEEE 802.10 FDDI frame. The VLAN information is carried in the Security Association Identifier (SAID) field of the 802.10 frame.

Dynamic Trunking Protocol

Trunk links on Catalyst switches can be manually configured for either ISL or 802.1Q mode. However, Cisco has implemented a proprietary point-to-point protocol called Dynamic Trunking Protocol (DTP) that will negotiate a common trunking mode between two switches. DTP is available in Catalyst supervisor engine software Release 4.2 and later. DTP negotiation should be disabled if a switch has a trunk link connected to a router because the router cannot participate in the DTP negotiation protocol.

NOTE A trunk link can be negotiated between two switches only if both switches belong to the same VLAN Trunking Protocol (VTP) management domain. VTP is discussed in the “VTP Configuration” section of this chapter. If the two switches are in different VTP domains and trunking is desired between them, the trunk links must be set to on or nonnegotiate mode. This setting will force the trunk to be established. These options are explained in the next section.

VLAN Trunk Configuration

By default, all switch ports are non-trunking and operate as access links until some intervention changes the mode. The sections that follow demonstrate the commands necessary to configure VLAN trunks on both an IOS-based and CLI-based switch.

VLAN Trunk Configuration on an IOS-Based Switch

Use the following commands to create a VLAN trunk link on an IOS-based switch:

```
Switch(config)# interface interface mod/port
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport trunk encapsulation {isl | dot1q}
Switch(config-if)# switchport trunk allowed vlan remove vlan-list
Switch(config-if)# switchport trunk allowed vlan add vlan-list
```

Individually, these commands place the switch port into trunking mode, using the encapsulation specified as either **isl** or **dot1q**. The last two commands define which VLANs can be trunked over the link. A list of VLANs is first removed from the trunk because all VLANs (1–1005) are trunked by default. Then, a list of VLANs can be added back into the trunk.

To view the trunking status on a switch port, use the **show interface int mod/port switchport** command.

VLAN Trunk Configuration on a CLI-Based Switch

To create a VLAN trunk link, use the **set trunk** CLI-based command. This command sets the trunking mode and any mode negotiation. The **set trunk** command also identifies the VLANs that will be transported over the trunk link. Trunk configuration uses the following command syntax:

```
Switch(enable) set trunk module/port [on | off | desirable | auto | nonegotiate]
vlan-range [isl | dot1q | dot10 | lane | negotiate]
```

Here, the trunk link is identified by its physical location as the switch module number and port number. The trunking mode can be set to any of the following:

- **on**—This setting places the port in permanent trunking mode. The corresponding switch port at the other end of the trunk should be similarly configured because negotiation is not allowed. The encapsulation or identification mode should also be manually configured.
- **off**—This setting places the port in permanent non-trunking mode. The port will attempt to convert the link to non-trunking mode.
- **desirable**—Selecting this port will actively attempt to convert the link into trunking mode. If the far end switch port is configured to **on**, **desirable**, or **auto** mode, trunking will be successfully negotiated.
- **auto**—The port will be willing to convert the link into trunking mode. If the far end switch port is configured to **on** or **desirable**, trunking will be negotiated. By default, all Fast Ethernet and Gigabit Ethernet links that are capable of negotiating using DTP are configured to this mode. Because of the passive negotiation behavior, the link will never become a trunk, if both ends of the link are left to the **auto** default.
- **nonegotiate**—The port is placed in permanent trunking mode, but no DTP frames are generated for negotiation. The far end switch port must be manually configured for trunking mode.

NOTE

Note that in all modes except **nonegotiate**, DTP frames are sent out every 30 seconds to keep neighboring switch ports informed of the link's mode. On critical trunk links in a network, manually configuring the trunking mode on both ends is best so that the link can never be negotiated to any other state.

By default, a switch will transport all VLANs (1–1000) over a trunk link, even if a VLAN range is specified in the **set trunk** command. There might be times when the trunk link should not carry all VLANs. For example, broadcasts are forwarded to every switch port on a VLAN—including the trunk link because it, too, is a member of the VLAN. If the VLAN doesn't extend past the far end of the trunk link, propagating broadcasts across the trunk makes no sense.

Therefore, to remove VLANs from a trunk link, use the following command:

```
Switch(enable) clear trunk module/port vlan-range
```

Then, if VLANs need to be added back to the trunk, they can be specified as the *vlan-range* in the **set trunk** command.

Lastly, the trunk encapsulation or identification mode is specified at the end of the **set trunk** command. These values are

- **isl**—VLANs are tagged by encapsulating each frame using the Cisco ISL protocol. This protocol is the default, if no value is specified.
- **dot1q**—VLANs are tagged in each frame using the IEEE 802.1Q standard protocol.
- **dot10**—VLANs are tagged on an FDDI switch port using the IEEE 802.10 protocol.
- **lane**—VLANs are identified on an ATM link using LAN Emulation.
- **negotiate**—On Fast and Gigabit Ethernet ports, the mode will be negotiated to select either ISL or IEEE 802.1Q. ISL is preferred, unless one end of the link is configured for **dot1q**.

To view and verify the trunk configuration on a switch, use the **show trunk** [*module/port*] command. Example 4-1 shows a sample output of trunk information.

Example 4-1 *show trunk Verifies Trunk Configuration on a Switch*

```
Switch> (enable) show trunk
Port      Mode      Encapsulation  Status      Native vlan
-----
2/1       auto      dot1q          trunking    1
3/1       auto      isl            trunking    1
3/2       desirable isl            trunking    1
```

continues

Example 4-1 `show trunk` Verifies Trunk Configuration on a Switch (Continued)

```

Port      Vlans allowed on trunk
-----
 2/1      1-1000
 3/1      1-1000
 3/2      1-1000
Port      Vlans allowed and active in management domain
-----
 2/1      1-10,20-35,100,201
 3/1      1,11-19,100,201
 3/2      1,11,15,100,201
Port      Vlans in spanning tree forwarding state and not pruned
-----
 2/1      1-10,20-35,100,201
 3/1      1000
 3/2      1000
Switch> (enable)

```

VLAN Trunking Protocol

As the previous sections have shown, VLAN configuration and trunking on a switch or a small group of switches is fairly easy and straightforward. Campus network environments, however, are usually made up of many interconnected switches. Configuring and managing a large number of switches, VLANs, and VLAN trunks can quickly get out of hand.

Cisco has developed a method to manage VLANs across the campus network. The VLAN Trunking Protocol (VTP) uses Layer 2 trunk frames to communicate VLAN information among a group of switches. VTP manages the addition, deletion, and renaming of VLANs across the network from a central point of control.

