CHAPTER 5

Creating and Maintaining VLANs

A virtual LAN (VLAN) is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to stations in the VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a router or bridge as shown in Figure 5-1. Because a VLAN is considered a separate logical network, it contains its own bridge Management Information Base (MIB) information and can support its own implementation of the Spanning Tree Protocol (STP).

This chapter describes how to create and maintain VLANs through the Cluster Management software and the command-line interface (CLI). It contains the following information:

- How to configure static-access ports without having the VLAN Trunk Protocol (VTP) database globally propagate VLAN configuration information.
- How VTP works and how to configure its domain name, modes, and version.
- How to add, modify, and remove VLANs with different media characteristics to and from the VTP database.
- How to configure Fast Ethernet and Gigabit Ethernet VLAN trunks on a switch. The switch supports IEEE 802.1Q trunking standards for transmitting VLAN traffic. This section describes how to configure the allowed-VLAN list, the native VLAN for untagged traffic, and two methods of load sharing.
- How to configure IEEE 802.1p class of service (CoS) port priorities for port forwarding untagged frames. You assign CoS to certain types of traffic to give them priority over other traffic.

Figure 5-1 VLANs as Logically Defined Networks

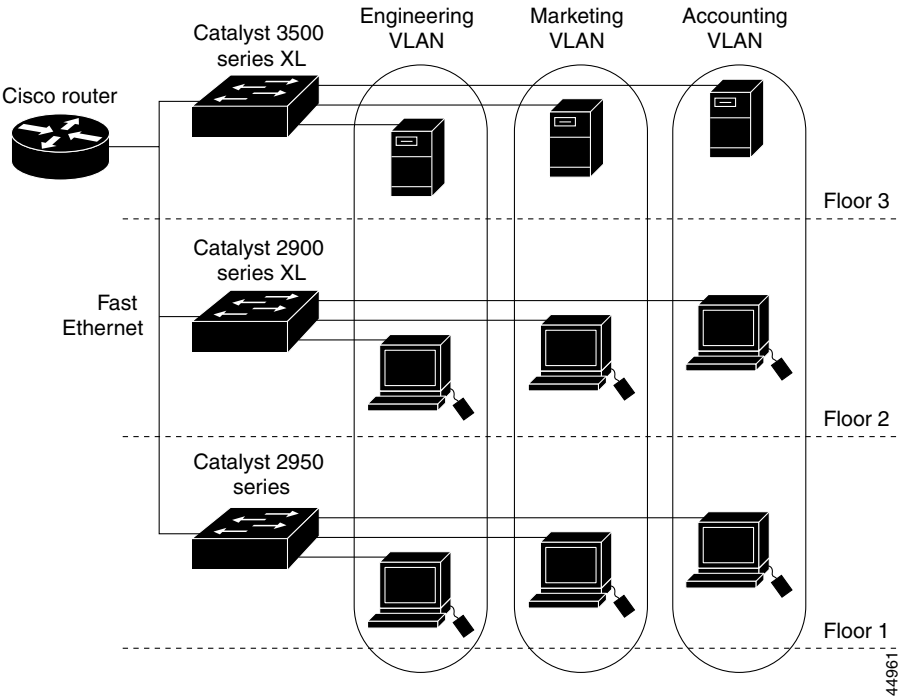

Number of Supported VLANs

Table 5-1 lists the number of supported VLANs on Catalyst 2950 switches.

Table 5-1 Number of Supported VLANs

Catalyst Switch	Number of Supported VLANs	Trunking Supported?
2950 switches with 16 MB of DRAM	64	Yes

VLANs are identified with a number between 1 and 1001. Regardless of the switch model, only 64 STP instances are supported.

The switches in Table 5-1 support IEEE 802.1Q trunking methods for transmitting VLAN traffic over 100BaseT, 100BaseFX, and Gigabit Ethernet ports.

VLAN Port Membership Modes

You configure a port to belong to a VLAN by assigning a membership mode that determines the kind of traffic the port carries and the number of VLANs it can belong to. Table 5-2 lists the membership modes and characteristics.

Table 5-2 Port Membership Modes

Membership Mode	VLAN Membership Characteristics
Static-access	A static-access port can belong to one VLAN and is manually assigned. By default, all ports are static-access ports assigned to VLAN 1.
Trunk (IEEE 802.1Q)	A trunk is a member of all VLANs in the VLAN database by default, but membership can be limited by configuring the allowed-VLAN list. VTP maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP exchanges VLAN configuration messages with other switches over trunk links.

When a port belongs to a VLAN, the switch learns and manages the addresses associated with the port on a per-VLAN basis. For more information, see the "Managing the MAC Address Tables" section on page 4-49.

VLAN Membership Combinations

You can configure your switch ports in various VLAN membership combinations as listed in Table 5-3.

Table 5-3 VLAN Combinations

Port Mode	VTP Required?	Configuration Procedure	Comments
Static-access ports	No	"Assigning Static-Access Ports to a VLAN" section on page 5-5	If you do not want to use VTP to globally propagate the VLAN configuration information, you can assign a static-access port to a VLAN and set the VTP mode to transparent to disable VTP.
Static-access and trunk ports	Recommended	"CLI: Configuring VTP Server Mode" section on page 5-14 Add, modify, or remove VLANs in the database as described in the "Configuring VLANs in the VTP Database" section on page 5-24 "CLI: Assigning Static-Access Ports to a VLAN" section on page 5-28 "Configuring a Trunk Port" section on page 5-31	Make sure to configure at least one trunk port on the switch and that this trunk port is connected to the trunk port of a second switch. Some restrictions apply to trunk ports. For more information, see the "Trunks Interacting with Other Features" section on page 5-30. You can change the VTP version on the switch. You can define the allowed-VLAN list and configure the native VLAN for untagged traffic on the trunk port.

Clusters, VLAN Membership, and the Management VLAN

This software release supports the grouping of switches into a cluster that can be managed as a single entity. The command switch is the single point of management for the cluster and cluster members.

Links among a command switch, cluster members, and candidate switches must be through ports that belong to the management VLAN. By default, the management VLAN is VLAN 1. If you are using SNMP or the Cluster Management Suite (CMS) to manage the switch, ensure that the port through

which you are connected to a switch is in the management VLAN. For information on configuring the management VLAN, see the "Changing the Management VLAN" section on page 3-34.

If you are configuring VLANs on a member switch, you might need to enter an extra command from the command-switch CLI to access the member switch. When configuring port parameters, for example, you can use the privileged EXEC **rcommand** command and the number of the member switch to display the member-switch CLI. Once you have accessed the member switch, command mode changes, and IOS commands operate as usual. Enter **exit** on the member switch in privileged EXEC mode to return to the command-switch CLI.

For more information about the **rcommand** command, refer to the *Catalyst 2950 Desktop Switch Command Reference*.

Assigning Static-Access Ports to a VLAN

By default, all ports are static-access ports assigned to the management VLAN, VLAN 1.

You can assign a static-access port to a VLAN without having VTP globally propagate VLAN configuration information (VTP is disabled). To assign a VLAN, you access the VLAN Membership window (Figure 5-2) by selecting **VLAN > VLAN Membership** from the menu bar and clicking the Assign VLANs tab.

Figure 5-2 VLAN Membership: Assign VLANs Tab

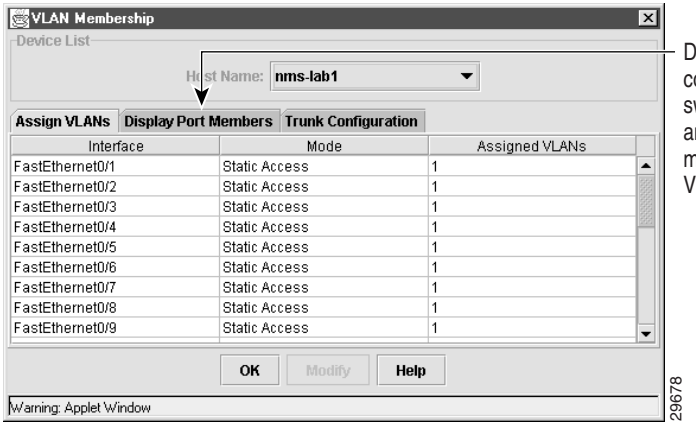

You configure the switch for VTP transparent mode, which disables VTP, by selecting **VLAN > VTP Management** from the menu bar and clicking the VTP Configuration tab (Figure 5-3).

