After completing this chapter, you will be able to perform the following tasks:

- Identify what a VLAN is and how it operates.
- Configure a VLAN to improve network performance.
- Identify what role the switch plays in the creation of VLANs.
- Identify how network devices communicate about VLANs.
- Describe the need and operation of the VLAN Trunking Protocol.
- Configure the Catalyst Switch for VLAN operation.

CHAPTER 3

Extending Switched Networks with Virtual LANs

The design and function of a bridged/switched network is to provide enhanced network services by segmenting the network into multiple collision domains. The fact remains, however, that without any other mechanism, the bridged/switched network is still a single broadcast domain. A broadcast domain is a group of devices that can receive one another's broadcast frames. For example, if device A sends a broadcast frame and that frame is received by devices B and C, all three devices are said to be in a common broadcast domain. Because broadcast frames are flooded out all ports on a bridge/switch (by default), the devices connected to the bridge/switch are in a common broadcast domain.

Controlling broadcast propagation throughout the network is important to reduce the amount of overhead associated with these frames. Routers, which operate at Layer 3 of the OSI model, provide broadcast domain segmentation for each interface. Switches can also provide broadcast domain segmentation using virtual LANs (VLANs). A VLAN is a group of switch ports, within a single or multiple switches, that is defined by the switch hardware and/or software as a single broadcast domain. A VLAN's goal is to group devices connected to a switch into logical broadcast domains to control the effect that broadcasts have on other connected devices. A VLAN can be characterized as a logical network.

The benefits of VLANs include the following:

- Security
- Segmentation
- Flexibility

VLANs enable you to group users into a common broadcast domain regardless of their physical location in the internetwork. Creating VLANs improves performance and security in the switched network by controlling broadcast propagation and requiring that communications between these broadcast be carried out by a Layer 3 device that is capable of implementing security features such as access control lists (ACLs).

In a broadcast environment, a broadcast sent out by a host on a single segment would propagate to all segments. In normal network operation, hosts frequently generate broadcast/multicast traffic. If hundreds or thousands of hosts each sent this type of traffic, it would saturate the bandwidth of the entire network, as shown in Figure 3-1. Also, without forcing some method of checking at an upper layer, all devices in the broadcast domain would be able to communicate via Layer 2. This severely limits the amount of security you can enforce on the network.

Figure 3-1 *Broadcast Propagation*

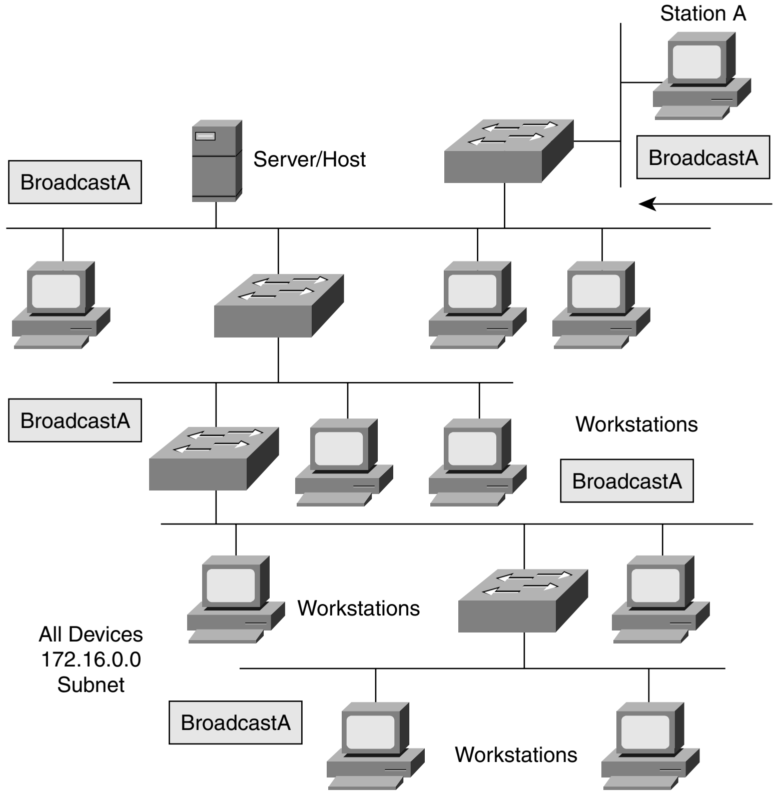

Before the introduction of switches and VLANs, internetworks were divided into multiple broadcast domains by connectivity through a router. Because routers do not forward broadcasts, each interface is in a different broadcast domain. Figure 3-2 shows an internetwork broken into multiple broadcast domains using routers. Notice that each segment is an individual IP subnet and that regardless of a workstation's function, its subnet is defined by its physical location.

A VLAN is a logical broadcast domain that can span multiple physical LAN segments. A VLAN can be designed to provide independent broadcast domains for stations logically segmented by functions, project teams, or applications, without regard to the users' physical location. Each switch port can be assigned to only one VLAN. Ports in a VLAN share broadcasts. Ports that do not belong to the same VLAN do not share broadcasts. This control of broadcast improves the internetwork's overall performance.

VLANs enable switches to create multiple broadcast domains within a switched environment, as illustrated in Figure 3-3.

Figure 3-2 *Multiple Broadcast Domains Using Routers*

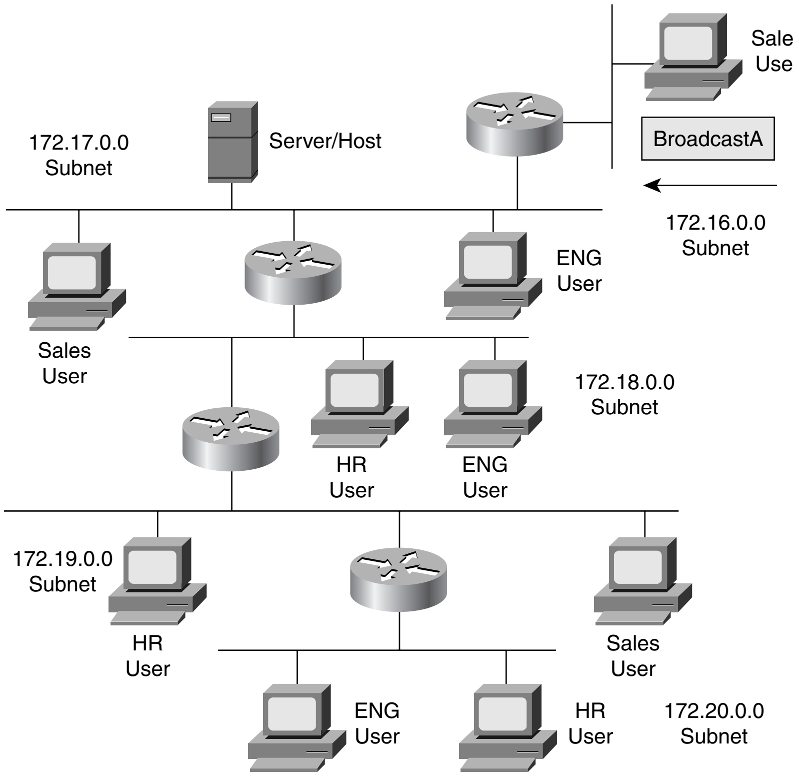

Notice that now all users in a given group (department in this example) are defined to be in the same VLAN. Any user in this VLAN receives a broadcast from any other member of the VLAN, while users of other VLANs do not receive these broadcasts. Each of the users in a given VLAN is also in the same IP subnet. This is different from the broadcast domains of Figure 3-2, in which the physical location of the device determines the broadcast domain. However, there is a similarity with a legacy, non-VLAN internetwork because a router is still needed to get from one broadcast domain to another, even if a VLAN is used to define the broadcast domain instead of a physical location. Therefore, the creation of VLANs does not eliminate the need for routers.

Within the switched internetwork, VLANs provide segmentation and organizational flexibility. Using VLAN technology, you can group switch ports and their connected users into logically defined communities of interest, such as coworkers in the same department, a cross-functional product team, or diverse user groups sharing the same network application.

A VLAN can exist on a single switch or span multiple switches. VLANs can include stations in a single building or multiple-building infrastructures. In rare and special cases, they can even connect across wide-area networks (WANs).