VTP Domains

VTP is organized into *management domains* or areas with common VLAN requirements. A switch can belong to only one VTP domain, in addition to sharing VLAN information with other switches in the domain. Similar to VLANs, switches in different VTP domains do not share VTP information.

Switches in a VTP domain advertise several attributes to their domain neighbors. Each advertisement contains information about the VTP management domain, VTP revision number, known VLANs, and specific VLAN parameters. When a VLAN is added to a switch in a management domain, other switches are notified of the new VLAN through VTP advertisements. In this way, all switches in a domain can prepare to receive traffic on their trunk ports using the new VLAN.

VTP Modes

To participate in a VTP management domain, each switch must be configured to operate in one of several modes. The VTP mode will determine how the switch processes and advertises VTP information. The following modes can be used:

- **Server mode**—VTP servers have full control over VLAN creation and modification for their domains. All VTP information is advertised to other switches in the domain, while all received VTP information is synchronized with the other switches. By default, a switch is in VTP server mode. Note that each VTP domain must have at least one server so that VLANs can be created, modified, or deleted, and so that VLAN information can be propagated.
- **Client mode**—VTP clients do not allow the administrator to create, change, or delete any VLANs. Instead, they listen to VTP advertisements from other switches and modify their VLAN configurations accordingly. In effect, this is a passive listening mode. Received VTP information is forwarded out trunk links to neighboring switches in the domain.
- **Transparent mode**—VTP transparent switches do not participate in VTP. While in transparent mode, a switch does not advertise its own VLAN configuration, and a switch does not synchronize its VLAN database with received advertisements. As well, in VTP version 1, a transparent mode switch does not even relay VTP information it receives to other switches. In VTP version 2, transparent switches do forward received VTP advertisements out of their trunk ports, acting as VTP relays.

NOTE

While a switch is in VTP transparent mode, a switch can create and delete VLANs that are local to itself. These VLAN changes, however, will not be propagated to any other switch.

VTP Advertisements

Each switch participating in VTP advertises VLANs, revision numbers, and VLAN parameters on its trunk ports to notify other switches in the management domain. VTP advertisements are sent as multicast frames. The switch intercepts frames sent to the VTP multicast address and processes them with its supervisory processor. VTP frames are forwarded out trunk links as a special case.

Because all switches in a management domain learn of new VLAN configuration changes, a VLAN need only be created and configured on just one VTP server switch in the domain.

By default, management domains are set to use non-secure advertisements without a password. A password can be added to set the domain to secure mode. The same password has to be configured on every switch in the domain so that all switches exchanging VTP information will use identical encryption methods.

The VTP advertisement process starts with configuration revision number 0 (*zero*). When subsequent changes are made, the revision number is incremented before advertisements are sent out. When listening switches receive an advertisement with a greater revision number than is locally stored, the advertisement will overwrite any stored VLAN information. Because of this, forcing any newly added network switches to have revision number zero is important. The VTP revision number is stored in NVRAM and is not altered by a power cycle of the switch. Therefore, the revision number can only be initialized to zero using one of the following methods:

- Change the VTP mode of the switch to *transparent* and then change the mode back to *server*.
- Change the VTP domain of the switch to a bogus name (a non-existent VTP domain) and then change the VTP domain back to the original name.
- Issue a **clear config all** command, which will clear the switch configuration *and* the VTP information stored in NVRAM. Power cycle the switch so that it boots up with a non-existent VTP domain name and a VTP revision number of zero. (*Use caution. This is the most drastic method because it will erase all configuration data.*)

If the VTP revision number is not reset to zero, a new server switch might advertise VLANs as non-existent or deleted. If the advertised revision number happens to be greater than previous legitimate advertisements, listening switches would overwrite good VLAN database entries with null or deleted VLAN status information. This is referred to as a *VTP synchronization problem*.

Advertisements can originate as requests from client-mode switches that want to learn about the VTP database at boot-up time. As well, advertisements can originate from server-mode switches as VLAN configuration changes occur.

VTP advertisements can occur in three forms:

- **Summary advertisements**—VTP domain servers will send summary advertisements every 300 seconds and every time a VLAN topology change occurs. The summary advertisement lists information about the management domain, including VTP version, domain name, configuration revision number, timestamp, MD5 encryption hash code, and the number of subset advertisements to follow. For VLAN configuration changes, summary advertisements are followed by one or more subset advertisements, with more specific VLAN configuration data. Figure 4-6 shows the summary advertisement format.

Figure 4-6 VTP Summary Advertisement Format

Version (1 byte)	Type (Summary Adv) (1 byte)	Number of subset advertisements to follow (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision Number (4 bytes)			
Updater Identity (originating IP address: 4 bytes)			
Update Timestamp (12 bytes)			
MD5 Digest hash code (16 bytes)			

- **Subset advertisements**—VTP domain servers will send subset advertisements after a VLAN configuration change occurs. These advertisements list the specific changes that have been performed, such as creation or deletion of a VLAN, suspending or activating a VLAN, changing the name of a VLAN, and changing the MTU of a VLAN. Subset advertisements can list the following VLAN parameters: status of the VLAN, VLAN type (like Ethernet or Token Ring), MTU, length of the VLAN name, VLAN number, SAID value, and the VLAN name. VLANs are listed individually in sequential subset advertisements. Figure 4-7 shows the VTP subset advertisement format.
- **Advertisement requests from clients**—A VTP client can request any lacking VLAN information. For example, a client switch might be reset and have its VLAN database cleared, its VTP domain membership might be changed, or it might hear a VTP summary advertisement with a higher revision number than it currently has. After a client advertisement request, the VTP domain servers respond with summary and subset advertisements. Figure 4-8 shows the advertisement request format.

Figure 4-7 VTP Subset Advertisement and VLAN Info Field Formats

VTP Subset Advertisement

0	1	2	3
Version (1 byte)	Type (Subset Adv) (1 byte)	Subset sequence number (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision Number (4 bytes)			
VLAN Info Field 1 (see below)			
VLAN Info Field ...			
VLAN Info Field N			

VTP VLAN Info Field

0	1	2	3
Info Length	VLAN Status	VLAN Type	VLAN Name Length
ISL VLAN ID		MTU Size	
802.10 SAID			
VLAN Name (padded with zeros to multiple of 4 bytes)			

Figure 4-8 VTP Advertisement Request Format

0	1	2	3
Version (1 byte)	Type (Adv request) (1 byte)	Reserved (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Starting advertisement to request			

Catalyst switches in server mode use a separate nonvolatile random-access memory (NVRAM) for VTP, different from the configuration NVRAM. All VTP information, including the VTP configuration revision number, is retained even when the switch power is off. In this manner, a switch is able to recover the last known VLAN configuration from its VTP database once it reboots.