You can also assign the port through the CLI on standalone, command, and member switches. If you are assigning a port on a cluster member to a VLAN, first log in to the member switch by using the privileged EXEC **rcommand** command. For more information on how to use this command, refer to the *Catalyst 2950 Desktop Switch Command Reference*.

Using the VLAN Trunk Protocol

VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP minimizes misconfigurations and configuration inconsistencies that can cause several problems, such as duplicate VLAN names, incorrect VLAN-type specifications, and security violations.

Before you create VLANs, you must decide whether to use VTP in your network. Using VTP, you can make configuration changes centrally on a single switch, such as a Catalyst 2950, 2900 XL, or 3500 XL switch, and have those changes automatically communicated to all the other switches in the network. Without VTP, you cannot send information about VLANs to other switches.

The VTP Domain

A VTP domain (also called a VLAN management domain) consists of one switch or several interconnected switches under the same administrative responsibility. A switch can be in only one VTP domain. You make global VLAN configuration changes for the domain by using the CLI, Cluster Management software, or Simple Network Management Protocol (SNMP).

By default, a Catalyst 2950, 2900 XL, or 3500 XL switch is in the no-management-domain state until it receives an advertisement for a domain over a trunk link (a link that carries the traffic of multiple VLANs) or until you configure a domain name. The default VTP mode is server mode, but VLAN information is not propagated over the network until a domain name is specified or learned.

If the switch receives a VTP advertisement over a trunk link, it inherits the domain name and configuration revision number. The switch then ignores advertisements with a different domain name or an earlier configuration revision number.

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 advertisements are sent over all trunk connections, including IEEE 802.1Q.

If you configure a switch for VTP transparent mode, you can create and modify VLANs, but the changes are not transmitted to other switches in the domain, and they affect only the individual switch.

For domain name and password configuration guidelines, see the "Domain Names" section on page 5-10.

VTP Modes and VTP Mode Transitions

You can configure a supported switch to be in one of the VTP modes listed in Table 5-4:

Table 5-4 VTP Modes

VTP Mode	Description
VTP server	In this mode, you can create, modify, and delete VLANs and specify other configuration parameters (such as VTP version) for the entire VTP domain. VTP servers advertise their VLAN configurations to other switches in the same VTP domain and synchronize their VLAN configurations with other switches based on advertisements received over trunk links. In VTP server mode, VLAN configurations are saved in nonvolatile RAM. VTP server is the default mode.
VTP client	In this mode, a VTP client behaves like a VTP server, but you cannot create, change, or delete VLANs on a VTP client. In VTP client mode, VLAN configurations are saved in nonvolatile RAM.
VTP transparent	In this mode, VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, transparent switches do forward VTP advertisements that they receive from other switches. You can create, modify, and delete VLANs on a switch in VTP transparent mode. In VTP transparent mode, VLAN configurations are saved in nonvolatile RAM, but they are not advertised to other switches.

The "VTP Configuration Guidelines" section on page 5-10 provides tips and caveats for configuring VTP.

VTP Advertisements

Each switch in the VTP domain sends periodic global configuration advertisements from each trunk port to a reserved multicast address. Neighboring switches receive these advertisements and update their VTP and VLAN configurations as necessary.

Note Because trunk ports send and receive VTP advertisements, you must ensure that at least one trunk port is configured on the switch and that this trunk port is connected to the trunk port of a second switch. Otherwise, the switch cannot receive any VTP advertisements.

VTP advertisements distribute the following global domain information in VTP advertisements:

- VTP domain name
- VTP configuration revision number
- Update identity and update timestamp
- MD5 digest

VTP advertisements distribute the following VLAN information for each configured VLAN:

- VLAN ID
- VLAN name
- VLAN type
- VLAN state
- Additional VLAN configuration information specific to the VLAN type

VTP Version 2

VTP version 2 supports the following features not supported in version 1:

- Token Ring support—VTP version 2 supports Token Ring LAN switching and VLANs (Token Ring Bridge Relay Function [TrBRF] and Token Ring Concentrator Relay Function [TrCRF]). For more information about Token Ring VLANs, see the "VLANs in the VTP Database" section on page 5-19.
- Unrecognized Type-Length-Value (TLV) support—A VTP server or client propagates configuration changes to its other trunks, even for TLVs it is not able to parse. The unrecognized TLV is saved in nonvolatile RAM when the switch is operating in VTP server mode.
- Version-Dependent Transparent Mode—In VTP version 1, a VTP transparent switch inspects VTP messages for the domain name and version and forwards a message only if the version and domain name match. Because only one domain is supported, VTP version 2 forwards VTP messages in transparent mode without checking the version and domain name.
- Consistency Checks—In VTP version 2, VLAN consistency checks (such as VLAN names and values) are performed only when you enter new information through the CLI, the Cluster Management software, or SNMP. Consistency checks are not performed when new information is obtained from a VTP message or when information is read from nonvolatile RAM. If the digest on a received VTP message is correct, its information is accepted without consistency checks.

VTP Configuration Guidelines

The following sections describe the guidelines you should follow when configuring the VTP domain name, password, and the VTP version number.

Domain Names

When configuring VTP for the first time, you must always assign a domain name. In addition, all switches in the VTP domain must be configured with the same domain name. Switches in VTP transparent mode do not exchange VTP messages with other switches, and you do not need to configure a VTP domain name for them.

Caution Do not configure a VTP domain if all switches are operating in VTP client mode. If you configure the domain, it is impossible to make changes to the VLAN configuration of that domain. Therefore, make sure you configure at least one switch in the VTP domain for VTP server mode.

Passwords

You can configure a password for the VTP domain, but it is not required. All domain switches must share the same password. Switches without a password or with the wrong password reject VTP advertisements.

Caution The domain does not function properly if you do not assign the same password to each switch in the domain.

If you configure a VTP password for a domain, a Catalyst 2950, 2900 XL, or 3500 XL switch that is booted without a VTP configuration does not accept VTP advertisements until you configure it with the correct password. After the configuration, the switch accepts the next VTP advertisement that uses the same password and domain name in the advertisement.

If you are adding a new switch to an existing network that has VTP capability, the new switch learns the domain name only after the applicable password has been configured on the switch.

VTP Version

Follow these guidelines when deciding which VTP version to implement:

- All switches in a VTP domain must run the same VTP version.
- A VTP version 2-capable switch can operate in the same VTP domain as a switch running VTP version 1 if version 2 is disabled on the version 2-capable switch (version 2 is disabled by default).

- Do not enable VTP version 2 on a switch unless all of the switches in the same VTP domain are version-2-capable. When you enable version 2 on a switch, all of the version-2-capable switches in the domain enable version 2. If there is a version 1-only switch, it will not exchange VTP information with switches with version 2 enabled.
- If there are Token Ring networks in your environment (TrBRF and TrCRF), you must enable VTP version 2 for Token Ring VLAN switching to function properly. To run Token Ring and Token Ring-Net, disable VTP version 2.

Default VTP Configuration

Table 5-5 shows the default VTP configuration.

Table 5-5 VTP Default Configuration

Feature	Default Value
VTP domain name	Null.
VTP mode	Server.
VTP version 2 enable state	Version 2 is disabled.
VTP password	None.

Configuring VTP

You can configure VTP by using the VTP Management window (Figure 5-3).

To display this window, select **VLAN > VTP Management** from the menu bar, and click the VTP Configuration tab.