Figure 3-3 *VLAN Overview*

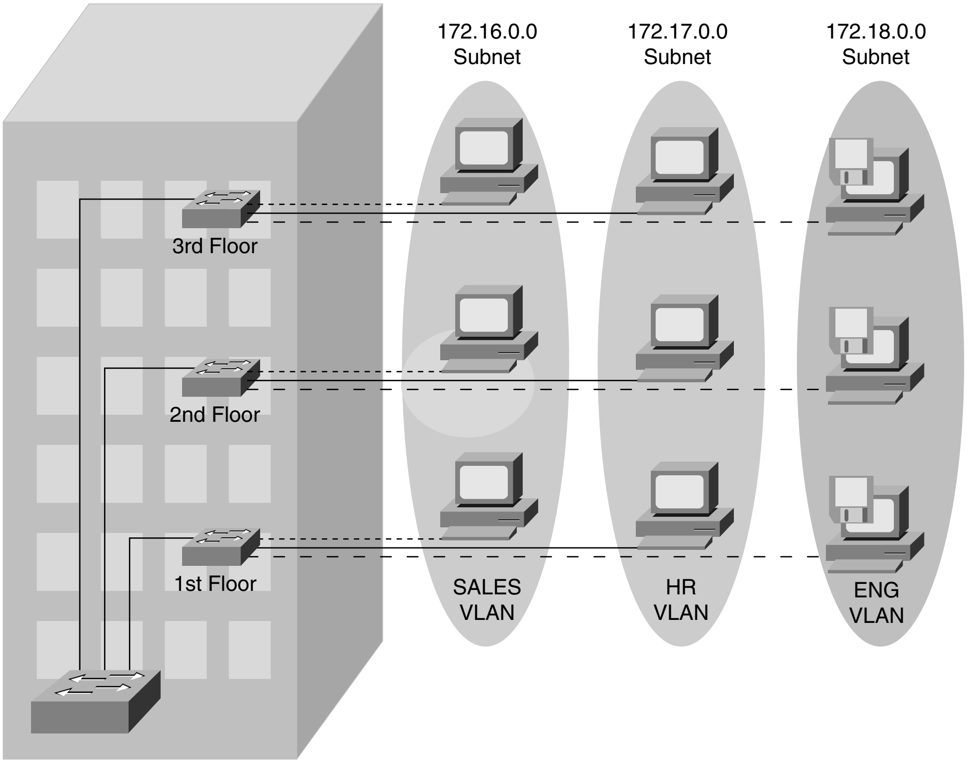

VLAN Concepts

As mentioned previously, prior to the VLAN, the only way to control broadcast traffic was through segmentation using routers. VLANs are an extension of a switched and routed internetwork. By having the ability to place segments (ports) in individual broadcast domains, you can control where a given broadcast is forwarded. The sections that follow expand on these concepts. Basically, each switch acts independently of other switches in the network. With the concept of VLANs, a level of interdependence is built into the switches themselves. The characteristics of a typical VLAN setup are as follows:

- Each logical VLAN is like a separate physical bridge.
- VLANs can span multiple switches.
- Trunk links carry traffic for multiple VLANs.

With VLANs, each switch can distinguish traffic from different broadcast domains. Each forwarding decision is based on which VLAN the packet came from; therefore, each VLAN acts like an individual bridge within a switch. To bridge/switch between switches, you must either connect each VLAN independently (that is, dedicate a port per VLAN) or have some method of maintaining and forwarding the VLAN information with the packets. A process called *trunking* allows this single connection. Figure 3-4 illustrates a typical VLAN setup in which multiple VLANs span two switches interconnected by a Fast Ethernet trunk.

Figure 3-4 *Multiple VLANs Can Span Multiple Switches*

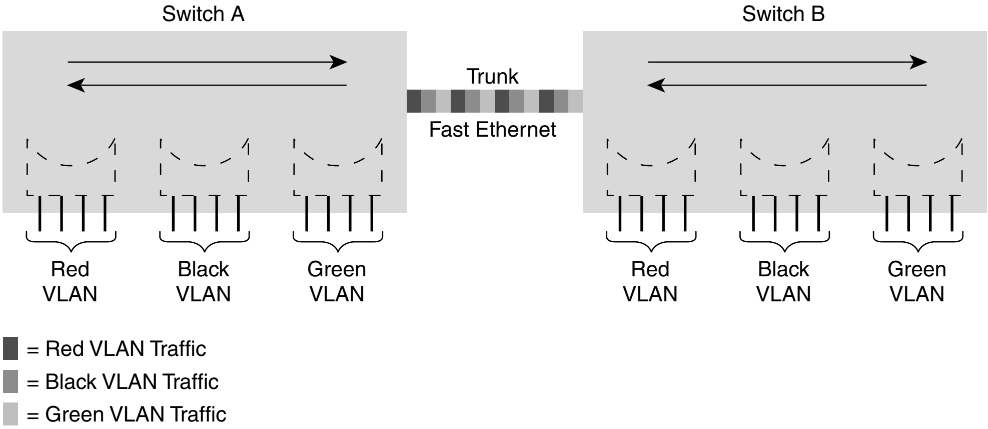

How VLANs Operate

A Catalyst switch operates in your network like a traditional bridge. Each VLAN configured on the switch implements address learning, forwarding/filtering decisions, and loop avoidance mechanisms as if it were a separate physical bridge. This VLAN might include several ports, possibly on multiple switches.

Internally, the Catalyst switch implements VLANs by restricting data forwarding to destination ports in the same VLAN as originating ports. In other words, when a frame arrives on a switch port, the Catalyst must retransmit the frame only to a port that belongs to the same VLAN as that of the incoming port. The implication is that a VLAN operating on a Catalyst switch limits transmission of unicast, multicast, and broadcast traffic. Flooded traffic originating from a particular VLAN floods out only other ports belonging to that VLAN. Each VLAN is an individual broadcast domain because a broadcast in a given VLAN will never reach any ports in other VLANs.

Normally, a port carries traffic only for the single VLAN to which it belongs. For a VLAN to span multiple switches on a single connection, a trunk is required to connect two switches. A trunk carries traffic for all VLANs by identifying the originating VLAN as the frame is carried between the switches. Figure 3-4 shows a Fast Ethernet trunk carrying multiple VLANs between the two switches. Most ports on Catalyst switches are capable of being trunk ports. Any port on a Catalyst 2950 can be a trunk port.

VLAN Membership Modes

VLANs are a Layer 2 implementation in your network's switching topology. Because they are implemented at the data link layer, they are protocol-independent. To put a given port (segment) into a VLAN, you must create a VLAN on the switch and then assign that port membership on the switch. After you define a port to a given VLAN, broadcast, multicast, and unicast traffic from that segment will be forwarded by the switches only to ports in the same VLAN. If you need to communicate between VLANs, you must add a router (or Layer 3 switch) and a Layer 3 protocol to your network.

The ports on a Layer 2 Catalyst switch, such as a 2950, all function as Layer 2 ports. In Cisco IOS Software, a Layer 2 port is known as a *switchport*. A switchport can either be a member of a single VLAN or be configured as a trunk link to carry traffic for multiple VLANs. When a port is in a single VLAN, the port is called an *access port*. Access ports are configured with a VLAN membership mode that determines to which VLAN they can belong. The membership modes follow:

- **Static**—When an administrator assigns a single VLAN to a port, it is called *static assignment*. By default, all Layer 2 switchports are statically assigned to VLAN 1 until an administrator changes this default configuration.
- **Dynamic**—The IOS Catalyst switch supports the dynamic assignment of a single VLAN to a port by using a VLAN Membership Policy Server (VMPS). The VMPS must be a Catalyst Operating System switch, such as a Catalyst 5500 or 6500, running the set-based operating system. An IOS-based Catalyst switch cannot operate as the VMPS. The VMPS contains a database that maps MAC addresses to VLAN assignment. When a frame arrives on a dynamic port, the switch queries the VMPS for the VLAN assignment based on the arriving frame's source MAC address.

A dynamic port can belong to only one VLAN at a time. Multiple hosts can be active on a dynamic port only if they all belong to the same VLAN. Figure 3-5 demonstrates the static and dynamic VLAN membership modes.

Figure 3-5 *VLAN Membership Modes*

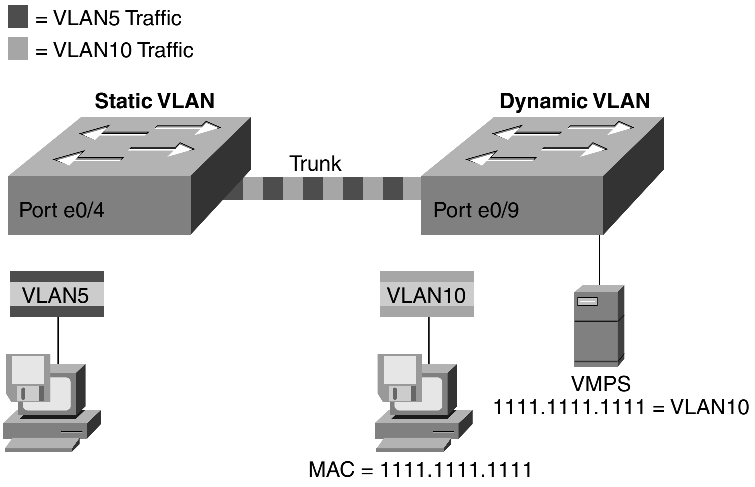

For an access port, the VLAN identity is not known by the sender or receiver attached to the access port. Frames going into and out of access ports are standard Ethernet frames, as discussed in Chapter 2, "Configuring Catalyst Switch Operations." The VLAN identity is used only within the switch to provide broadcast domain boundaries.

Trunk Links

A broadcast domain must sometimes exist on more than one switch in the network. To accomplish this, one switch must send frames to another switch and indicate which VLAN a particular frame belongs to. On Cisco switches, a *trunk link* is created to accomplish this VLAN identification. ISL and IEEE 802.1Q are different methods of putting a VLAN identifier in a Layer 2 frame.

A trunk link is the other type of Layer 2 port supported on Cisco switches. When a trunk port is configured, it begins marking frames as they exit the port to indicate which VLAN each frame is associated with. The trunk port can also read the markings, called tags, as they enter the trunk port. This enables the switch to send a frame only to the ports for the given VLAN associated with the incoming frame.