VTP Configuration

Before VLANs can be configured, VTP must be configured. By default, every switch will operate in VTP server mode for the management domain *NULL*, with no password or secure mode. The following sections discuss the commands and considerations that should be used to configure a switch for VTP operation.

Configuring a VTP Management Domain

Before a switch is added into a network, the VTP management domain should be identified. If this switch is the first one on the network, the management domain will need to be created. Otherwise, the switch may have to join an existing management domain with other existing switches.

Configuring a VTP Management Domain on an IOS-Based Switch

The following command can be used to assign a switch to a management domain, where the *domain-name* is a text string up to 32 characters long.

```
Switch# vlan database  
Switch(vlan)# vtp domain domain-name
```

Configuring a VTP Management Domain on a CLI-Based Switch

Similar to the command to assign a switch to a management domain on an IOS-based switch, the following command does the same for a CLI-based switch:

```
Switch(enable) set vtp [domain domain-name]
```

Configuring the VTP Mode

Next, the VTP mode needs to be chosen for the new switch. The three VTP modes of operation and their guidelines for use are as follows:

- **Server mode**—Server mode can be used on any switch in a management domain, even if other server and client switches are in use. This mode provides some redundancy in the event of a server failure in the domain. However, each VTP management domain must have at least one server. The first server defined in a network also defines the management domain that will be used by future VTP servers and clients. Server mode is the default VTP mode.
- **Client mode**—If other switches are in the management domain, a new switch should be configured for client mode operation. In this way, the switch will learn any existing VTP information from a server.

If this switch will be used as a redundant server, it should start out in client mode to learn all VTP information from reliable sources. If the switch was initially configured for server mode instead, it might propagate incorrect information to the other domain switches. Once the switch has learned the current VTP information, it can be reconfigured for server mode.

- **Transparent mode**—This mode is used if a switch is not going to share VLAN information with any other switch in the network. VLANs can still be created, deleted, and renamed on the transparent switch. However, they will not be advertised to other neighboring switches. VTP advertisements received by a transparent switch will be forwarded on to other switches on trunk links.
- Keeping switches in transparent mode can eliminate the chance for duplicate, overlapping VLANs in a large network with many network administrators. For example, two administrators might configure VLANs on switches in their respective areas, but use the same VLAN identification or VLAN number. Even though the two VLANs have different meanings and purposes, they could overlap if both administrators advertised them using VTP.

Configuring the VTP Mode on an IOS-Based Switch

On an IOS-based switch, the VTP mode can be configured with the following sequence of commands:

```
Switch# vlan database
Switch(vlan)# vtp domain domain-name
Switch(vlan)# vtp {server | client | transparent}
Switch(vlan)# vtp password password
```

Configuring the VTP Mode on a CLI-Based Switch

On a CLI-based switch, the VTP mode can be configured with the following command:

```
Switch(enable) set vtp [domain domain-name] [mode {server | client | transparent}]
[passwd password]
```

If the domain is operating in secure mode, a password can be included in the command line. The password must be a string of 8 to 64 characters.

Configuring the VTP Version

Two versions of VTP are available for use in a management domain. Catalyst switches are capable of running either VTP version 1 or VTP version 2. Within a management domain, the two versions are not interoperable. Therefore, the same VTP version must be configured on each switch in a domain. VTP version 1 is the default protocol on a switch.

If a switch is capable of running VTP version 2, however, a switch may coexist with other version 1 switches, as long as its VTP version 2 is not enabled. This situation becomes important if you want to use version 2 in a domain. Then, only one server mode switch needs to have VTP version 2 enabled. The new version number is propagated to all other version 2-capable switches in the domain, causing them all to automatically enable version 2 for use.

By default, VTP version 1 is enabled. Version 2 can be enabled or disabled using the **v2** option.

The two versions of VTP differ in the features they support. VTP version 2 offers the following additional features over version 1:

- **Version-dependent transparent mode**—VTP version 1 in transparent mode matches the VTP version and domain name before forwarding the information to other switches using VTP. VTP version 2 in transparent mode differs by forwarding the VTP messages without checking the version number. Because only one domain is supported in a switch, the domain name doesn't have to be checked.
- **Consistency checks**—VTP version 2 performs consistency checks on the VTP and VLAN parameters entered from the CLI or by Simple Network Management Protocol (SNMP). This checking helps prevent errors in such things as VLAN names and numbers from being propagated to other switches in the domain. However, no consistency checks are performed on VTP messages that are received on trunk links or on configuration and database data that is read from NVRAM.
- **Token Ring support**—VTP version 2 supports the use of Token Ring switching and Token Ring VLANs. (If Token Ring switching is being used, VTP version 2 must be enabled.) The section “Token Ring VLANs” later in this chapter discusses these topics in detail.
- **Unrecognized Type-Length-Value (TLV) support**—VTP version 2 switches will propagate received configuration change messages out other trunk links, even if the switch supervisor is not able to parse or understand the message. For example, a VTP advertisement contains a *Type* field to denote what type of VTP message is being sent. VTP message type 1 is a summary advertisement, and message type 2 is a subset advertisement. An extension to VTP could be in use that utilizes other message types and other message length values. Instead of dropping the unrecognized VTP message, version 2 will still propagate the information and keep a copy in NVRAM.

Configuring the VTP Version on an IOS-Based Switch

On an IOS-based switch, the VTP version number is configured using the following commands:

```
Switch# vlan database  
Switch(vlan)# vtp v2-mode
```

Configuring the VTP Version on a CLI-Based Switch

On a CLI-based switch, the VTP version number is configured using the following command:

```
Switch(enable) set vtp v2 enable
```

NOTE Many separate VTP options can be given in a single **set vtp** command, if desired. Consider the following example:

```
Switch(enable) set vtp domain mayberry mode server password opie9 v2 enable
```

Here, one command has set the VTP management domain name to **mayberry**, enabled VTP server mode, set the VTP password to **opie9**, and enabled VTP version 2.

VTP Status

The current VTP parameters for a management domain can be displayed. On an IOS-based switch, use the **show vtp status** command. On a CLI-based switch, use the **show vtp domain** command. Example 4-2 demonstrates some sample output of this command on a CLI-based switch.

Example 4-2 **show vtp domain** Reveals VTP Parameters for a Management Domain

```
Switch> show vtp domain
Domain Name                Domain Index VTP Version Local Mode Password
-----
mydomain                    1           2           server      -

Vlan-count Max-vlan-storage Config Revision Notifications
-----
15          1023             7           disabled

Last Updater    V2 Mode Pruning PruneEligible on Vlans
-----
192.168.1.4     enabled disabled 2-1000

Switch>
```

VTP message and error counters can also be displayed with the **show vtp counters** IOS-based command and the **show vtp statistics** CLI-based command. This command can be used for basic VTP troubleshooting to see if the switch is interacting with other VTP nodes in the domain. Example 4-3 demonstrates some sample output from the **show vtp statistics** command.