Figure 5-3 VTP Management: VTP Configuration Tab

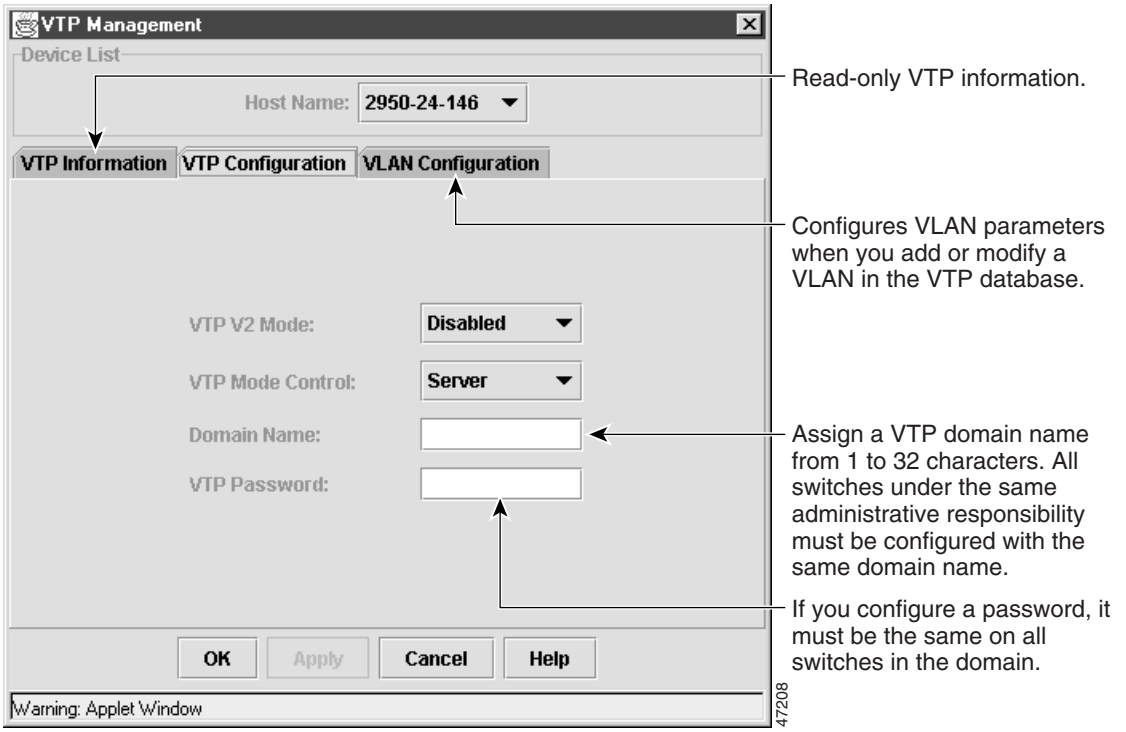

After you configure VTP, you must configure a trunk port so that the switch can send and receive VTP advertisements. For more information, see the "How VLAN Trunks Work" section on page 5-29.

You can also configure VTP through the CLI on standalone, command, and member switches by entering commands in the VLAN database command mode. If you are configuring VTP on a cluster member switch to a VLAN, first log in to the member switch by using the privileged EXEC **rcommand** command. For more information on how to use this command, refer to the *Catalyst 2950 Desktop Switch Command Reference*.

When you enter the **exit** command in VLAN database mode, it applies all the commands that you entered. VTP messages are sent to other switches in the VTP domain, and you are returned to privileged EXEC mode.

Note The Cisco IOS **end** and Ctrl-Z commands are not supported in VLAN database mode.

CLI: Configuring VTP Server Mode

When a switch is in VTP server mode, you can change the VLAN configuration and have it propagated throughout the network.

Beginning in privileged EXEC mode, follow these steps to configure the switch for VTP server mode:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN database mode.
Step 2	**vtp domain** *domain-name*	Configure a VTP administrative-domain name. The name can be from 1 to 32 characters. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.
Step 3	**vtp password** *password-value*	(Optional) Set a password for the VTP domain. The password can be from 8 to 64 characters. If you configure a VTP password, the VTP domain does not function properly if you do not assign the same password to each switch in the domain.
Step 4	**vtp server**	Configure the switch for VTP server mode (the default).
Step 5	**exit**	Return to privileged EXEC mode.
Step 6	**show vtp status**	Verify the VTP configuration. In the display, check the VTP Operating Mode and the VTP Domain Name fields.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Configuring VTP Client Mode

When a switch is in VTP client mode, you cannot change its VLAN configuration. The client switch receives VTP updates from a VTP server in the VTP domain and then modifies its configuration accordingly.

Caution Do not configure a VTP domain name if all switches are operating in VTP client mode. If you do so, it is impossible to make changes to the VLAN configuration of that domain. Therefore, make sure you configure at least one switch as the VTP server.

Beginning in privileged EXEC mode, follow these steps to configure the switch for VTP client mode:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN database mode.
Step 2	**vtp client**	Configure the switch for VTP client mode. The default setting is VTP server.
Step 3	**vtp domain** *domain-name*	Configure a VTP administrative-domain name. The name can be from 1 to 32 characters. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.
Step 4	**vtp password** *password-value*	(Optional) Set a password for the VTP domain. The password can be from 8 to 64 characters. If you configure a VTP password, the VTP domain does not function properly if you do not assign the same password to each switch in the domain.

	Command	Purpose
Step 5	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 6	**show vtp status**	Verify the VTP configuration. In the display, check the VTP Operating Mode field.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Disabling VTP (VTP Transparent Mode)

When you configure the switch for VTP transparent mode, you disable VTP on the switch. The switch then does not send VTP updates and does not act on VTP updates received from other switches. However, a VTP transparent switch does forward received VTP advertisements on all of its trunk links.

Beginning in privileged EXEC mode, follow these steps to configure the switch for VTP transparent mode:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN database mode.
Step 2	**vtp transparent**	Configure the switch for VTP transparent mode. The default setting is VTP server. This step disables VTP on the switch.
Step 3	**exit**	Return to privileged EXEC mode.
Step 4	**show vtp status**	Verify the VTP configuration. In the display, check the VTP Operating Mode field.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Enabling VTP Version 2

VTP version 2 is disabled by default on VTP version 2-capable switches. When you enable VTP version 2 on a switch, every VTP version 2-capable switch in the VTP domain enables version 2.

Caution VTP version 1 and VTP version 2 are not interoperable on switches in the same VTP domain. Every switch in the VTP domain must use the same VTP version. Do not enable VTP version 2 unless every switch in the VTP domain supports version 2.

Note In a Token Ring environment, you must enable VTP version 2 for Token Ring VLAN switching to function properly.

For more information on VTP version configuration guidelines, see the "VTP Version" section on page 5-11.

Beginning in privileged EXEC mode, follow these steps to enable VTP version 2:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**vtp v2-mode**	Enable VTP version 2 on the switch. VTP version 2 is disabled by default on VTP version 2-capable switches.
Step 3	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 4	**show vtp status**	Verify that VTP version 2 is enabled. In the display, check the VTP V2 Mode field.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Disabling VTP Version 2

Beginning in privileged EXEC mode, follow these steps to disable VTP version 2:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**no vtp v2-mode**	Disable VTP version 2.
Step 3	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 4	**show vtp status**	Verify that VTP version 2 is disabled. In the display, check the VTP V2 Mode field.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Monitoring VTP

You monitor VTP by displaying its configuration information: the domain name, the current VTP revision, and the number of VLANs. You can also display statistics about the advertisements sent and received by the switch.

Beginning in privileged EXEC mode, follow these steps to monitor VTP activity:

	Command	Purpose
Step 1	**show vtp status**	Display the VTP switch configuration information.
Step 2	**show vtp counters**	Display counters about VTP messages being sent and received.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

VLANs in the VTP Database

You can set the following parameters when you add a new VLAN to or modify an existing VLAN in the VTP database:

- VLAN ID
- VLAN name
- VLAN type (Ethernet, Fiber Distributed Data Interface [FDDI], FDDI network entity title [NET], TrBRF, or TrCRF, Token Ring, Token Ring-Net)
- VLAN state (active or suspended)
- Maximum transmission unit (MTU) for the VLAN
- Security Association Identifier (SAID)
- Bridge identification number for TrBRF VLANs
- Ring number for FDDI and TrCRF VLANs
- Parent VLAN number for TrCRF VLANs
- Spanning Tree Protocol (STP) type for TrCRF VLANs
- VLAN number to use when translating from one VLAN type to another

The "Default VLAN Configuration" section on page 5-21 lists the default values and possible ranges for each VLAN media type.