The main purpose of trunking is to carry traffic between switches and maintain the VLAN information. Unlike an access link, the trunk link does not belong to a single VLAN but instead can carry traffic from several VLANs over a point-to-point link between two devices that understand the protocol. Because a trunk is typically a point-to-point connection between two switches, it is very efficient and highly recommended that it runs in full-duplex mode. Figure 3-6 shows trunk links between switches carrying traffic for multiple VLANS.

Figure 3-6 *Trunking*

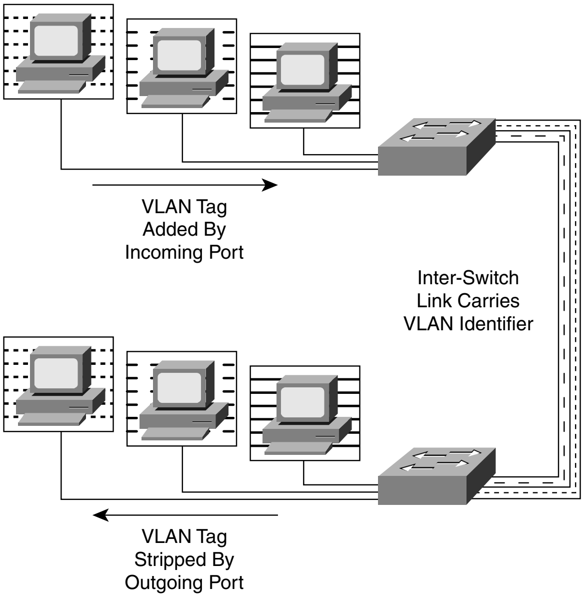

Two forms of trunking are used for Cisco switches on Ethernet networks:

- An IEEE industry standard called IEEE 802.1Q. This is a frame-tagging mechanism that adds a VLAN identifier to the frame by inserting a tag at Layer 2.
- Another form of trunking on Cisco switches is called Inter-Switch Link (ISL), which is a Cisco proprietary trunking mechanism. ISL uses a frame encapsulation method that adds a header to identify the VLAN.

802.1Q Trunking

The IEEE 802.1Q protocol interconnects VLANs between multiple switches, routers, and servers. With 802.1Q, a network administrator can define a VLAN topology to span multiple physical devices. If you examine Figure 3-6, you can see that VLANs 1, 2, and 3 are physically attached to different switches; however, because the trunk link carries traffic for all of these VLANs, all the users in a given VLAN are in the same broadcast domain.

Cisco switches support IEEE 802.1Q for FastEthernet and GigabitEthernet interfaces. An 802.1Q trunk link provides VLAN identification by adding a 4-byte tag to an Ethernet Frame as it leaves a trunk port. Because the frame has been changed, a new frame check sequence (FCS) must also be computed and added to the frame. Figure 3-7 shows a frame entering an access port and leaving a trunk port with a tag.

NOTE On Cisco switches, 802.1Q is usually referred to as *dot1Q* after the IEEE standard number.

Figure 3-7 *802.1Q Frame Tagging*

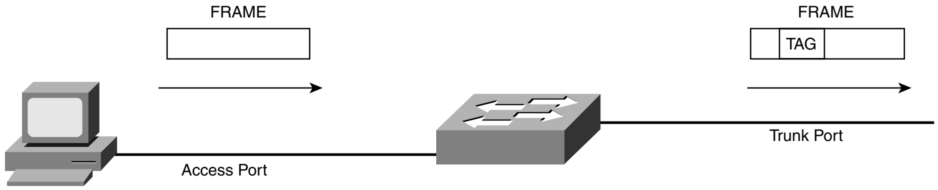

The 4-byte tag is inserted into the frame immediately following the source address field and is composed of two separate 2-byte sections—the Tag Protocol ID (TPID) field and the Tag Control Information (TCI) field.

The TPID field, for Ethernet frames, is always the hexadecimal value 8100 (0x8100). You might recall from Chapter 2 that any value over 05DC following the source address is an Ethertype field. The value 0x8100 tells an 802.1Q-compliant device that this is a tagged frame and to use the next 2 bytes for 802.1Q information.

In the next 2-byte TCI field, the first 3 bits of the TCI are referred to as the Priority bits. These bits indicate the priority of the frame for quality of service (QoS) reasons. For example, an IP telephone will mark any voice traffic with a priority of 5, indicating to a switch running QoS that the frame should be sent through the network as fast as possible. In some implementations, the priority bits are the information in the TCI that are most important to the user. For this reason, 802.1Q is sometimes referred to as 802.1p. These are the same frames with an emphasis on different part of the TCI.

The next bit in the TCI field is called a Canonical Format Identifier (CFI). This is a 1-bit field that, when off, indicates that the device should read the information in a field canonically (right-to-left or low-order bits first). The reason for this bit is that 802.1Q can be used for Token Ring or Ethernet frames. An Ethernet device reads canonically, but Token Ring devices read in a noncanonical form. For an Ethernet frame, this value will always be 0, but if the tag is in a Token Ring frame, it will be 1. For this reason, the CFI is sometimes referred to as the *Token Ring Encapsulation Flag*.

The last 12 bits in the CFI are the VLAN ID. This allows for the identification of 4096 unique VLANs. Figure 3-8 shows the components of the 802.1Q tag.

Figure 3-8 *802.1Q Tag Components*

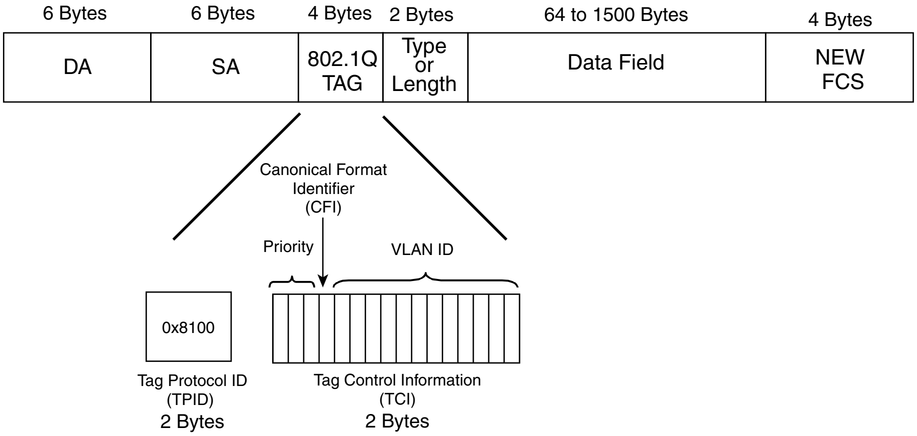

With 802.1Q, a trunk link can tag frames between devices that understand the protocol. This allows for multiple VLANs to exist on a single topology. Because 802.1Q is defined as a type of Ethernet frame, it does not require that every device on a link speaks the 802.1Q protocol. Because Ethernet is a shared media and more than two device could be connected on this media, all devices on the link must still be capable of communicating even if they do not speak the 802.1Q protocol. For this reason, 802.1Q also defines a Native VLAN. A trunk port on a switch is defined to be in a Native VLAN, and the 802.1Q trunk will not tag frames that are going out the port that came in on any port that belongs to the same VLAN

that is the Native VLAN on the switch. Any Ethernet device would be capable of reading frames for the Native VLANs. The Native VLAN is important on an 802.1Q trunk link. If both sides of the link do not agree on the Native VLAN, the trunk will not operate properly. Figure 3-9 shows how frames on the Native VLAN are not tagged out trunk links.

Figure 3-9 *Native VLAN*

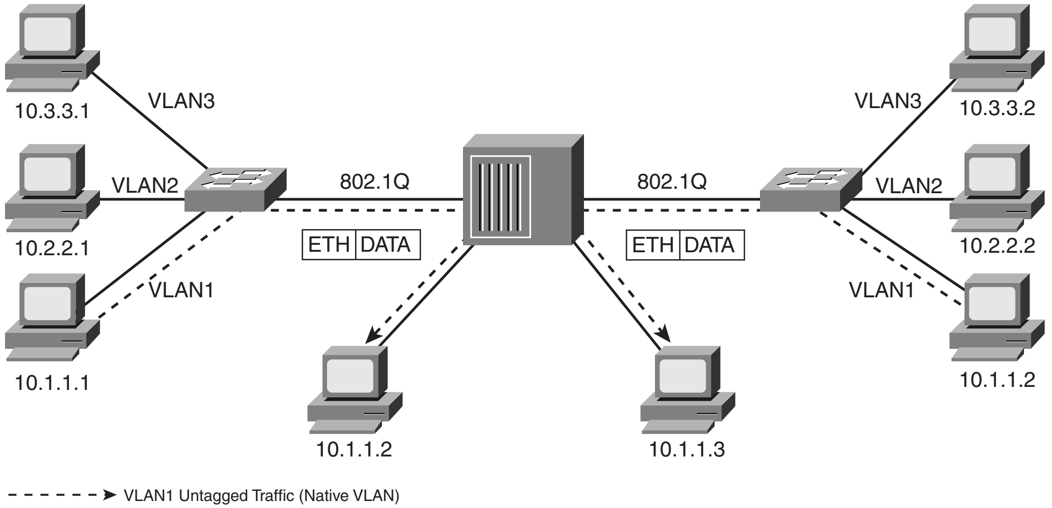

Inter-Switch Links

Inter-Switch Link (ISL) tagging accomplishes the same task as 802.1Q trunking but uses a different frame format. ISL trunks are Cisco proprietary and define only a point-to-point connection between two devices, typically switches. The name *Inter-Switch Link* hints at this design. ISL frame tagging uses a low-latency mechanism for multiplexing traffic from multiple VLANs on a single physical path. ISL has been implemented for connections among switches, routers, and network interface cards (NICs) used on nodes such as servers. To support the ISL feature, each connecting device must be ISL-configured. A router that is ISL-configured can allow inter-VLAN communications. A non-ISL device that receives ISL-encapsulated Ethernet frames will most likely consider them protocol errors because of the format and size of the frames.