Example 4-3 show vtp statistics Reveals VTP Message and Error Counters

```

Switch> show vtp statistics
VTP statistics:
summary advts received      8
subset advts received      11
request advts received      0
summary advts transmitted   1
subset advts transmitted   1
request advts transmitted   0
No of config revision errors 0
No of config digest errors  0

VTP pruning statistics:

Trunk      Join Transmitted  Join Received  Summary advts received from
-----      -----      -----      -----
3/1
Switch>

```


VTP Pruning

Recall that by definition, a switch must forward broadcast frames out all available ports in the broadcast domain because broadcasts are destined everywhere there is a listener. Multicast frames, unless forwarded by more intelligent means, follow the same pattern. (Multicast switching is covered in detail in Chapter 11, “Configuring Multicast Networks.”)

In addition, frames destined for an address that the switch has not yet learned or has forgotten (the MAC address has aged out of the address table) must be forwarded out all ports in an attempt to find the destination. These frames are referred to as *unknown unicast*.

When forwarding frames out all ports in a broadcast domain or VLAN, trunk ports are included. By default, a trunk link transports traffic from all VLANs, unless specific VLANs are removed from the trunk with the **clear trunk** command. Generally, in a network with several switches, trunk links are enabled between switches and VTP is used to manage the propagation of VLAN information. This scenario causes the trunk links between switches to carry traffic from *all* VLANs—not just from the specific VLANs created.


Consider the network shown in Figure 4-9. When end user HostPC in VLAN 3 sends a broadcast, Catalyst switch C forwards the frame out all VLAN 3 ports, including the trunk link to Catalyst A. Catalyst A, in turn, forwards the broadcast on to Catalysts B and D over those trunk links. Catalysts B and D forward the broadcast out only their access links that have been configured for VLAN 3. If Catalysts B and D don’t have any users in VLAN 3, forwarding that broadcast frame to them would consume bandwidth on the trunk links and processor resources in both switches, only to have switches B and D discard the frames.

Figure 4-9 *Flooding in a Catalyst Switch Network*

VTP pruning makes more efficient use of trunk bandwidth by reducing unnecessary flooded traffic. Broadcast and unknown unicast frames on a VLAN are forwarded over a trunk link only if the switch on the receiving end of the trunk has ports in that VLAN. VTP pruning occurs as an extension to VTP version 1, using an additional VTP message type. When a Catalyst switch has a port associated with a VLAN, the switch sends an advertisement to its neighbor switches that it has active ports on that VLAN. The neighbors keep this information, enabling them to decide if flooded traffic from a VLAN should use a trunk port or not.

Figure 4-10 shows the network from Figure 4-9 with VTP pruning enabled. Because Catalyst B has not advertised its use of VLAN 3, Catalyst A will choose not to flood VLAN 3 traffic to it over the trunk link. Catalyst D has advertised the need for VLAN 3, so traffic will be flooded to it.

Figure 4-10 Flooding in a Catalyst Switch Network Using VTP Pruning



Enabling VTP Pruning on an IOS-Based Switch

By default, VTP pruning is disabled on IOS-based switches. In the VLAN database configuration mode, the **vtp pruning** command can be used to enable pruning.

Enabling VTP Pruning on a CLI-Based Switch

VTP pruning is enabled using the **set vtp pruning enable** command. If this command is used on a VTP server, pruning is enabled for the entire management domain. By default, VTP pruning is disabled. When pruning is enabled with this command, all VLANs become eligible for pruning on all trunk links, if needed. The default list of pruning eligibility can be modified. Like VLAN trunking, you can first clear VLANs from the eligibility list using the **clear vtp pruneeligible vlan-range** command. Then, specify the VLANs that can be pruned with the **set vtp pruneeligible vlan-range** command.

NOTE

By default, VLANs 2–1000 are eligible for pruning. VLAN 1 has a special meaning because it is normally used as a management VLAN and is never eligible for pruning. In addition, VLANs 1001–1005 are never eligible for pruning.

Example 4-4 shows the use of these commands. VTP pruning is enabled, and VLAN 6 should be eligible for pruning, while VLANs 5, 7, 8, 9, and 10 are not eligible for pruning. Notice how VLANs 1 and 1001–1005 are never eligible for pruning.

Example 4-4 *Enabling VTP Pruning*

```
Switch(enable) set vtp pruning enable
Switch(enable) clear vtp pruneeligible 5-10
Switch(enable) set vtp pruneeligible 6
```

The pruning status on the switch and its VLANs can be displayed with the **show vtp domain** command. Example 4-5 shows some sample output from this command:

Example 4-5 *show vtp domain Command Output Displays VTP Pruning Status on a Switch and Its VLANs*

```
Switch> show vtp domain
Domain Name                Domain Index VTP Version Local Mode Password
-----
accounting                  1           2           server      -

Vlan-count Max-vlan-storage Config Revision Notifications
-----
3           1023           2           disabled

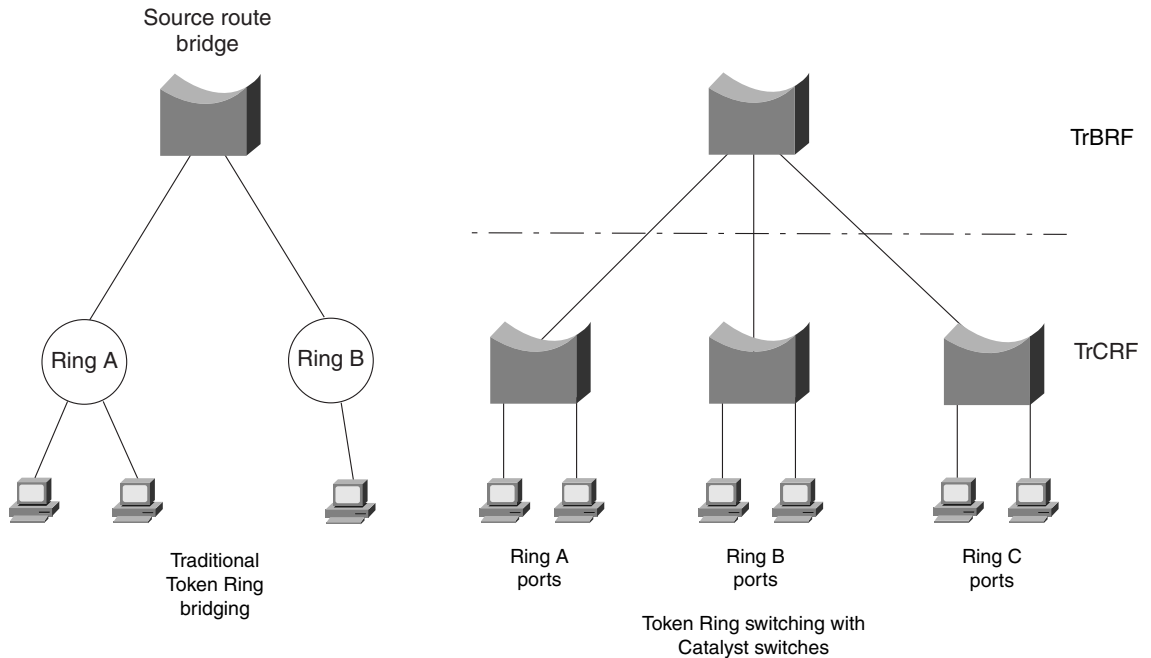
Last Updater V2 Mode Pruning PruneEligible on Vlans
-----
172.16.4.1   disabled enabled 6,11-2000
Switch>
```