Token Ring VLANs

Although the 2950, 2900 XL, and 3500 XL switches do not support Token Ring connections, a remote device such as a Catalyst 5000 series switch with Token Ring connections could be managed from one of the supported switches. Switches running this IOS release advertise information about the following Token Ring VLANs when running VTP version 2:

- Token Ring TrBRF VLANs
- Token Ring TrCRF VLANs

For more information on configuring Token Ring VLANs, see the *Catalyst 5000 Series Software Configuration Guide*.

VLAN Configuration Guidelines

Follow these guidelines when creating and modifying VLANs in your network:

- A maximum of 250 VLANs can be active on supported switches, but some models only support 64 VLANs. (The Catalyst 2950 switches support 64 VLANs.) If VTP reports that there are 254 active VLANs, 4 of the active VLANs (1002 to 1005) are reserved for Token Ring and FDDI.
- Before you can create a VLAN, the switch must be in VTP server mode or VTP transparent mode. For information on configuring VTP, see the "Configuring VTP" section on page 5-12.
- Switches running this IOS release do not support Token Ring or FDDI media. The switch does not forward FDDI, FDDI-Net, TrCRF, or TrBRF traffic, but it does propagate the VLAN configuration through VTP.

Default VLAN Configuration

Table 5-6 through Table 5-10 shows the default configuration for the different VLAN media types.

Note Catalyst 2950 switches support Ethernet interfaces exclusively. Because FDDI and Token Ring VLANs are not locally supported, you configure FDDI and Token Ring media-specific characteristics only for VTP global advertisements to other switches.

Table 5-6 Ethernet VLAN Defaults and Ranges

Parameter	Default	Range
VLAN ID	1	1–1005
VLAN name	VLAN*xxxx*, where *xxxx* is the VLAN ID	No range
802.10 SAID	100000+VLAN ID	1–4294967294
MTU size	1500	1500–18190
Translational bridge 1	0	0–1005
Translational bridge 2	0	0–1005
VLAN state	active	active, suspend

Table 5-7 FDDI VLAN Defaults and Ranges

Parameter	Default	Range
VLAN ID	1002	1–1005
VLAN name	VLAN*xxxx*, where *xxxx* is the VLAN ID	No range
802.10 SAID	100000+VLAN ID	1–4294967294
MTU size	1500	1500–18190
Ring number	None	1–4095
Parent VLAN	0	0–1005

Table 5-7 FDDI VLAN Defaults and Ranges (continued)

Parameter	Default	Range
Translational bridge 1	0	0–1005
Translational bridge 2	0	0–1005
VLAN state	active	active, suspend

Table 5-8 FDDI-Net VLAN Defaults and Ranges

Parameter	Default	Range
VLAN ID	1004	1–1005
VLAN name	VLAN*xxxx*, where *xxxx* is the VLAN ID	No range
802.10 SAID	100000+VLAN ID	1–4294967294
MTU size	1500	1500–18190
Bridge number	0	0–15
STP type	ieee	auto, ibm, ieee
Translational bridge 1	0	0–1005
Translational bridge 2	0	0–1005
VLAN state	active	active, suspend

Table 5-9 Token Ring (TrBRF) VLAN Defaults and Ranges

Parameter	Default	Range
VLAN ID	1005	1–1005
VLAN name	VLAN*xxxx*, where *xxxx* is the VLAN ID	No range
802.10 SAID	100000+VLAN ID	1–4294967294
MTU size	VTPv1 1500; VTPv2 4472	1500–18190
Bridge number	VTPv1 0; VTPv2 user-specified	0–15

Table 5-9 Token Ring (TrBRF) VLAN Defaults and Ranges (continued)

Parameter	Default	Range
STP type	ibm	auto, ibm, ieee
Translational bridge 1	0	0–1005
Translational bridge 2	0	0–1005
VLAN state	active	active, suspend

Table 5-10 Token Ring (TrCRF) VLAN Defaults and Ranges

Parameter	Default	Range
VLAN ID	1003	1–1005
VLAN name	VLAN*xxxx*, where *xxxx* is the VLAN ID	No range
802.10 SAID	100000+VLAN ID	1–4294967294
Ring Number	VTPv1 default 0; VTPv2 user-specified	1–4095
Parent VLAN	VTPv1 default 0; VTPv2 user-specified	0–1005
MTU size	VTPv1 default 1500; VTPv2 default 4472	1500–18190
Translational bridge 1	0	0–1005
Translational bridge 2	0	0–1005
VLAN state	active	active, suspend
Bridge mode	srb	srb, srt
ARE max hops	7	0–13
STE max hops	7	0–13
Backup CRF	disabled	disable; enable

Configuring VLANs in the VTP Database

You can use the VTP Management window (Figure 5-4) or the CLI to add, modify or remove VLAN configurations in the VTP database. VTP globally propagates these VLAN changes throughout the VTP domain.

To display this window, select **VLAN > VTP Management** from the menu bar, and click the VLAN Configuration tab. Click **Help** to for more information on using this window.

Figure 5-4 VTP Management: VLAN Configuration Tab

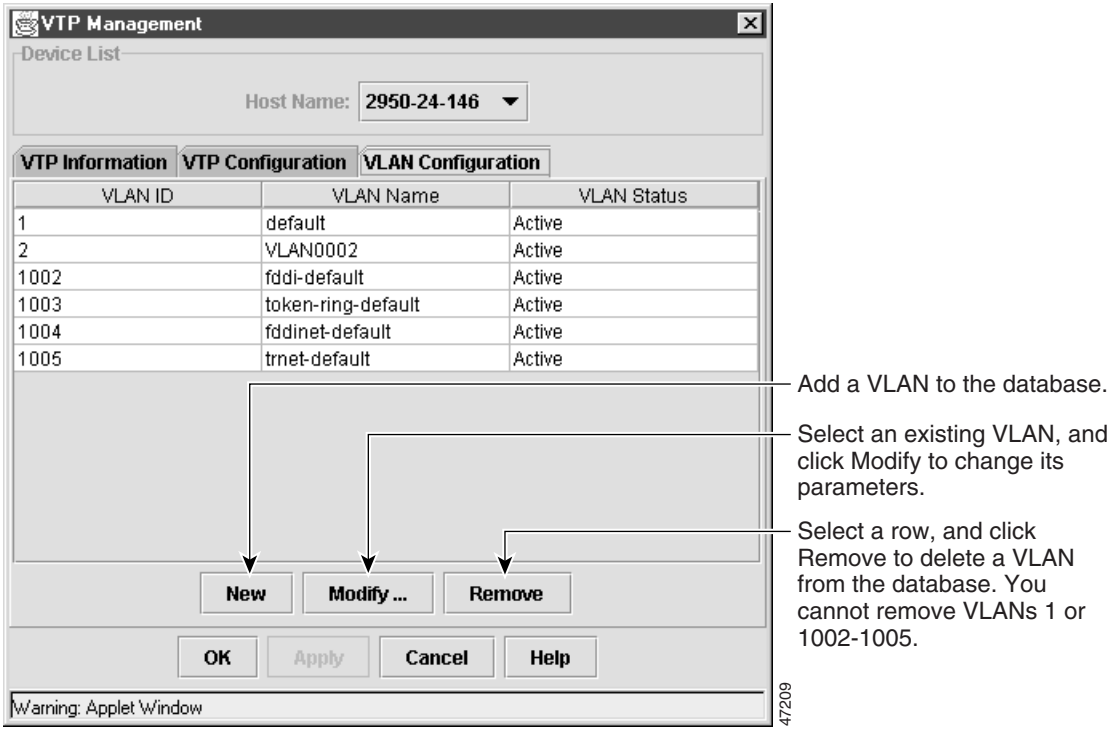

You use the CLI **vlan database** command mode to add, change, and delete VLANs. In VTP server or transparent mode, commands to add, change, and delete VLANs are written to the file vlan.dat, and you can display them by entering the

privileged EXEC mode **show vlan** command. The vlan.dat file is stored in nonvolatile memory. The vlan.dat file is upgraded automatically, but you cannot return to an earlier version of Cisco IOS after you upgrade to this release.

Caution You can cause inconsistency in the VLAN database if you attempt to manually delete the vlan.dat file. If you want to modify the VLAN configuration or VTP, use the VLAN database commands described in the *Catalyst 2950 Desktop Switch Command Reference*.