ISL functions at Layer 2 of the OSI model like 802.1Q, but it differs by encapsulating the entire Layer 2 Ethernet frame inside an ISL header and trailer. Because ISL encapsulates the entire frame, it is protocol-independent and can carry any type of Layer 2 frame or upper-layer protocol between the switches. ISL has the following characteristics:

- Performed with application-specific integrated circuits (ASIC)
- Not intrusive to client stations; client does not see the ISL header
- Effective between switches, routers and switches, and switches and servers with ISL NICs

The following section addresses ISL tagging and ISL encapsulation.

ISL Encapsulation

ISL functions at OSI Layer 2 by encapsulating a data frame with a new (ISL) header and an additional (ISL) cyclic redundancy check (CRC). ISL-encapsulated frames are passed over trunk lines. ISL is protocol-independent because the data frame might carry any data-link protocol.

Ports configured as ISL trunks encapsulate each frame with a 26-byte ISL header and a 4-byte CRC before sending it out the trunk port. Because ISL technology is implemented in ASICs, frames are tagged with low latency. The ISL header supports 10 bits for ISL identification. Each bit can be one of two values, 2^{10}, allowing for the 1024 unique VLANs. The number of actual VLANs supported by a switch depends on the switch hardware. Figure 3-10 illustrates a typical ISL-encapsulated data frame.

Figure 3-10 *ISL Encapsulation*

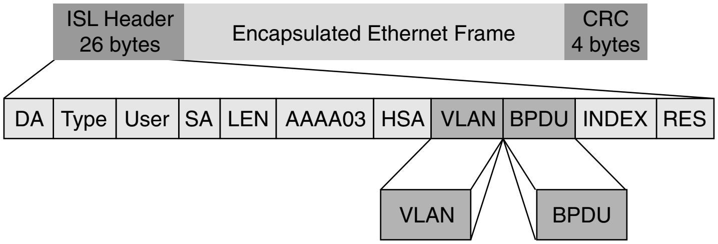

As illustrated in Figure 3-10, the ISL frame header contains the following information fields:

- **DA**—40-bit multicast destination address.
- **Type**—4-bit descriptor of the encapsulated frame types—Ethernet (0000), Token Ring (0001), FDDI (0010), and ATM (0011).
- **User**—4-bit descriptor used as the type field extension or to define Ethernet priorities. This is a binary value from 0, the lowest priority, to 3, the highest priority.
- **SA**—48-bit source MAC address of the transmitting Catalyst switch.
- **LEN**—16-bit frame-length descriptor minus DA type, user, SA, LEN, and CRC.
- **AAAA03**—Standard SNAP 802.2 LLC header.
- **HSA**—First 3 bytes of SA (manufacturer's ID or organizational unique ID).
- **VLAN**—15-bit VLAN ID. Only the lower 10 bits are used for 1024 VLANs.
- **BPDU**—1-bit descriptor identifying whether the frame is a Spanning Tree bridge protocol data unit (BPDU). Also set if the encapsulated frame is a Cisco Discovery Protocol (CDP) frame.

- **INDEX**—16-bit descriptor that identifies the transmitting port ID. Used for diagnostics.
- **RES**—16-bit reserved field used for additional information, such as Token Ring and Fiber Distributed Data Interface (FDDI) frame Frame Check (FC) field.

VLAN Trunking Protocol

To provide VLAN connectivity throughout the switched network, VLANs must be configured on each switch. If you are going to trunk VLAN10 from Switch A to Switch C through Switch B, as shown in Figure 3-11, VLAN10 must exist on Switch B even though none of the access ports on that switch are in VLAN10.

Figure 3-11 *Purpose for VTP*

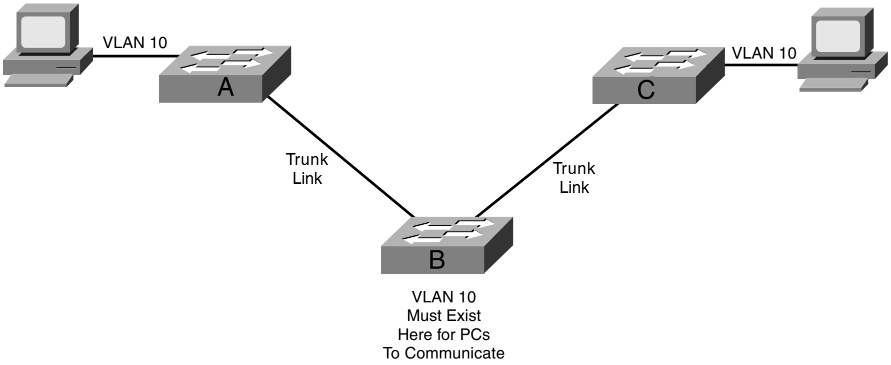

To ensure that a VLAN exists between every pair of Trunked switches, an administrator must manually create all the needed VLANs on each of the switches individually. Cisco's VLAN Trunking Protocol (VTP) provides an easier method for maintaining consistent VLAN configuration throughout the switched network.

VTP is a protocol used to distribute and synchronize identifying information about VLANs configured throughout a switched network. Configurations made to a single VTP server are propagated across trunk links to all connected switches in the network. VTP enables switched network solutions to scale to large sizes by reducing the network's manual configuration needs.

VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency throughout a common administrative domain by managing the additions, deletions, and name changes of VLANs across networks. VTP minimizes misconfigurations and configuration inconsistencies that can cause problems, such as duplicate VLAN names or incorrect VLAN-type specifications.

A VTP domain is one switch or several interconnected switches sharing the same VTP environment. A switch can be configured only in one VTP domain.

By default, a Catalyst switch is in the no-management-domain (or null domain) state until it is configured with a domain or receives an advertisement for a domain over a trunk link. Configuration changes made to the VLANs on a single VTP server switch are propagated across Trunk links to all trunk-connected switches in the network.

Figure 3-12 illustrates how VLAN configuration information is propagated from switch to switch.

Figure 3-12 *VTP Operation*

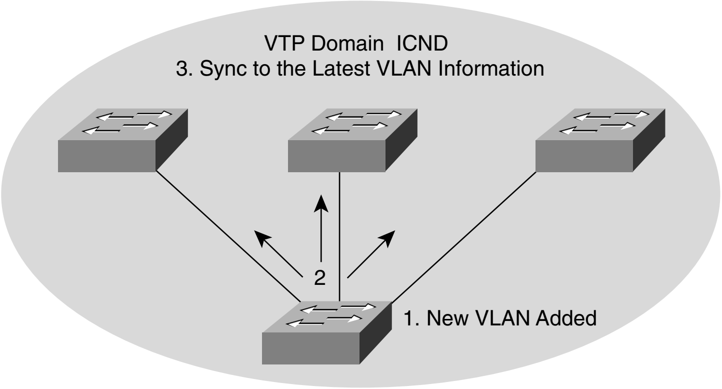

Figure 3-12 shows a VLAN added to the switched network. The steps illustrated in the figure are as follows:

1. A new VLAN is added. At this point, VTP makes your job easier.
2. The VTP advertisement is sent to the other switches in the VTP domain.
3. The new VLAN is added to the other switch configurations. The result is consistent VLAN configuration.

VTP Modes

VTP operates in one of three modes:

- Server mode
- Client mode
- Transparent mode

The default VTP mode is server mode, but VLANs are not propagated over the network until a management domain name is specified or learned and trunking has been established.

A Catalyst switch operating in the VTP server mode can create, modify, and delete VLANs and other configuration parameters for the entire VTP domain. In server mode, VLAN configurations are saved in the Catalyst nonvolatile random-access memory (NVRAM). When you make a change to the VLAN configuration on a VTP server, the change is propagated to all switches in the VTP domain. VTP messages are transmitted out all trunk connections, such as ISL.

A device operating as a VTP client cannot create, change, or delete VLANs. A switch in client mode does send VTP messages, however. A VTP client does not save VLAN configurations in nonvolatile memory.

In both client and server mode, the switches synchronize their VLAN configuration to that of the switch with the highest revision number in the VTP domain.

A switch operating in VTP transparent mode does not create VTP advertisements or synchronize its VLAN configuration with information received from other switches in the management domain. A switch in transparent mode forwards VTP advertisements received from other switches that are part of the same management domain. A switch configured in VTP transparent mode can create, delete, and modify VLANs, but the changes are not transmitted to other switches in the domain; they affect only the local switch. Table 3-1 offers a comparative overview of the three VTP modes.