Token Ring VLANs

This section discusses VLANs as they are applied to Token Ring networks. Only the Catalyst 5000 and the Catalyst 3900 switches support Token Ring—both using CLI-based commands.

Recall from the discussion in Chapter 3, “Basic Switch and Port Configuration,” the basic topology of Token Ring networks. End stations are connected to multistation access units (MSAUs), which interconnect with other MSAUs to form a ring. Multiple rings can be interconnected by bridges for segmentation and frame forwarding using source-route bridging and the RIF information. Figure 4-11 shows a typical Token Ring network with two rings and a source-route bridge in the left half. The right half of the figure shows a similar network topology with three rings and a source-route bridge, using the Token Ring switching features of Catalyst switches. The functionality of rings and bridges are performed within the switches, using Token Ring switching functions.

Figure 4-11 *Token Ring Networks: Traditional and Switched*



Token Ring switching follows the same topology, but performs the various functions within the switch. Where groups of end stations are connected by MAUs in a ring, the IEEE has defined the Concentrator Relay Function (CRF). The function of a multiport bridge to connect individual rings is defined as the Bridge Relay Function (BRF). These functions as performed by Catalyst switches are further described in the sections that follow.

TrBRF

A Catalyst switch connects *logical Token Ring Concentrator Relay Functions (TrCRFs)* with a logical multiport bridge, or *Token Ring Bridge Relay Function (TrBRF)*. In the hierarchy of bridged Token Rings, each TrCRF must be connected to a *parent TrBRF*. The hierarchical structure of Token Ring VLANs is shown in the right half of Figure 4-11.

By default, the TrBRF interconnects only TrCRFs located on the local switch. However, if trunking is used with ISL encapsulation, the TrBRF can extend to TrCRFs located on other Catalyst switches.

Each TrBRF exists as a special VLAN within a Catalyst switch. A switch can support many TrBRFs, but only one VLAN can be assigned to each TrBRF. By default, one TrBRF is defined as “trbrf-default” on VLAN 1005. Each TrBRF can operate as either a source-route bridge

(SRB), a source-route transparent (SRT) bridge, or both as a mixed mode. Additionally, each TrBRF runs a separate instance of either the IBM or IEEE Spanning-Tree Protocol to prevent bridging loops. The Spanning-Tree Protocol is covered in Chapter 5.

NOTE

To create and use Token Ring VLANs, VTP version 2 must be enabled on all Catalyst switches in the Token Ring domain. Enabling VTP version 2 was discussed in the previous “VTP Configuration” section.

To define a TrBRF on a Catalyst switch, use the following command:

```
Switch(enable) set vlan vlan-num [name name] type trbrf bridge bridge-num [stp {ieee | ibm}]
```

The only two required fields for a TrBRF are the VLAN number and the bridge number. Bridge numbers are defined by a single byte value, as a hexadecimal number from 0x1 to 0xf. The default Spanning-Tree Protocol is **ibm**. Notice that the type of bridging is not defined with the TrBRF, although the TrBRF performs the actual bridging function. Instead, the type of bridging is defined at the TrCRF. This way, multiple TrCRFs can be connected by a single parent TrBRF, each bridged with the desired method.

TrCRF

In a Catalyst switch, individual Token Ring ports can be connected to a logical ring, or Token Ring VLAN, by assigning them with identical ring numbers. Internally, the Catalyst performs the TrCRF to maintain the ring connectivity. Frame forwarding between ports on a common ring is performed with source-route switching, using either MAC addresses or route descriptors.

The TrCRF can be confined within a single switch or can be spread across multiple switches, depending on the topology and switch configuration. When a TrCRF is contained completely within a switch, it is referred to as an *undistributed TrCRF*. However, a TrCRF can be distributed across multiple switches if ISL trunking is enabled between switches and TrCRFs with identical VLAN numbers are defined.

By default, one TrCRF is defined on every Catalyst switch as “trcrf-default” on VLAN 1003. The default TrCRF is also assigned to the default TrBRF on VLAN 1005. If ISL trunking is in use, every Token Ring port on every switch will be defined to the same distributed TrCRF. Because only one TrBRF is defined by default, no bridging will occur. Instead, source-route switching will be performed to forward frames between switch ports within the TrCRF.

To define a TrCRF on a Catalyst switch, use the following command:

```
Switch(enable) set vlan vlan-num [name name] type trcrf {ring hex-ring-num | decring decimal-ring-num} parent vlan-num
```

Both the ring number and the parent VLAN number must be specified to define a TrCRF. The ring number can be defined as a hexadecimal value of 0x1 to 0xfff with the **ring** option, or as a decimal value of 1 to 4095 with the **decring** option. The parent VLAN number must match the VLAN number assigned to the parent TrBRF.

On the Catalyst 5000, a single TrCRF can be distributed across multiple switches. To enable this feature, use the **set tokenring distrib-crf enable** command.

After a TrCRF VLAN has been created, switch ports can be assigned to it. As with Ethernet switching, use the following command to assign ports to a VLAN:

```
Switch(enable) set vlan vlan-num mod-num/port-num
```

To view the current Token Ring VLAN configuration, use the **show vlan** command. The output of which is demonstrated in Example 4-6.