You use the interface configuration command mode to define the port membership mode and add and remove ports from VLAN. The results of these commands are written to the running-configuration file, and you can display the file by entering the privileged EXEC mode **show running-config command**.

Note VLANs can be configured to support a number of parameters that are not discussed in detail in this section. For complete information on the commands and parameters that control VLAN configuration, refer to the *Catalyst 2950 Desktop Switch Command Reference*.

CLI: Adding an VLAN

Each VLAN has a unique, 4-digit ID that can be a number from 1 to 1001. To add a VLAN to the VLAN database, assign a number and name to the VLAN. For the list of default parameters that are assigned when you add a VLAN, see the "Default VLAN Configuration" section on page 5-21.

If you do not specify the VLAN type, the VLAN is an Ethernet VLAN.

Beginning in privileged EXEC mode, follow these steps to add an Ethernet VLAN:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN database mode.
Step 2	**vlan** *vlan-id* **name** *vlan-name*	Add an Ethernet VLAN by assigning a number to it. If no name is entered for the VLAN, the default is to append the *vlan-id* to the word VLAN. For example, VLAN0004 could be a default VLAN name.
Step 3	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 4	**show vlan name** *vlan-name*	Verify the VLAN configuration.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Modifying a VLAN

Beginning in privileged EXEC mode, follow these steps to modify an Ethernet VLAN:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**vlan** *vlan-id* **mtu** *mtu-size*	Identify the VLAN, and change the MTU size.
Step 3	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 4	**show vlan** *vlan-id*	Verify the VLAN configuration.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Deleting a VLAN

When you delete a VLAN from a switch that is in VTP server mode, the VLAN is removed from all switches in the VTP domain. When you delete a VLAN from a switch that is in VTP transparent mode, the VLAN is deleted only on that specific switch.

You cannot delete the default VLANs for the different media types: Ethernet VLAN 1 and FDDI or Token Ring VLANs 1002 to 1005.

Caution When you delete a VLAN, any ports assigned to that VLAN become inactive. They remain associated with the VLAN (and thus inactive) until you assign them to a new VLAN.

Beginning in privileged EXEC mode, follow these steps to delete a VLAN on the switch:

	Command	Purpose
Step 1	**vlan database**	Enter VLAN configuration mode.
Step 2	**no vlan** *vlan-id*	Remove the VLAN by using the VLAN ID.
Step 3	**exit**	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 4	**show vlan brief**	Verify the VLAN removal.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Assigning Static-Access Ports to a VLAN

By default, all ports are static-access ports assigned to VLAN 1, which is the default management VLAN. If you are assigning a port on a cluster member switch to a VLAN, first log in to the member switch by using the privileged EXEC **rcommand** command. For more information on how to use this command, refer to the *Cisco IOS Desktop Switching Command Reference* (online only).

Beginning in privileged EXEC mode, follow these steps to assign a port to a VLAN in the VTP database:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface*	Enter interface configuration mode, and define the interface to be added to the VLAN.
Step 3	**switchport mode access**	Define the VLAN membership mode for this port.
Step 4	**switchport access vlan 3**	Assign the port to the VLAN.
Step 5	**exit**	Return to privileged EXEC mode.
Step 6	**show interface** *interface-id* **switchport**	Verify the VLAN configuration. In the display, check the Operation Mode, Access Mode VLAN, and the Priority for Untagged Frames fields.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

How VLAN Trunks Work

A trunk is a point-to-point link that transmits and receives traffic between switches or between switches and routers. Trunks carry the traffic of multiple VLANs and can extend VLANs across an entire network.

Figure 5-5 shows a network of switches that are connected by 802.1Q trunks.

Figure 5-5 Catalyst 2950, 2900 XL, and 3500 XL Switches in a 802.1Q Trunking Environment

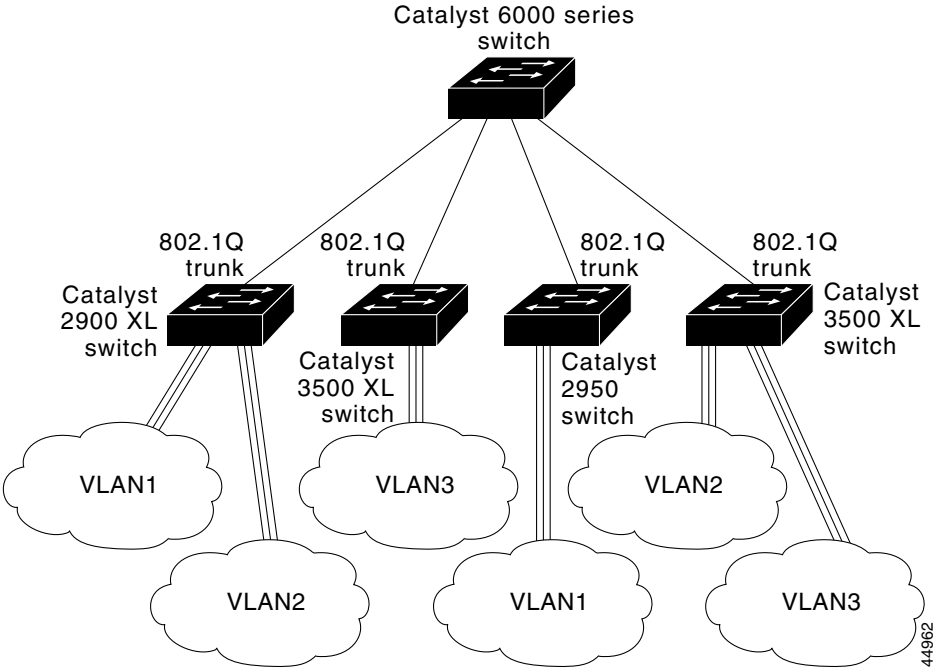

IEEE 802.1Q Configuration Considerations

IEEE 802.1Q trunks impose some limitations on the trunking strategy for a network. The following restrictions apply when using 802.1Q trunks:

- Make sure the native VLAN for a 802.1Q trunk is the same on both ends of the trunk link. If the native VLAN on one end of the trunk is different from the native VLAN on the other end, spanning-tree loops might result.
- Disabling STP on the native VLAN of a 802.1Q trunk without disabling STP on every VLAN in the network can potentially cause STP loops. We recommend that you leave STP enabled on the native VLAN of a 802.1Q trunk or disable STP on every VLAN in the network. Make sure your network is loop-free before disabling STP.

Trunks Interacting with Other Features

IEEE 802.1Q trunking interacts with other switch features as described in Table 5-11.

Table 5-11 Trunks Interacting with Other Features

Switch Feature	Trunk Port Interaction
Port monitoring	A trunk port cannot be a monitor port. A static-access port can monitor the traffic of its VLAN on a trunk port.

Table 5-11 Trunks Interacting with Other Features (continued)

Switch Feature	Trunk Port Interaction
Secure ports	A trunk port cannot be a secure port.
Port grouping	802.1Q trunks can be grouped into EtherChannel port groups, but all trunks in the group must have the same configuration. When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of the following parameters, the switch propagates the setting you entered to all ports in the group: • Allowed-VLAN list • STP path cost for each VLAN • STP port priority for each VLAN • STP Port Fast setting • Trunk status: if one port in a port group ceases to be a trunk, all port cease to be trunks.

Configuring a Trunk Port

You configure trunk ports by using the Assign VLANs (Figure 5-2) and Trunk Configuration (Figure 5-6) tabs of the VLAN Membership window.

To display this window, select **VLAN > VLAN Membership** from the menu bar. Then click the Assign VLANs tab or the Trunk Configuration tab.