Table 3-1 *VTP Modes*

Server Mode	Client Mode	Transparent Mode
Sends/forwards VTP advertisements.	Sends/forwards VTP advertisements.	Forwards VTP advertisements.
Synchronizes VLAN configuration information with other switches.	Synchronizes VLAN configuration information with other switches.	*Does not* synchronize VLAN configuration information with other switches.
VLAN configurations are saved in NVRAM.	VLAN configurations *are not* saved in NVRAM.	VLAN configurations are saved in NVRAM.
Catalyst switch can create VLANs.	Catalyst switch *cannot* create VLANs.	Catalyst switch can create VLANs.
Catalyst switch can modify VLANs.	Catalyst switch *cannot* modify VLANs.	Catalyst switch can modify VLANs.
Catalyst switch can delete VLANs.	Catalyst switch *cannot* delete VLANs.	Catalyst switch can delete VLANs.

When setting up VTP on a switch, choosing the appropriate mode is important. Because VTP is a simple and dangerous tool, it can overwrite VLAN configurations on some switches and create network problems. The next section further explains this phenomenon. Nevertheless, you must be aware that the mode you choose can eliminate the chance of these problems:

- Choose server mode for the switch that you will use to create, change, or delete VLANs. The server will propagate this information to other switches that are configured as servers or clients.
- Set client mode on any switch where you do not want to create, change, or delete VLANS.
- Use transparent mode on a switch that needs to pass VTP advertisements to other switches but also needs the capability to have its VLANs independently administered.

How VTP Works

VTP advertisements are flooded throughout the management domain every five minutes or whenever a change occurs in VLAN configurations. VTP advertisements are sent over a factory default VLAN (VLAN 1) using multicast frames. Included in a VTP advertisement is a configuration revision number. A higher configuration revision number indicates that the VLAN information being advertised is more current than the stored information.

A device that receives VTP advertisements must check various parameters before incorporating the received VLAN information.

First, the management domain name and the password, which can be configured to prevent unauthorized switches from altering the VTP domain, must match those configured in the local switch before information can be used.

Next, if the configuration revision number indicates that the message was created after the configuration currently in use, the switch overwrites its VLAN database with the advertised VLAN information. To reset the configuration revision number on a Catalyst switch, you must either change the switch mode to transparent then back to server or client with the command **vtp mode** [**server** | **client** | **transparent**] in global configuration mode, or change the VTP domain name and then set it back using the command **vtp domain** *name* in global configuration mode. Example 3-1 demonstrates changing the mode and then setting it back to reset the configuration revision number. The command **show vtp status** is executed before and after the change to show the configuration number being reset.

Example 3-1 *Resetting a Switches VTP Configuration Revision Number*

```
Switch#show vtp status
VTP Version                     : 2
Configuration Revision          : 5
Maximum VLANs supported locally : 250
Number of existing VLANs        : 10
VTP Operating Mode              : Server
VTP Domain Name                 : switch_domain_1
VTP Pruning Mode                : Disabled
```

continues

Example 3-1 *Resetting a Switches VTP Configuration Revision Number (Continued)*

```
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x1E 0xED 0x19 0x49 0x0F 0x37 0x65 0x64
Configuration last modified by 192.168.255.21 at 3-1-93 00:02:39
Local updater ID is 192.168.255.21 on interface Vl1 (lowest numbered VLAN interface
found)
Switch#config t
P2_2950(config)#vtp mode transparent
Setting device to VTP TRANSPARENT mode.
Switch(config)#vtp mode server
Setting device to VTP SERVER mode
Switch(config)#end

Switch#show vtp status
VTP Version                     : 2
Configuration Revision          : 0
Maximum VLANs supported locally : 250
Number of existing VLANs        : 10
VTP Operating Mode              : Server
VTP Domain Name                 : switch_domain_1
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x1E 0xED 0x19 0x49 0x0F 0x37 0x65 0x64
Configuration last modified by 192.168.255.21 at 3-1-93 00:02:39
Local updater ID is 192.168.255.21 on interface Vl1 (lowest numbered VLAN interface
found)
```

NOTE Underscores are used in the VTP domain name because a domain name cannot contain spaces.

One of the most critical components of VTP is the configuration revision number. Each time a VTP server modifies its VLAN information, it increments the configuration revision number by one. The VTP server then sends out a VTP advertisement with the new configuration revision number. If the configuration revision number being advertised is higher than the number stored on the other switches in the VTP domain, the other switches will overwrite their VLAN configurations with the new information being advertised. The configuration revision number in VTP transparent mode is always 0. Figure 3-13 illustrates how VTP operates in a switched network.

CAUTION The overwrite process would mean that the VTP server with the highest revision number determines the overall VLAN configuration for the domain. For example, if you deleted all VLANs on a VTP server and that server had the higher revision number, the other devices in the VTP domain would also delete their VLANs. This could create a loss of connectivity.

Figure 3-13 *VTP Operation*

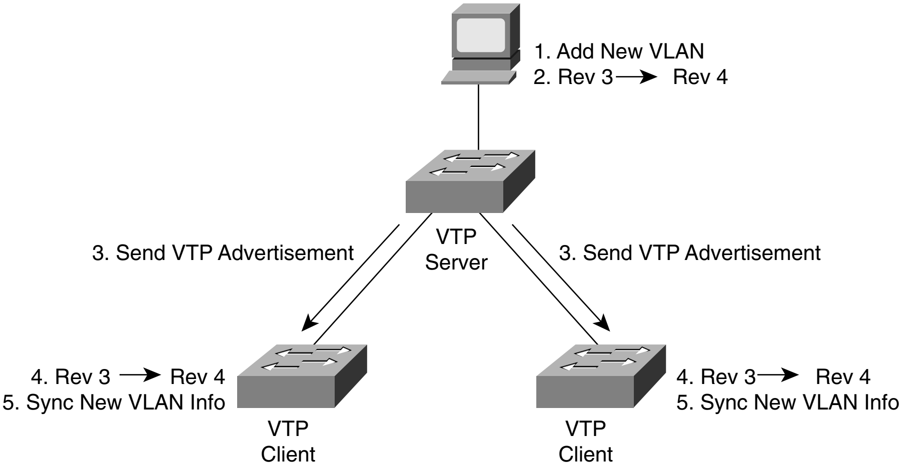

VTP Pruning

Because ISL trunk lines carry VLAN traffic for all VLANs by default, some traffic might be needlessly flooded across links that do not need to carry that traffic. VTP pruning uses VLAN advertisements to determine when a trunk connection is flooding traffic needlessly.

By default, a trunk connection carries traffic for all VLANs in the VTP management domain. Often, some switches in an enterprise network do not have local ports configured in each VLAN. In Figure 3-14, Switches 1 and 4 support ports statically configured in VLAN10. As illustrated, with VTP pruning enabled, when Station A sends a broadcast, the broadcast is flooded only toward any switch with ports assigned to VLAN10. As a result, broadcast traffic from Station A is not forwarded to Switches 3, 5, and 6 because traffic for VLAN10 has been pruned on the links indicated on Switches 2 and 4. Pruning must be enabled on one VTP server, and it will be propagated to all other switches in the VTP domain.

VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices.

NOTE Because VLAN1 is the management VLAN and is used for administrative functions such as VTP advertisements, it will not be pruned from a trunk line by VTP pruning.

Figure 3-14 *VTP Pruning*

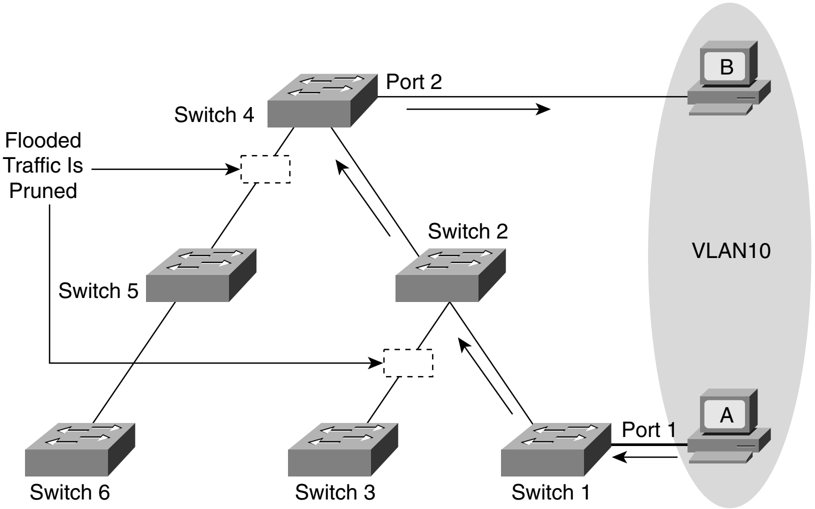

Per-VLAN Spanning Tree

One of the things that must be considered with VLANs is the function of the Spanning Tree Protocol (STP). STP is designed to prevent loops in a switch/bridged topology to eliminate the endless propagation of broadcast around the loop. With VLANs, there are multiple broadcast domains to be considered. Because each broadcast domain is like a unique bridged internetwork, you must consider how STP will operate.