Example 4-6 show vlan Command Output Displays the Current Token Ring VLAN Configuration

```
Switch(enable) show vlan
```

VLAN Name	Status	Mod/Ports, Vlans
1 default	active	1/1-2
800 brf800	active	801,802
801Floor_1	active	
802Floor_2	active	
1002 fddi-default	active	
1003 trcrf-default	active	2/1-16
1004 fddinet-default	active	
1005 trbrf-default	active	1003

VLAN	Type	SAID	MTU	Parent	RingNo	BrdgNo	Stp	BrdgMode	Trans1	Trans2
1	enet	100001	1500	-	-	-	-	-	0	0
800	trbrf	100800	4472	-	-	0x2	ibm	-	0	0
801	trcrf	100201	4472	800	0x01	-	srb	-	0	0
802	trcrf	100202	4472	800	0x02	-	srb	-	0	0
1002	fddi	101002	1500	-	0x0	-	-	-	0	0
1003	trcrf	101003	4472	1005	0xccc	-	srb	-	0	0
1004	fdnet	101004	1500	-	-	0x0	ieee	-	0	0
1005	trbrf	101005	4472	-	-	0xf	ibm	-	0	0

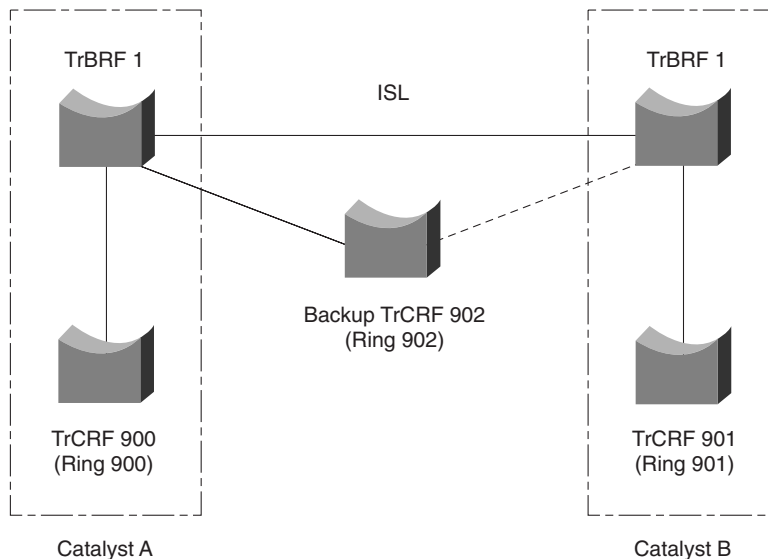
Notice in this example that one TrBRF has been assigned to VLAN 800, and two TrCRFs to VLANs 801 and 802. TrBRF on VLAN 800 shows the two active TrCRFs. In the lower group of output lines, the two TrCRFs on VLANs 801 and 802 show the parent TrBRF as 800 with source-route bridging enabled. Even though the defaults are not being used in this example, the default TrBRF is listed as VLAN 1005, and the default TrCRF is VLAN 1003, assigned to the default TrBRF.

TrCRF Redundancy

Catalyst switches also offer a form of redundancy for Token Ring switching. When two switches are connected by a common TrBRF *and* ISL trunking is enabled, connectivity between the TrCRFs in the switches could be disrupted if the ISL trunk link fails. A *backup TrCRF* can be used to provide a backup path in this case.

For each TrBRF, a single backup TrCRF can be defined with a single port from each connected switch. Only one of the TrCRF ports will be active at all times, while the other ports will be disabled. If the ISL trunk link goes down (along with the common TrBRF), the backup TrCRF links will come up and pass traffic between switches. A backup TrCRF is shown in Figure 4-12. The backup TrCRF is first defined on all switches, assigned to the TrBRF connecting the switches, marked as a backup TrCRF, and then assigned to one port on each switch.

Figure 4-12 A Backup TrCRF



To enable a backup TrCRF, first define a TrCRF that spans between switches. Then assign one port from each switch to the backup TrCRF. Finally, use the **set vlan *vlan-num* backupcrf on** command to enable the backup TrCRF function.

VTP and Token Ring VLANs

Using VTP in a Token Ring network domain will simplify VLAN administration, just as it does for Ethernet. TrCRF information will be propagated to all switches in a management domain.

As well, VTP pruning can also be performed on Token Ring VLANs. Both the default TrBRF (VLAN 1005) and the default TrCRF (VLAN 1003) are always pruning ineligible. VTP pruning is configured on a per-TrBRF basis. When a TrBRF is made pruning-eligible, all TrCRFs connected to it are also made pruning-eligible.

Duplicate Ring Protocol (DRiP)

Catalyst switches also have a mechanism to monitor the use of TrCRFs or ring numbers within a domain of switches. The *Duplicate Ring Protocol (DRiP)* collects and maintains the status of TrCRFs that are interconnected by TrBRFs. This information is used for the following purposes:

- Preventing duplicate ring numbers from being assigned to TrCRFs.
- Filtering All-Routes Explorer (ARE) frames from reentering TrCRFs that they have already visited.
- Operating the backup TrCRF function when an ISL trunk link goes down.

Every switch participating in Token Ring switching sends a DRiP advertisement out all ISL trunk ports every 30 seconds. Advertisements are sent to multicast address 01:00:0C:CC:CC:CC and are sent only on the default VLAN 1. When a switch receives the multicast advertisements, the switch does not forward the advertisements on to other switches over ISL links unless the advertisements contain new information. As well, a switch compares advertisements to the information in its own configuration. If it detects that a TrCRF has already been configured elsewhere, the local TrCRF configuration will be denied.

Foundation Summary

The Foundation Summary is a collection of tables and figures that provides a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final preparation before the exam, these tables and figures will hopefully be a convenient way to review the day before the exam.

Figure 4-13 *VTP Summary Advertisement Format*

Version (1 byte)	Type (Summary Adv) (1 byte)	Number of subset advertisements to follow (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision Number (4 bytes)			
Updater Identity (originating IP address; 4 bytes)			
Update Timestamp (12 bytes)			
MD5 Digest hash code (16 bytes)			

Figure 4-14 *VTP Subset Advertisement and VLAN Info Field Formats***VTP Subset Advertisement**

0	1	2	3
Version (1 byte)	Type (Subset Adv) (1 byte)	Subset sequence number (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision Number (4 bytes)			
VLAN Info Field 1 (see below)			
VLAN Info Field ...			
VLAN Info Field N			

VTP VLAN Info Field

0	1	2	3
Info Length	VLAN Status	VLAN Type	VLAN Name Length
ISL VLAN ID		MTU Size	
802.10 SAID			
VLAN Name (padded with zeros to multiple of 4 bytes)			

Figure 4-15 *VTP Subset Request Format*

0	1	2	3
Version (1 byte)	Type (Adv request) (1 byte)	Reserved (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Starting advertisement to request			