Figure 5-6 VLAN Membership: Trunk Configuration Tab

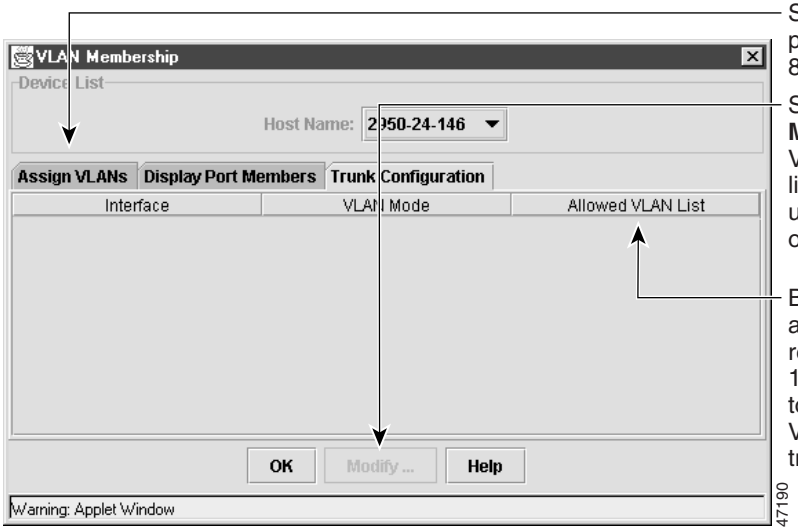

Select this tab to change the port membership mode to 802.1Q trunk.

Select a row or rows, and click **Modify** to change the allowed-VLAN list, the pruning-eligible list, or the native VLAN for untagged traffic (802.1Q trunks only).

By default, VLANs 1-1005 are allowed on each trunk. You can remove VLANs (except VLAN 1002-1005) from the allowed list to prevent traffic from those VLANs from passing over the trunk.

You can also configure a trunk port through the CLI on standalone, command, and member switches. If you are assigning a port on a cluster member switch to a VLAN, first log in to the member switch by using the privileged EXEC **rcommand** command. For more information on how to use this command, refer to the *Catalyst 2950 Desktop Switch Command Reference*.

CLI: Configuring a Trunk Port

For information on trunk port interactions with other features, see the "Trunks Interacting with Other Features" section on page 5-30.

Note Because trunk ports send and receive VTP advertisements, you must ensure that at least one trunk port is configured on the switch and that this trunk port is connected to the trunk port of a second switch. Otherwise, the switch cannot receive any VTP advertisements.

Beginning in privileged EXEC mode, follow these steps to configure a port as a 802.1Q trunk port:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface_id*	Enter the interface configuration mode and the port to be configured for trunking.
Step 3	**switchport mode trunk**	Configure the port as a VLAN trunk.
Step 4	**switchport trunk encapsulation** {**dot1q**}	Configure the port to support 802.1Q encapsulation. You must configure each end of the link with the same encapsulation type.
Step 5	**end**	Return to privileged EXEC mode.
Step 6	**show interface** *interface-id* **switchport**	Verify your entries. In the display, check the Operational Mode and the Operational Trunking Encapsulation fields.
Step 7	**copy running-config startup-config**	Save the configuration.

Note This software release does not support trunk negotiation through the Dynamic Trunk Protocol (DTP), formerly known as Dynamic ISL (DISL). If you are connecting a trunk port to a Catalyst 5000 switch or other DTP device, use the non-negotiate option on the DTP-capable device so that the switch port does not generate DTP frames.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Disabling a Trunk Port

You can disable trunking on a port by returning it to its default static-access mode.

Beginning in privileged EXEC mode, follow these steps to disable trunking on a port:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface_id*	Enter the interface configuration mode and the port to be added to the VLAN.
Step 3	**no switchport mode**	Return the port to its default static-access mode.
Step 4	**end**	Return to privileged EXEC.
Step 5	**show interface** *interface-id* **switchport**	Verify your entries. In the display, check the Negotiation of Trunking field.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Defining the Allowed VLANs on a Trunk

By default, a trunk port sends to and receives traffic from all VLANs in the VLAN database. All VLANs, 1 to 1005, are allowed on each trunk. However, you can remove VLANs from the allowed list, preventing traffic from those VLANs from passing over the trunk. To restrict the traffic a trunk carries, use the **remove** *vlan-list* parameter to remove specific VLANs from the allowed list.

A trunk port can become a member of a VLAN if the VLAN is enabled, if VTP knows of the VLAN, and if the VLAN is in the allowed list for the port. When VTP detects a newly enabled VLAN and the VLAN is in the allowed list for a trunk port, the trunk port automatically becomes a member of the enabled VLAN. When VTP detects a new VLAN and the VLAN is not in the allowed list for a trunk port, the trunk port does not become a member of the new VLAN.

Beginning in privileged EXEC mode, follow these steps to modify the allowed list of a 802.1Q trunk:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface_id*	Enter interface configuration mode and the port to be added to the VLAN.
Step 3	**switchport mode trunk**	Configure VLAN membership mode for trunks.
Step 4	**switchport trunk allowed vlan remove** *vlan-list*	Define the VLANs that are *not* allowed to transmit and receive on the port. The *vlan-list* parameter is a range of VLAN IDs Separate nonconsecutive VLAN IDs with a comma and no spaces; use a hyphen to designate a range of IDs. Valid IDs are from 2 to 1001.
Step 5	**end**	Return to privileged EXEC.
Step 6	**show interface** *interface-id* **switchport allowed-vlan**	Verify your entries.
Step 7	**copy running-config startup-config**	Save the configuration.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Configuring the Native VLAN for Untagged Traffic

A trunk port configured with 802.1Q tagging can receive both tagged and untagged traffic. By default, the switch forwards untagged traffic with the native VLAN configured for the port. The native VLAN is VLAN 1 by default.

Note The native VLAN can be assigned any VLAN ID, and it is not dependent on the management VLAN.

For information about 802.1Q configuration issues, see the "IEEE 802.1Q Configuration Considerations" section on page 5-30.

Beginning in privileged EXEC mode, follow these steps to configure the native VLAN on a 802.1Q trunk:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface-id*	Enter interface configuration mode, and define the interface that is configured as the 802.1Q trunk.
Step 3	**switchport trunk native vlan** *vlan-id*	Configure the VLAN that is sending and receiving untagged traffic on the trunk port. Valid IDs are from 1 to 1001.
Step 4	**show interface** *interface-id* **switchport**	Verify your settings.

If a packet has a VLAN ID the same as the outgoing port native VLAN ID, the packet is transmitted untagged; otherwise, the switch transmits the packet with a tag.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

Configuring IEEE 802.1p Class of Service

The Catalyst 2950 switches provide QoS-based 802.1p class of service (CoS) values. QoS uses classification and scheduling to transmit network traffic from the switch in a predictable manner. QoS classifies frames by assigning priority-indexed CoS values to them and gives preference to higher-priority traffic such as telephone calls.

How Class of Service Works

Before you set up 802.1p CoS on a Catalyst 2950, 2900 XL, and 3500 XL switch that operates with the Catalyst 6000 family of switches, refer to the Catalyst 6000 documentation. There are differences in the 802.1p implementation, and they should be understood to ensure compatibility.

Port Priority

Frames received from users in the administratively-defined VLANs are classified or *tagged* for transmission to other devices. Based on rules you define, a unique identifier (the tag) is inserted in each frame header before it is forwarded. The tag is examined and understood by each device before any broadcasts or transmissions to other switches, routers, or end stations. When the frame reaches the last switch or router, the tag is removed before the frame is transmitted to the target end station. VLANs that are assigned on trunk or access ports without identification or a tag are called *native* or *untagged* frames.

For IEEE 802.1Q frames with tag information, the priority value from the header frame is used. For native frames, the default priority of the input port is used.

Port Scheduling

Each port on the switch has a single receive queue buffer (the *ingress* port) for incoming traffic. When an untagged frame arrives, it is assigned the value of the port as its port default priority. You assign this value by using the CLI or CMS software. A tagged frame continues to use its assigned CoS value when it passes through the ingress port.

CoS configures each transmit port (the *egress* port) with a normal-priority transmit queue and a high-priority transmit queue, depending on the frame tag or the port information. Frames in the normal-priority queue are forwarded only after frames in the high-priority queue are forwarded.

Table 5-12 shows the two categories of switch transmit queues.