The 802.1Q standard defines one unique Spanning Tree instance to be used by all VLANs in the network. STP runs on the Native VLAN so that it can communicate with both 802.1Q and non-802.1Q compatible switches. This single instance of STP is often referred to as 802.1Q Mono Spanning Tree or Common Spanning Tree (CST). A single spanning tree lacks flexibility in how the links are used in the network topology. Cisco implements a protocol known as Per-VLAN Spanning Tree Plus (PVST+) that is compatible with 802.1Q CST but allows a separate spanning tree to be constructed for each VLAN. There is only one active path for each spanning tree; however, in a Cisco network, the active path can be different for each VLAN.

NOTE The term Mono Spanning Tree is typically not used anymore because the IEEE 802.1s standard has now defined a Multiple Spanning Tree (MST) protocol that uses the same acronym.

Because a trunk link carries traffic for more than one broadcast domain and switches are typically connected together via trunk links, it is possible to define multiple Spanning Tree topologies for a given network. With PVST+, a root bridge and STP topology can be defined for each VLAN. This is accomplished by exchanging BPDUs for each VLAN operating on the switches. By configuring a different root or port cost based on VLANs, switches could utilize all the links to pass traffic without creating a bridge loop. Using PVST+, administrators can use ISL or 802.1Q to maintain redundant links and load balance traffic between parallel links using the Spanning Tree Protocol. Figure 3-15 shows an example of load balancing using PVST+.

Figure 3-15 *PVST Load Balancing*

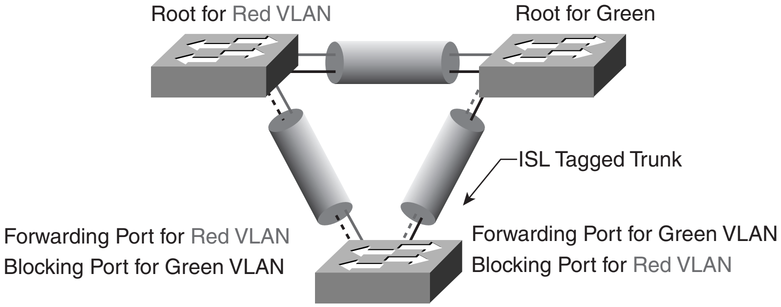
Root for Red VLAN
Root for Green
ISL Tagged Trunk
Forwarding Port for Red VLAN
Blocking Port for Green VLAN
Forwarding Port for Green VLAN
Blocking Port for Red VLAN

Cisco developed PVST+ to allow running several STP instances, even over an 802.1Q network by using a tunneling mechanism. PVST+ utilizes Cisco devices to connect to a Mono Spanning Tree zone, typically another vendor's 802.1Q-based network, to a PVST+ zone, typically a Cisco ISL-based network. No specific configuration is needed to achieve this. PVST+ provides support for 802.1Q trunks and the mapping of multiple spanning trees to the single spanning tree of standard 802.1Q switches running Mono Spanning Tree.

The PVST+ architecture distinguishes three types of regions:

- A PVST region (PVST switches using ISL only)
- A PVST+ region (PVST+ using ISL and/or 802.1Q between Cisco switches)
- A Mono Spanning Tree region (Common or Mono Spanning Tree using 802.1Q and exchanging BPDUs on the Native VLAN only between a Cisco and Non-Cisco switches using 802.1Q)

Each region consists of a homogenous type of switch. You can connect a PVST region to a PVST+ region using ISL ports. You can also connect a PVST+ region to a Mono Spanning Tree region using 802.1Q ports.

At the boundary between a PVST region and a PVST+ region, the mapping of Spanning Tree is one-to-one. At the boundary between a Mono Spanning Tree region and a PVST+ region, the Spanning Tree in the Mono Spanning Tree region maps to one PVST in the PVST+ region. The one it maps to is the CST. The CST is the PVST of the Native VLAN (VLAN 1 by default).

On a 802.1Q trunk, BPDUs can be sent or received only by the Native VLAN. Using PVST+, Cisco can send its PVST BPDUs as tagged frames using a Cisco multicast address as the destination. When a non-Cisco switch receives the multicast, it is flooded (but not interpreted as a BPDU, thus maintaining the integrity of CST). Because it is flooded, it will eventually reach Cisco switches on the other side of the CST domain. This allows the PVST fames to be tunneled through the MST region. Tunneling means that the BPDUs are flooded through the Mono Spanning Tree region along the single spanning tree present in the Mono Spanning Tree region.

PVST+ networks must be in a tree-like structure for proper STP operation.

Section 1 Quiz

1 Which of the following are valid Layer 2 (switchport) types?

A Static Access Port

B Trunk Port

C Common Spanning Tree Port

D Root Port

E Dynamic Access Port

2 True or False: A Catalyst 2950 cannot act as a VMPS server and it cannot have ports with dynamically assigned VLANs.

3 VLANs provide which of the following? (Choose all that apply.)

A Security

B Redundancy

C Segmentation

D Loop prevention

E Collision domains

F Flexibility

G All of the above

4 How many bytes of overhead are added to a frame for 802.1Q frame tagging and ISL frame tagging?

A 4 bytes ISL & 30 bytes 802.1Q

B 8 bytes ISL & 26 bytes 802.1Q

C 26 bytes ISL & 4 bytes 802.1Q

D 30 bytes ISL & 8 bytes 802.1Q

E 30 bytes ISL & 4 bytes 802.1Q

F None. The frame tagging is performed in Hardware ASICs.

5 Which of the following is true for VTP? (Choose the best answers.)

A VTP is required for proper VLAN operation.

B VTP eases VLAN creation for a switched network.

C Changes need to be made to only one VTP server in a domain.

D A switch in VTP transparent mode will increment its configuration revision number but will not synchronize with other switches.

E VTP will run only across trunk links.

F VTP requires a domain name to operate.

G The default VTP domain name for all Catalyst switches is Cisco.

H VTP is completely safe and will never cause problems in your network.

VLAN Configuration

This section discusses the guidelines for configuring VLANs on the Catalyst switch. You will learn the steps to configure VLANs, how to enable VTP domains, how to define a trunk, how to create a VLAN, and how to verify proper VLAN operation.

You should remember several facts before you begin VLAN configuration:

- The maximum number of VLANs that can operate on a switch is switch-dependent.
- VLAN1 is one of the factory default VLANs.
- Cisco Discovery Protocol (CDP) and VTP advertisements are sent on VLAN1.
- The switch must be in VTP server mode or transparent mode to create, add, or delete VLANs.

VLAN Configuration Guidelines

The Catalyst switches have a factory default configuration in which various default VLANs are preconfigured. One of the default VLANs is VLAN1, which is used for CDP and VTP advertisements. The VLAN1 interface on a switch is also in the default VLAN1. As you'll recall, the switch requires an IP address for management purposes—for example, to allow Telnet connections into the switch, or to use the Visual Switch Manager (VSM) via an HTTP browser to configure the switch.

Before you can create a VLAN, the switch must be in VTP server mode or VTP transparent mode. If you want to propagate the VLAN to other switches in the domain, use server mode.

VLAN Configuration Steps

Before you create VLANs, you must decide whether to use VTP to maintain global VLAN configuration information for your network.

To allow VLANs to span multiple Catalyst switches on a single link, you must configure trunks to interconnect the switches.

By default, a switch is in VTP server mode so that VLANs can be added, changed, or deleted. If the switch is set to VTP client mode, VLANs cannot be added, changed, or deleted from that switch.

VLAN membership on the switch ports is assigned manually on a port-by-port basis. When you assign switch ports to VLANs using this method, it is known as port-based, or static, VLAN membership.

The following sections elaborate on the details of the steps to configure VLANs.

VTP Configuration Guidelines

The default VTP configuration parameters for the Catalyst switch are as follows:

- VTP domain name: None
- VTP mode: Server
- VTP password: None
- VTP pruning: Disabled
- VTP trap: Disabled

The VTP domain name can be specified by the administrator or learned across a configured trunk line from a server with a domain name configured. By default, the domain name is not set.

By default, the switch is set to the VTP server mode.

A password can be set for the VTP management domain. The password entered must be the same for all switches in the domain. If you configure a VTP password, VTP does not function properly unless you assign the same password to each switch in the domain.

VTP pruning eligibility is one VLAN parameter advertised by the VTP protocol. Enabling or disabling VTP pruning on a VTP server propagates the change throughout the management domain. Enabling or disabling VTP pruning on a VTP server affects the entire management domain.

VTP trap is disabled by default. If you enable this feature, it causes an SNMP message to be generated every time a new VTP message is sent.

CAUTION When adding a new switch to an existing domain, you should verify that the configuration revision number for the switch is 0 to prevent the new switch from propagating incorrect VLAN information. Example 3-1, in the "How VTP Works" section, demonstrated one method for resetting the VTP configuration revision number on the new switch.

Configuring VTP

Use the **vtp** global configuration command to specify the operating mode, domain name, password, generation of traps, and pruning capabilities of VTP. The syntax for this command is as follows:

```
switch(config)# vtp { [mode {server | transparent | client}] [domain domain-name]
  [password password] [pruning {enable | disable}]}
```

To verify a recent configuration change, or to just view the VTP configuration information, use the **show vtp status** privileged EXEC command, as demonstrated in Example 3-2. Also displayed is the IP address of the device that last modified the configuration and a time stamp showing when the modification was made. VTP has two versions:

- VTP version 1 only supports Ethernet.
- VTP version 2 supports Ethernet and Token Ring.