Table 4-2 *VLAN Configuration Commands*

Task	IOS-Based Command	CLI-Based Command
Create VLAN	vlan database vlan <i>vlan-num</i> name <i>vlan-name</i>	set vlan <i>vlan-num</i> [name <i>name</i>]
Assign port to VLAN	interface <i>interface module/number</i> switchport mode access switchport access vlan <i>vlan-num</i>	set vlan <i>vlan-num mod-num/port-list</i>
Display VLANs	show vlan	show vlan
Configure trunk	interface <i>interface mod/port</i> switchport mode trunk switchport trunk encapsulation { isl dot1q } switchport trunk allowed vlan remove <i>vlan-list</i> switchport trunk allowed vlan add <i>vlan-list</i>	set trunk <i>module/port</i> [on off desirable auto nonegotiate] <i>vlan-range</i> [isl dot1q dot10 lane negotiate] clear trunk <i>module/port vlan-range</i>
Display trunks	show interface <i>mod/num</i> switchport	show trunk

Table 4-3 *VTP Configuration Commands*

Task	IOS-Based Command	CLI-Based Command
Configure VTP domain	vlan database vtp domain <i>domain-name</i>	set vtp [domain <i>domain-name</i>]
Configure VTP mode	vlan database vtp domain <i>domain-name</i> vtp { server client transparent } vtp password <i>password</i>	set vtp [domain <i>domain-name</i>] [mode { server client transparent }] [passwd <i>password</i>]
Configure VTP version	vlan database vtp v2-mode	set vtp v2 enable
Display VTP status	show vtp status show vtp counters	show vtp domain show vtp statistics
VTP pruning	vtp pruning	set vtp pruning enable set vtp pruneeligible <i>vlan-range</i> clear vtp pruneeligible <i>vlan-range</i>

Table 4-4 *Token Ring VLAN Configuration Commands*

Task	IOS-Based Command	CLI-Based Command
Define TrBRF	N/A	set vlan <i>vlan-num</i> [name <i>name</i>] type trbrf bridge <i>bridge-num</i> [stp { <i>ieee</i> <i>ibm</i> }]
Define TrCRF	N/A	set vlan <i>vlan-num</i> [name <i>name</i>] type trcrf { ring <i>hex-ring-num</i> decring <i>decimal-ring-</i> <i>num</i> } parent <i>vlan-num</i>
Enable distributed TrCRF	N/A	set tokenring distrib-crf enable
Assign Token Ring ports to TrCRF	N/A	set vlan <i>vlan-num</i> <i>mod-num/port-num</i>
Enable backup TrCRF	N/A	set vlan <i>vlan-num</i> backupcrf on

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however, they are designed to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to these questions can be found in Appendix A, on page 477.

- 1 What is a VLAN? When is it used?

- 2 When a VLAN is configured on a Catalyst switch port, in how much of the campus network will the VLAN number be unique and significant?

- 3 What are two types of VLANs, in terms of spanning areas of the campus network?

- 4 What is the Catalyst CLI-based switch command to configure ports 4/11 and 5/1 through 5/24 for VLAN 2?

- 5** Generally speaking, what must be configured (both switch and end user device) for a port-based VLAN?

- 6** What is the default VLAN on all ports of a Catalyst switch?

- 7** What are the components of a Token Ring VLAN?

- 8** What is a trunk link?

- 9** What methods of Ethernet VLAN frame identification can be used on a Catalyst switch?

- 10** What is the difference between these two trunking methods? How many bytes are added to trunked frames for VLAN identification in each method?

- 11** What is the purpose of Dynamic Trunking Protocol (DTP)?

- 12 What CLI-based commands are needed to configure a Catalyst switch trunk port 1/1 to transport only VLANs 100, 200–205, and 300 using IEEE 802.1Q? (Assume that trunking is enabled and active on the port already.)

- 13 What VTP modes can a Catalyst switch be configured for? Can VLANs be created in each of the modes?

- 14 Two neighboring switch trunk ports are set to *auto* mode with *ISL* trunking mode. What will the resulting trunk mode become?

- 15 How many VTP management domains can a Catalyst switch participate in? How many VTP servers can a management domain have?

- 16 What CLI-based command can be used on a Catalyst switch to verify exactly what VLANs will be transported over a trunk link?

- 17 What conditions must exist for two Catalyst switches to be in the same VTP management domains?

- 18** What are the types of VTP messages or advertisements used by Catalyst switches? What field in these messages determines if a switch should use and record VLAN data in the messages?

- 19** What CLI-based command can be used to configure a Catalyst switch to become a VTP server for the domain “engineering”? The domain should be secured with the password “secret123.”

- 20** What is the purpose of VTP pruning?

- 21** Which VLAN numbers are never eligible for VTP pruning? Why?

- 22** What commands can be used to make only VLANs 300 and 400 eligible for VTP pruning?

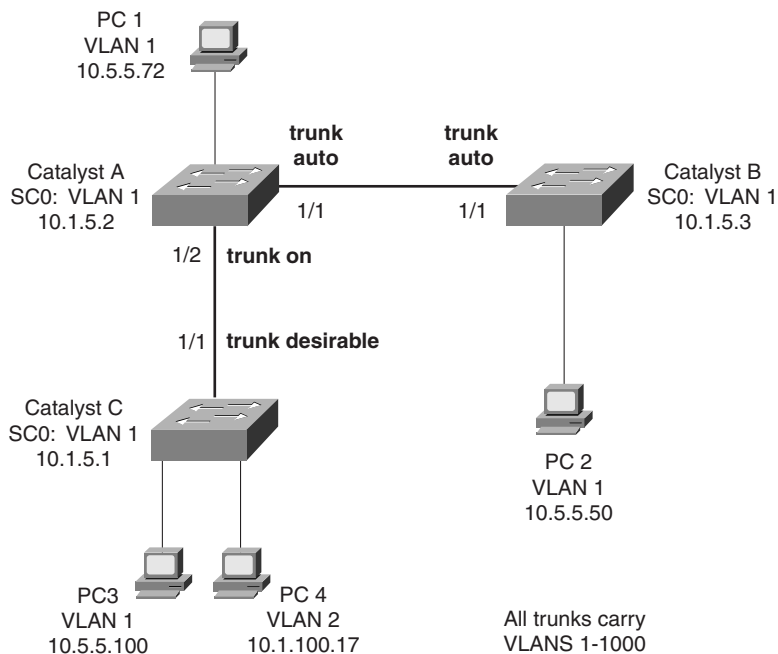
- 23** What are the steps needed to establish Token Ring switching with VLANs?

Scenarios

Scenario 4-1

Consider the network shown in Figure 4-16 and answer the questions that follow.