Table 5-12 Transmit Queue Information

Transmit queue category[1]	Transmit Queues
2950 switches (802.1p user priority)	There are four priority queues. The frames are forwarded to appropriate queues based on priority-to-queue mapping as defined by the user.
2900 XL switches, 2900 XL Ethernet modules (802.1p user priority)	Frames with a priority value of 0 through 3 are sent to a normal-priority queue. Frames with a priority value of 4 through 7 are sent to a high-priority queue.
3500 XL switches, Gigabit Ethernet modules (802.1p user priority)	Frames with a priority value of 0 through 3 are sent to a normal-priority queue. Frames with a priority value of 4 through 7 are sent to a high-priority queue.

1. Catalyst 2900 XL switches with 4 MB of DRAM and the WS-X2914-XL and the WS-X2922-XL modules only have one transmit queue and do not support QoS.

CLI: Configuring the CoS Port Priorities

Beginning in privileged EXEC mode, follow these steps to set the port priority for untagged (native) Ethernet frames:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**interface** *interface*	Enter the interface to be configured.
Step 3	**switchport priority default** *default-priority-id*	Set the port priority on the interface. Frames are forwarded to appropriate queues as per CoS to queue mapping.

	Command	Purpose
Step 4	**end**	Return to privileged EXEC mode.
Step 5	**show interface** *interface-id* **switchport**	Verify your entries. In the display, check the Priority for Untagged Frames field.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CoS and WRR

The Catalyst 2950 switches support four CoS queues for each egress port. For each queue, you can specify the following types of scheduling:

- Strict priority scheduling

  Strict priority scheduling is based on the priority of queues. Packets can have priorities from 0 to 7, 7 being the highest. Packets in the high-priority queue always transmit first, and packets in the low-priority queue do not transmit until all the high-priority queues become empty.

- Weighted round-robin (WRR) scheduling

  WRR scheduling requires you to specify a number that indicates the importance (weight) of the queue relative to the other CoS queues. WRR scheduling prevents the low-priority queues from being completely neglected during periods of high-priority traffic. The WRR scheduler transmits some packets from each queue in turn. The number of packets it transmits corresponds to the relative importance of the queue. For example, if one queue has a weight of 3 and another has a weight of 4, then three packets are transmitted from the first queue for every four that are transmitted from the second queue. By using this scheduling, low-priority queues have the opportunity to transmit packets even though the high-priority queues are not empty.

Use the CoS and WRR window (Figure 5-7) to assign priorities to the queues and to enable the WRR scheduler. To display this window, select **Device > CoS & WRR** from the menu bar.

You can use this window to perform the following tasks:

- Enable or disable WRR
- Assign packets to queues based on priority

- Assign relative weights to the output queues

Use the **CoS** tab on the CoS and WRR window (Figure 5-7) to view the default settings. If you want to reassign a priority, open the list under that priority, and select a different queue number.

Figure 5-7 Modify CoS Settings

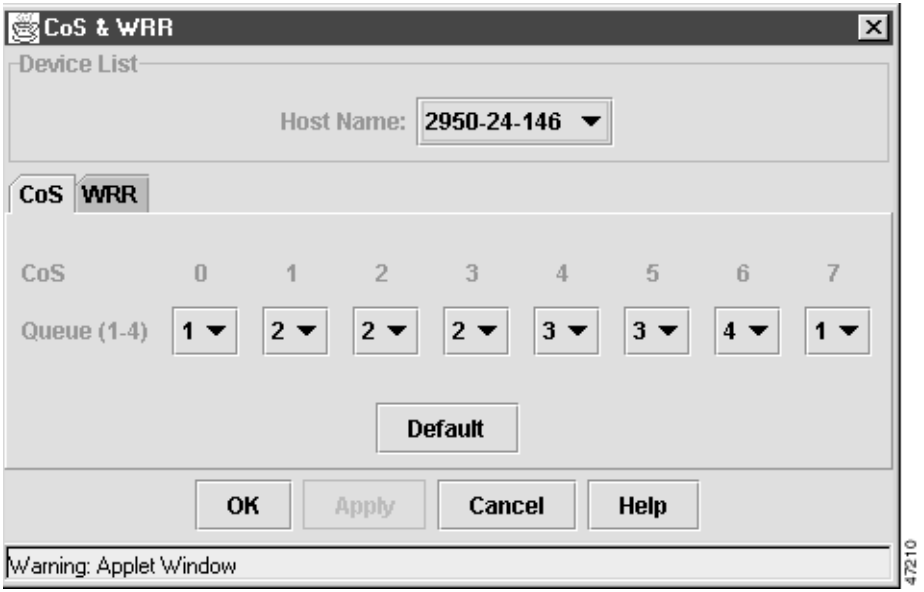

Use the **WRR** tab on the CoS and WRR window (Figure 5-8) to view the current settings. If WRR scheduler is disabled, all the fields will be blank.

If the WRR priority box is checked, WRR is enabled. You can assign a weighted number from 0 to 255 in the field below each queue number, as shown in Figure 5-8.

Figure 5-8 Modify WRR Settings

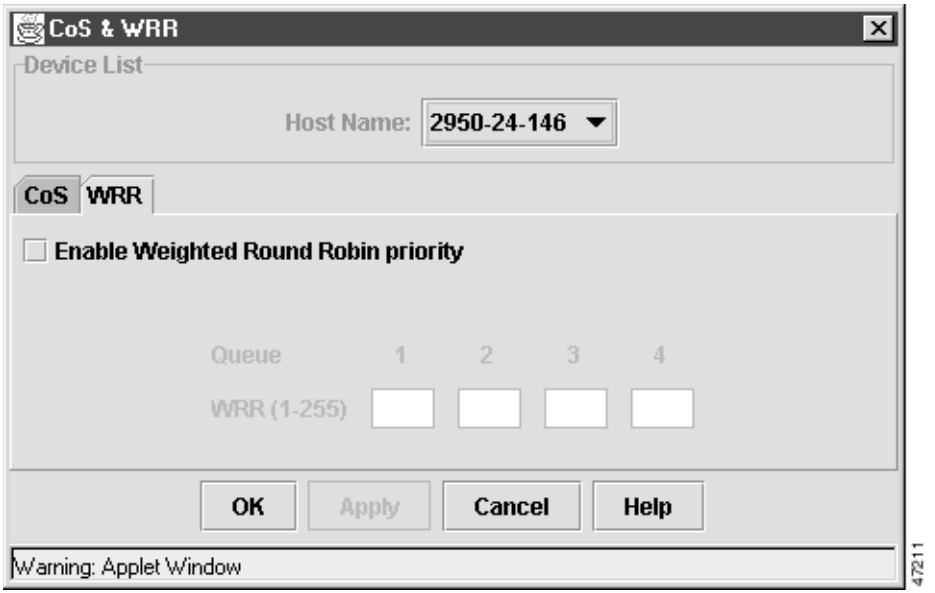

CLI: Configuring CoS Priority Queues

Beginning in privileged EXEC mode, follow these steps to configure the CoS priority queues:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**wrr-queue cos-map** *qid cos1..cosn*	Specify the queue id of the CoS priority queue. (Ranges are 1 to 4 where 1 is the lowest CoS priority queue.) Specify the CoS values that are mapped to queue id. Default values are as follows: CoS Value — CoS Priority Queues 0, 1 — 1 2, 3 — 2 4, 5 — 3 6, 7 — 4
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show cos-map**	Display the mapping of the CoS priority queues.

To disable the new CoS settings and return to default settings, use the **no wrr-queue cos-map** command.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

CLI: Configuring WRR

Beginning in privileged EXEC mode, follow these steps to configure the weighted round robin priority:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode.
Step 2	**wrr-queue bandwidth** *weight1...weight4*	Assign WRR weights to the four CoS queues. (Ranges for the WRR values are 1 to 255.)
Step 3	**end**	Return to privileged EXEC mode.
Step 4	**show wrr-queue bandwidth**	Display the WRR bandwidth allocation for the CoS priority queues.

To disable the WRR scheduler and enable the strict priority scheduler, use the **no wrr-queue bandwidth** command.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

Load Sharing Using STP

Load sharing divides the bandwidth supplied by parallel trunks connecting switches. To avoid loops, STP normally blocks all but one parallel link between switches. With load sharing, you divide the traffic between the links according to which VLAN the traffic belongs.

You configure load sharing on trunk ports by using STP port priorities or STP path costs. For load sharing using STP port priorities, both load-sharing links must be connected to the same switch. For load sharing using STP path costs, each load-sharing link can be connected to the same switch or to two different switches.