Example 3-2 **show vtp status** *Output*

```
Switch#show vtp status
VTP Version                     : 2
Configuration Revision          : 5
Maximum VLANs supported locally : 250
Number of existing VLANs        : 10
VTP Operating Mode              : Server
VTP Domain Name                 : switch_domain_1
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x1E 0xED 0x19 0x49 0x0F 0x37 0x65 0x64
Configuration last modified by 192.168.255.21 at 3-1-93 00:02:39
Local updater ID is 192.168.255.21 on interface Vl1 (lowest numbered
  VLAN interface found)
```

Trunk Line Configuration

Use the command **switchport mode trunk** at the interface configuration mode to set a port to trunk. On the Catalyst 2950, this enables 802.1Q trunking. On other Cisco IOS Software-based switches, such as the 3550, 4500, or 6500, you will need to choose an encapsulation method before you can enable trunking. The command **switchport trunk encapsulation** [**isl** | **dot1q**] chooses an encapsulation mode.

The Catalyst IOS switches also support Dynamic Trunking Protocol (DTP), which manages automatic trunk negotiation. The **switchport mode** command specifies a Layer 2 ports operation:

```
switch(config-if)# switchport mode [trunk | access | dynamic
  [desirable | auto | nonegotiate]]
```

The options for the **switchport mode** command are as follows:

- **trunk**—Configures the port to permanent trunk mode and negotiates with the connected device on the other side to convert the link to trunk mode. If multiple trunk encapsulations are available, the encapsulation must be chosen before this command will work.
- **access**—Disables port trunk mode and negotiates with the connected device to convert the link to nontrunk. This port will belong to only the configured access VLAN.
- **dynamic desirable**—Triggers the port to negotiate the link from nontrunk to trunk mode. The port negotiates to a trunk port if the connected device is in the **trunk, dynamic desirable**, or **dynamic auto** state. Otherwise, the port becomes a nontrunk port. This is the default for IOS switch ports
- **dynamic auto**—Enables the port to become a trunk only if the connected device has the state set to **trunk** or **dynamic desirable**.
- **nonnegotiate**—Configures the port to permanent trunk mode. No negotiation takes place with the partner. The other side must be **trunk** or **nonegotiate** for the trunk to work. You must also specify the encapsulation before choosing this mode.

Verifying Trunk Line Configuration

To verify a trunk configuration, use the command **show interface switchport** or **show interface trunk** privileged EXEC command. The syntax for the **show interface switchport** and privileged EXEC command is as follows:

```
switch(config)# show interface [type module/port] switchport
```

The syntax for the **show interface trunk** and privileged EXEC command is as follows:

```
switch(config)# show interface [type module/port] trunk
```

These commands display the trunk parameters, as demonstrated in Example 3-3.

Example 3-3 **show interface trunk** *and* **show interface switchport** *Output*

```
Switch#show interface trunk

Port      Mode         Encapsulation  Status        Native vlan
Fa0/1     on           802.1q         trunking      1
Gi0/1     on           802.1q         trunking      1

Port      Vlans allowed on trunk
```

Example 3-3 **show interface trunk** *and* **show interface switchport** *Output (Continued)*

```
Fa0/1     1-4094
Gi0/1     1-4094

Port      Vlans allowed and active in management domain
Fa0/1     1,101,202,303,404,505
Gi0/1     1,101,202,303,404,505

Port      Vlans in spanning tree forwarding state and not pruned
Fa0/1     1,101,202,303,404,505
Gi0/1     1,101,202,303,404,505

Switch#show interfaces fastEthernet 0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001

Protected: false

Voice VLAN: none (Inactive)
Appliance trust: none
```

Adding a VLAN

Use the **vlan** global configuration command to configure a VLAN. The syntax for the **vlan** global configuration command is as follows:

```
Switch(config)#vlan number
Switch(config-vlan)#[name | mtu | shutdown | exit]
```

Each VLAN has a unique four-digit ID that can be a number from 0001 to 4096. To add a VLAN to the VLAN database, assign a number and name to the VLAN. After creating the VLAN, you will be in VLAN configuration mode. In this mode, use the **name** command to give the VLAN a name. VLAN1, VLAN1002, VLAN1003, VLAN1004, and VLAN1005 are the factory default VLANs. These VLANs exist on all Catalyst switches and are used as default VLANs for other topologies, such as Token Ring and FDDI. None of the default VLANs can be modified or deleted.

To add an Ethernet VLAN, you must specify at least a VLAN number. If no VLAN name is entered for the VLAN, the default is to append the VLAN number to the word *VLAN*. For example, *VLAN0404* could be a default name for *VLAN404* if no name is assigned.

Remember that to add, change, or delete VLANs, the switch must be in VTP server or transparent mode.

Verifying a VLAN/Modifying VLAN Parameters

When the VLAN is configured, the parameters for that VLAN should be confirmed to ensure validity. To verify the VLAN's parameters, use the **show vlan id** *vlan#* privileged EXEC command to display information about a particular VLAN. Use **show vlan** to show all configured VLANs.

The **show vlan** command output in Example 3-4 also shows which switch ports are assigned to the VLAN.

Example 3-4 **show vlan** *Output*

```
Switch#show vlan

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/2, Fa0/3, Fa0/4, Fa0/6
                                                Fa0/7, Fa0/8, Fa0/9, Fa0/10
                                                Fa0/11, Fa0/12, Gi0/2
101  VLAN0101                         active
202  VLAN0202                         active
303  VLAN0303                         active
404  VLAN0404                         active
505  VLAN0505                         active
986  VLAN0986                         active
1002 fddi-default                     active
1003 token-ring-default               active
1004 fddinet-default                  active
1005 trnet-default                    active

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
1    enet  100001     1500  -      -      -        -    -        0      0
101  enet  100101     1500  -      -      -        -    -        0      0
202  enet  100202     1500  -      -      -        -    -        0      0
303  enet  100303     1500  -      -      -        -    -        0      0

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
404  enet  100404     1500  -      -      -        -    -        0      0
505  enet  100505     1500  -      -      -        -    -        0      0
986  enet  100986     1500  -      -      -        -    -        0      0
1002 fddi  101002     1500  -      -      -        -    -        0      0
1003 tr    101003     1500  -      -      -        -    -        0      0
```

Example 3-4 **show vlan** *Output (Continued)*

```
1004 fdnet 101004     1500  -      -      -        ieee -        0      0
1005 trnet 101005     1500  -      -      -        ibm  -        0      0

Remote SPAN VLANs
------------------------------------------------------------------------------


Primary Secondary Type             Ports
------- --------- ---------------- ------------------------------------------
```

Other VLAN parameters shown in Example 3-4 include the following:

- Type (default is Ethernet)
- Security Association ID (SAID), which is used for the FDDI trunk
- Maximum transmission unit (MTU, where the default is 1500 for Ethernet VLAN)
- Other parameters used for Token Ring or FDDI VLANs

To modify an existing VLAN parameter (such as the VLAN name), use the same command syntax used to add a VLAN.

In Example 3-5, the VLAN name for VLAN986 is changed to CSR_VLAN.

Example 3-5 *Change VLAN Name*

```
Switch# config t
Enter configuration commands, one per line. End with CNTL/Z
Switch(config)#vlan 986
Switch(config-vlan)#name CSR_VLAN
```

Use the **show vlan id 986** command, as demonstrated in Example 3-6, to verify the change.

Example 3-6 *Verify VLAN Change*

```
Switch# show vlan id 986

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
986  CSR_VLAN                         active    Fa0/1, Gi0/1

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
986  enet  100986     1500  -      -      -        -    -        0      0

Remote SPAN VLAN
----------------
Disabled

Primary Secondary Type             Ports
------- --------- ---------------- ------------------------------------------
```

Assigning Ports to a VLAN

After creating a VLAN, you can statically assign a port or a number of ports to that VLAN. A port can belong to only one VLAN at a time.

Configure the VLAN port assignment from the interface configuration mode using the interface command **switchport access vlan** number, as shown in the following syntax:

```
Switch(config-if)#switchport access vlan [1-4096 | dynamic]
```

dynamic means that the Catalyst switch queries a VMPS for VLAN information based on a MAC address. A number in the range of 1 to 4096 would represent the VLAN assignment for the port.

By default, all ports are members of the default VLAN—VLAN1.

Use the **show vlan brief** privileged EXEC command to display the VLAN assignment for all switch ports, as demonstrated in Example 3-7.