Figure 4-16 Diagram for Scenario 4-1

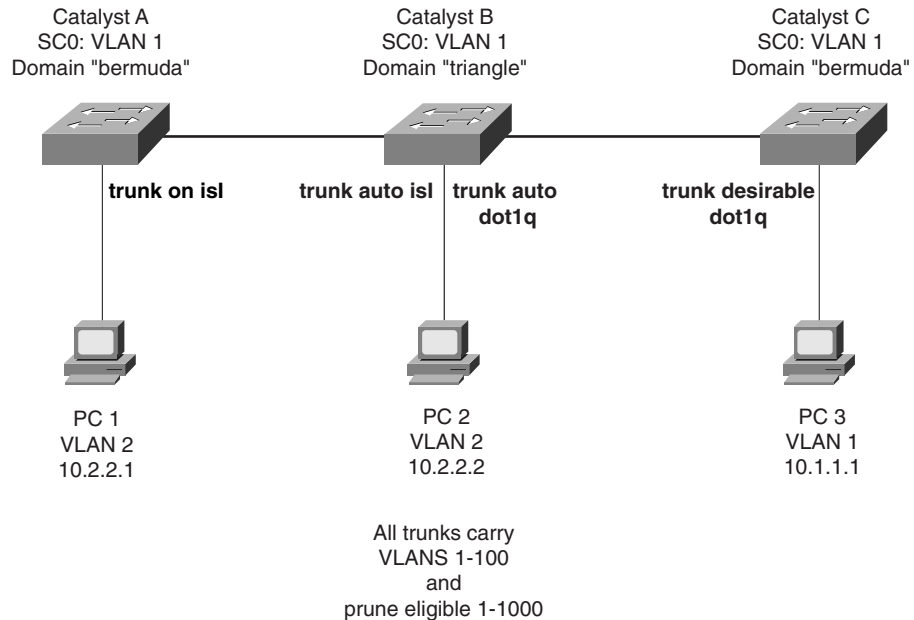


- 1 What is the mode of the link between Catalyst A and Catalyst B?
- 2 Now suppose the network administrator types the command **set trunk 1/1 nonegotiate** on Catalyst B. What will the link mode be now?
- 3 Catalyst B has been given the command **set trunk 1/1 on**. What is the link mode now?
- 4 What is the mode of the link between Catalyst A and Catalyst C?
- 5 Assume that all links between Catalyst switches are in trunking mode, transporting VLANs 1-1000. Can PC-2 ping PC-4?
- 6 Suppose PC-1 begins to generate a broadcast storm. Where would the effects of this storm be experienced in this network? Consider both devices and links. Will PC-4 receive the broadcasts?

Scenario 4-2

See the diagram shown in Figure 4-17 and answer the questions that follow.

Figure 4-17 Diagram for Scenario 4-2



- 1 What is the mode of the link between Catalyst A and Catalyst B?
- 2 Can Catalyst A **ping** Catalyst C? Can PC-1 **ping** PC-2? Why or why not?
- 3 Suppose Catalyst B's VTP domain is now changed to "bermuda." Can Catalyst A **ping** Catalyst C?
- 4 Which Catalyst switches will receive a broadcast from PC-1?
- 5 Where will VLAN1 be pruned? Why?
- 6 Now suppose Catalyst A is a VTP server, Catalyst C is a VTP client, and Catalyst B is configured for VTP transparent mode. All switches are in the "bermuda" management domain. If VLAN14 is created on Catalyst A, which switches will also get VLAN14 created via VTP?
- 7 If VLAN15 is created on Catalyst B, what other switches will also create VLAN15 via VTP?
- 8 If VLAN16 is created on Catalyst C, what will happen?

Scenarios Answers

Scenario Answers 4-1

- 1 The link is still an access link, with no trunking established because both switches are set to *auto* mode. The switches are each passively waiting for the other to initiate trunking.
- 2 Trunking is still not established. Catalyst A is waiting to be asked to trunk, and Catalyst B is set to *nonegotiate*. Catalyst B will never try to negotiate trunking.
- 3 Trunking has finally been established.
- 4 Trunking. Catalyst A expects trunking on the link, while Catalyst C actively tries to negotiate trunking.
- 5 No. The two PC devices are connected to different VLANs. Without a router or Layer 3 device connecting the VLANs, no traffic will cross between them.
- 6 All hosts on VLAN1 (PC-1, PC-2, and PC-3) will experience the broadcast storm. All trunk links between switches will transport the broadcast frames. In addition, all switch supervisor CPUs will receive and process the broadcasts, because each switch has its SC0 port assigned to VLAN1. (For this reason, it is recommended to reserve VLAN1 for management traffic only. User-generated broadcasts can overload the switch supervisor to the extent that it can no longer keep track of its management protocols like VTP, CDP, and so forth. Instead, all user traffic should be kept off VLAN1.)

Scenario Answers 4-2

- 1 The link is still an access link, with no trunking established. The two switches would have negotiated trunking, but the switches are configured for different VTP management domains. Neighboring switches must be in the same domain for trunking to be negotiated.
- 2 Catalyst A can **ping** Catalyst B. The SC0 ports on both switches are configured for the same VLAN. Because trunking has not been established between Catalyst A and Catalyst B (due to domain name conflicts), the link is still an access link. Fortunately, the access link has defaulted to VLAN1 so that the two SC0 ports on VLAN1 can communicate.

PC-1 cannot **ping** PC-2, however. Both PCs are in VLAN2, but VLAN2 is not being transported between switches because the trunk link has not been established.
- 3 Yes. Trunk links are now negotiated or established between all switches.
- 4 Catalyst A and Catalyst B. Because Catalyst C has no ports in VLAN2 (where PC-1 resides), VLAN2 will be pruned by Catalyst B and will not cross the trunk link to Catalyst C.

- 5 VLAN1 will not be pruned at all. Although VLAN1 is present on all switches, it is not pruned because VLAN1 is ineligible for pruning by definition. Remember that VLAN1 is usually used for management traffic and should be kept intact so that no switches become isolated.
- 6 Only Catalyst C will create VLAN14 in response to VTP advertisements. Catalyst B in transparent mode will only relay the VTP information without interpreting the information.
- 7 Only Catalyst B will create VLAN15. Because it is in transparent mode, no VLAN activity will be advertised to other neighboring switches. However, Catalyst B is allowed to create, delete, and rename VLANs freely. These VLANs are significant only to the local switch.
- 8 Catalyst C will not allow any VLANs to be created, unless they are learned from a VTP server in the “bermuda” domain. Because it is in VTP client mode, no VLAN changes can be performed from the console.