You can change STP port parameters by using the Port Parameters tab of the Spanning Tree Protocol window or by using the CLI. To display this window, select **Device > Spanning-Tree Protocol** from the menu bar. Then click the **Port Parameters** tab.

For more information about the STP window, see the "Configuring the Spanning Tree Protocol" section on page 4-80, or consult the online help in the application.

Load Sharing Using STP Port Priorities

When two ports on the same switch form a loop, the STP port priority setting determines which port is enabled and which port is in standby mode. You can set the priorities on a parallel trunk port so that the port carries all the traffic for a given VLAN. The trunk port with the higher priority (lower values) for a VLAN is forwarding traffic for that VLAN. The trunk port with the lower priority (higher values) for the same VLAN remains in a blocking state for that VLAN. One trunk port transmits or receives all traffic for the VLAN.

Figure 5-9 shows two trunks connecting supported switches. In this example, the switches are configured as follows:

- VLANs 8 through 10 are assigned a port priority of 10 on trunk 1.
- VLANs 3 through 6 retain the default port priority of 128 on trunk 1.
- VLANs 3 through 6 are assigned a port priority of 10 on trunk 2.
- VLANs 8 through 10 retain the default port priority of 128 on trunk 2.

In this way, trunk 1 carries traffic for VLANs 8 through 10, and trunk 2 carries traffic for VLANs 3 through 6. If the active trunk fails, the trunk with the lower priority takes over and carries the traffic for all of the VLANs. No duplication of traffic occurs over any trunk port.

Figure 5-9 Load Sharing by Using STP Port Priorities

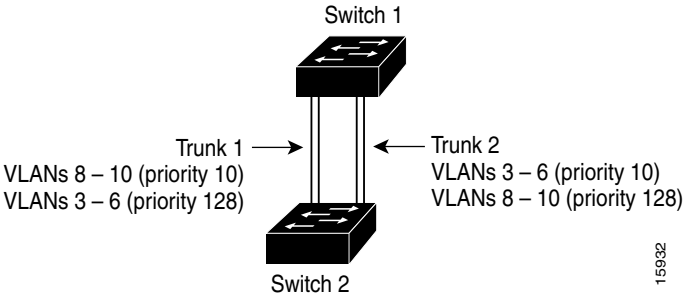

CLI: Configuring STP Port Priorities and Load Sharing

Beginning in privileged EXEC mode, follow these steps to configure the network shown in Figure 5-9:

	Command	Purpose
Step 1	**vlan database**	On Switch 1, enter VLAN configuration mode.
Step 2	**vtp domain** *domain-name*	Configure a VTP administrative domain. The domain name can be from 1 to 32 characters.
Step 3	**vtp server**	Configure Switch 1 as the VTP server.
Step 4	**exit**	Return to privileged EXEC mode.
Step 5	**show vtp status**	Verify the VTP configuration on both Switch 1 and Switch 2. In the display, check the VTP Operating Mode and the VTP Domain Name fields.
Step 6	**show vlan**	Verify that the VLANs exist in the database on Switch 1.
Step 7	**configure terminal**	Enter global configuration mode.
Step 8	**interface fa0/1**	Enter interface configuration mode, and define Fa0/1 as the interface to be configured as a trunk.
Step 9	**switchport mode trunk**	Configure the port as a trunk port.
Step 10	**end**	Return to privilege EXEC mode.
Step 11	**show interface fa0/1 switchport**	Verify the VLAN configuration.
Step 12		Repeat Steps 7 through 11 on Switch 1 for interface Fa0/2.
Step 13		Repeat Steps 7 through 11 on Switch 2 to configure the trunk ports on interface Fa0/1 and Fa0/2.

	Command	Purpose
Step 14	**show vlan**	When the trunk links come up, VTP passes the VTP and VLAN information to Switch 2. Verify the Switch 2 has learned the VLAN configuration.
Step 15	**configure terminal**	Enter global configuration mode on Switch 1.
Step 16	**interface fa0/1**	Enter interface configuration mode, and define the interface to set the STP port priority.
Step 17	**spanning-tree vlan 8 9 10 port-priority 10**	Assign the port priority of 10 for VLANs 8, 9, and 10.
Step 18	**end**	Return to global configuration mode.
Step 19	**interface fa0/2**	Enter interface configuration mode, and define the interface to set the STP port priority.
Step 20	**spanning-tree vlan 3 4 5 6 port priority 10**	Assign the port priority of 10 for VLANs 3, 4, 5, and 6.
Step 21	**exit**	Return to privileged EXEC mode.
Step 22	**show running-config**	Verify your entries.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation.

Load Sharing Using STP Path Cost

You can configure parallel trunks to share VLAN traffic by setting different path costs on a trunk and associating the path costs with different sets of VLANs. The VLANs keep the traffic separate, because no loops exist, STP does not disable the ports, and redundancy is maintained in the event of a lost link.

In Figure 5-10, trunk ports 1 and 2 are 100BaseT ports. The path costs for the VLANs are assigned as follows:

- VLANs 2 through 4 are assigned a path cost of 30 on trunk port 1.
- VLANs 8 through 10 retain the default 100BaseT path cost on trunk port 1 of 19.
- VLANs 8 through 10 are assigned a path cost of 30 on trunk port 2.
- VLANs 2 through 4 retain the default 100BaseT path cost on trunk port 2 of 19.

Figure 5-10 Load-Sharing Trunks with Traffic Distributed by Path Cost

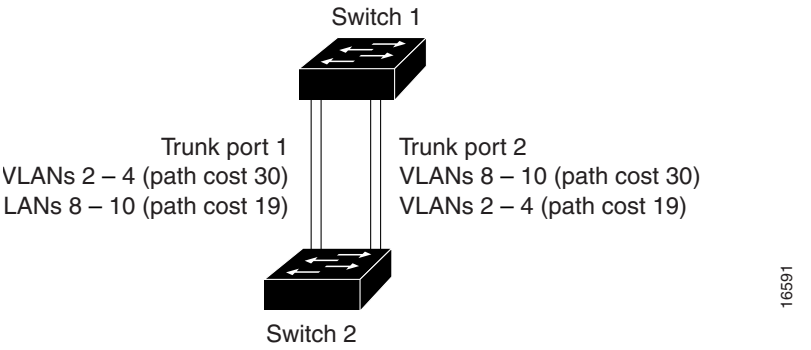

CLI: Configuring STP Path Costs and Load Sharing

Beginning in privileged EXEC mode, follow these steps to configure the network shown in Figure 5-10:

	Command	Purpose
Step 1	**configure terminal**	Enter global configuration mode on Switch 1.
Step 2	**interface fa0/1**	Enter interface configuration mode, and define Fa0/1 as the interface to be configured as a trunk.
Step 3	**switchport mode trunk**	Configure the port as a trunk port.
Step 4	**end**	Return to global configuration mode.
Step 5		Repeat Steps 2 through 4 on Switch 1 interface Fa0/2.
Step 6	**show running-config**	Verify your entries. In the display, make sure that interface Fa0/1 and Fa0/2 are configured as trunk ports.
Step 7	**show vlan**	When the trunk links come up, Switch 1 receives the VTP information from the other switches. Verify that Switch 1 has learned the VLAN configuration.
Step 8	**configure terminal**	Enter global configuration mode.
Step 9	**interface fa0/1**	Enter interface configuration mode, and define Fa0/1 as the interface to set the STP cost.
Step 10	**spanning-tree vlan 2 3 4 cost 30**	Set the spanning-tree path cost to 30 for VLANs 2, 3, and 4.
Step 11	**end**	Return to global configuration mode.
Step 12		Repeat Steps 9 through 11 on Switch 1 interface Fa0/2, and set the spanning-tree path cost to 30 for VLANs 8, 9, and 10.

	Command	Purpose
Step 13	**exit**	Return to privileged EXEC mode.
Step 14	**show running-config**	Verify your entries. In the display, verify that the path costs are set correctly for interface Fa0/1 and Fa0/2.

The "Finding More Information About IOS Commands" section on page 4-1 contains the path to the complete IOS documentation set.