Example 3-7 *Displaying VLAN Assignments and Membership for All Switch Ports*

```
Switch#show vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/2, Fa0/3, Fa0/4, Fa0/6
                                                Fa0/7, Fa0/8, Fa0/9, Fa0/10
                                                Fa0/11, Fa0/12, Gi0/2
101  VLAN0101                         active
202  VLAN0202                         active
303  VLAN0303                         active
404  VLAN0404                         active
505  VLAN0505                         active
986  CSR_VLAN                         active
1002 fddi-default                     active
1003 token-ring-default               active
1004 fddinet-default                  active
1005 trnet-default                    active
```

Displaying Spanning Tree Protocol Configuration Status

Use the **show spanning-tree** privileged EXEC command to display the switch's Spanning Tree Protocol configuration status, as demonstrated in Example 3-8. The basic syntax for the **show spanning-tree** privileged EXEC command is as follows:

```
Switch# show spanning-tree [vlan number]
```

Example 3-8 **show spanning-tree** *Output*

```
Switch# show spanning-tree vlan 1
VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority    0
             Address     0005.00a9.2401
```

Example 3-8 **show spanning-tree** *Output (Continued)*

```
             Cost        8
             Port        13 (GigabitEthernet0/1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32769  (priority 32768 sys-id-ext 1)
             Address     000b.5f2a.5a40
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Port ID                     Designated                 Port ID
Name             Prio.Nbr      Cost Sts      Cost Bridge ID             Prio.Nbr
---------------- -------- --------- --- --------- -------------------- --------
Fa0/1            128.1          100 FWD         8 32769 000b.5f2a.5a40 128.1
Gi0/1            128.13           4 FWD         4 32768 0005.3104.c000  32.65
```

Example 3-8 displays various spanning tree information for VLAN1, including the following:

- Port Fa0/1 and G0/1 are in the forwarding state for VLAN1.
- The root bridge for VLAN1 has a bridge priority of 0 with a MAC address of 0005.00a9.2401.
- The switch is running the IEEE 802.1D Spanning Tree Protocol.

Recall that a Catalyst switch can support a separate Spanning Tree instance per VLAN. This allows for load balancing between switches. For example, one switch can be the root for VLAN1, and another switch can be the root for VLAN2.

VLAN Command Summary

Table 3-2 lists the commands covered in this chapter and briefly describes each command's function.

Table 3-2 *VLAN Command Summary*

Command	Description
vtp mode [**server** | **client** | **transparent**]	In global configuration mode, this command sets the operational VTP mode for the switch. The default is **server**.
vtp domain *name*	In global configuration mode, this command assigns a VTP domain name, which allows the switch to send VTP advertisements out trunk links. The default is NULL, which would allow a switch to join the first domain it received an update from.
show vtp status	Displays VTP status information including configuration revision number, domain name, and switch mode.

continues

Table 3-2 *VLAN Command Summary (Continued)*

switchport mode [**trunk** | **access** | **dynamic** [**desirable** | **auto** | **nonegotiate**]]	In interface configuration mode, this configures the behavior of the interface. **Trunk** mode will force frame tagging. **Dynamic** mode can become a trunk if it negotiates with the other side of the link. **Access** mode is a nontrunk port.
switchport trunk encapsulation [**isl** | **dot1q**]	Used in interface configuration mode to specify a trunking protocol. For some switches, before you can set an interface to **trunk** mode, you must first specify the encapsulation.
show interface [*type module/port*] **trunk**	Displays trunking information about the active or specified trunk links on the switch.
show interface [*type module/port*] **switchport**	Displays Layer 2 configuration and operational parameters of the switch. This includes VLAN membership and trunking status.
vlan *number*	In global configuration mode, this command defines a VLAN and puts the switch into VLAN configuration mode. In VLAN configuration mode, commands such as **name** can be used to further define the VLAN.
show vlan [**id** *vlan#*]	Displays VLAN information. The **id** option allows you to specify a particular VLAN.
switchport access vlan [*1-4096* | **dynamic**]	In interface configuration mode, this command assigns an **access** port to a VLAN or makes it a dynamic port.
show vlan brief	Displays a brief table of the VLANs, including the port membership for each VLAN.
show spanning-tree [**vlan** *number*]	Displays Spanning Tree information for the switch or a VLAN if the **vlan** option is used.

Section 2 Quiz

1 A VLAN can be created or modified on a switch in which of the following VTP modes? (Choose all that apply.)

A Server

B Access

C Client

D Transparent

E Root

2 Regarding VTP configuration revision numbers, which of the following statements are true? (Choose all that apply.)

A A transparent switch will always have a higher configuration revision number than any other switch on the network.

B VTP configuration revision numbers are changed on a switch when a VLAN is created, deleted, or modified.

C If a switch with a higher configuration revision number is added to an existing network with the same VTP domain name, it will have no effect on the VLANs on all the functioning switches.

D VTP configuration revision numbers can be reset to 0 by changing the VTP to transparent mode and then back to server or client.

E You can view a switch's current VTP configuration revision number by issuing the command **show vtp status**.

3 Choose the commands that force an IOS switch to perform trunking on a FastEthernet interface 0/12. (Choose the best answer.)

A **set trunk on**

B **interface Fa0/12 trunk on**

C **switchport mode trunk**

D **interface Fa0/12 mode trunk**

E **interface Fa0/12 then switchport mode trunk**

4 Which of the following is the default mode for a Layer 2 port on an IOS switch?

A **switchport mode access**

B **switchport mode dynamic auto**

C **switchport mode nonegotiate**

D **switchport mode dynamic desirable**

E **switchport mode trunk**

5 Which of the following commands you can use to see which VLAN a port is assigned to? (Choose all that apply.)

A **show interface trunk**

B **show interface** *type slot/port*

C **show vtp status**

D **show interface status**

E **show vlan brief**

F **show interface** *type slot/port* **switchport**

Summary

This chapter discussed how VLANs operate to provide more effective networks by controlling broadcasts in your network. To configure VLANs on a Catalyst switch, you must first configure VTP to administer VLANs. Therefore, you learned how VTP operates and how it is configured. You also learned how to create a trunk link to carry all VLAN traffic, and how to configure a VLAN. Finally, this chapter discussed the verification of Spanning Tree operations, including the following:

- How VLANs operate
- How to configure VTP
- How to configure a trunk
- How to configure a VLAN
- How to verify Spanning Tree operations

Case Study

Now that Ann has used switches to segment the network using switches, the network performance has noticeably improved. However, some of the servers are having some CPU utilization issues. After some research by the vendor who installed the servers, it has been determined that the problem is the amount of broadcast traffic. It seems that one of the servers runs an application that uses broadcasts to locate and poll all of its clients on the network. These broadcasts are affecting both servers and clients throughout the network, but it is more noticeable on the servers. Because of this, Ann has decided to implement VLANs. Based on the following requirements, what steps should Ann take in creating her VLANs? Figure 3-16 shows the layout of the switched network and location of the servers.

Figure 3-16 *International Widgets Ltd. Switched Network Diagram*

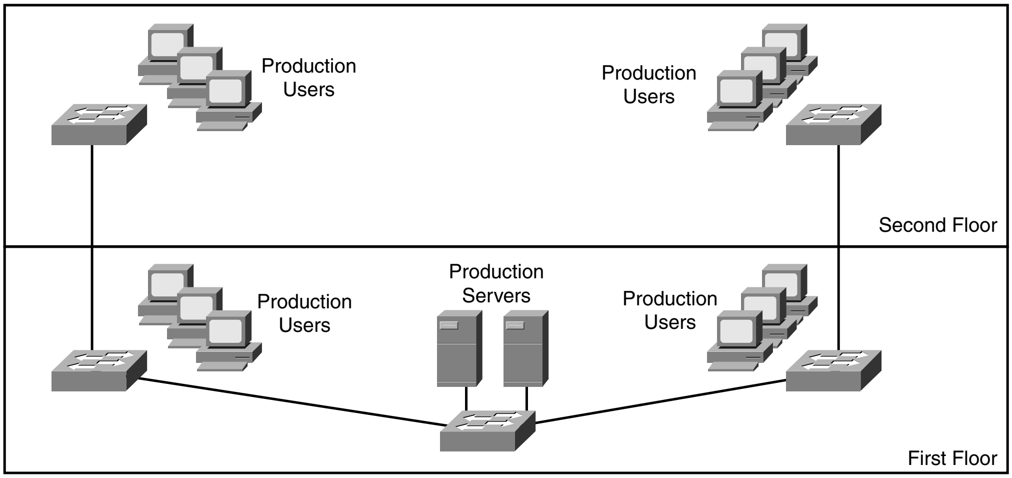

Ann has five servers. One server for production uses an all network broadcast to communicate with its clients. Those clients are located on both floors of the building, as shown in Figure 3-16. Of the other four servers, all use TCP/IP to communicate with various departments all over the company. It has been decided that for clients not using the production server, PCs and servers will be placed in a VLAN base at the location:

1 How many VLANs will Ann need and where will they need to be located in relation to the switches?

2 Do any of the switches have multiple VLANs on them? If so, what will Ann need to configure to ensure that multiple VLANs can pass between the switches?

3 In the future, Ann might need to create VLANs that will need to be used on some or all of the switches. To ensure that all VLANs exist on all trunked switches, what should Ann do?

Review Questions

1 VLANs allow for the creation of what in switched networks?

2 What are the two types of VLAN port assignments?

3 What type of port is capable of carrying all VLAN traffic?

4 What mechanism is used by switches to provide inter-switch communication between devices about which VLAN a packet originated from?

5 What is the purpose of VTP?

6 What is the default VTP mode for a Catalyst Switch?

7 Assume that a Catalyst switch is being added to your network. The switch needs to learn VLANs from the other switches in the network. You are not sure of the current VTP configuration and are fearful that it might overwrite your current VLAN information. How could you prevent the switch from accidentally overwriting the VLANs in your VTP domain?

8 What is unique about the Native VLAN on an IEEE 802.1Q trunk link?

9 List all the steps required to configure a VLAN on a Catalyst switch port.

10 Which command would you use to view the Spanning Tree configuration for VLAN9 on a Catalyst switch?