

3.0 VLANs

3.0.1.1 Introduction

Network performance is a key factor in the productivity of an organization. One of the technologies used to improve network performance is the separation of large broadcast domains into smaller ones. By design, routers will block broadcast traffic at an interface. However, routers normally have a limited number of LAN interfaces. A router's primary role is to move information between networks, not provide network access to end devices.

The role of providing access into a LAN is normally reserved for an access layer switch. A virtual local area network (VLAN) can be created on a Layer 2 switch to reduce the size of broadcast domains, similar to a Layer 3 device. VLANs are commonly incorporated into network design making it easier for a network to support the goals of an organization. While VLANs are primarily used within switched local area networks, modern implementations of VLANs allow them to span MANs and WANs.

This chapter will cover how to configure, manage, and troubleshoot VLANs and VLAN trunks. It will also examine security considerations and strategies relating to VLANs and trunks, and best practices for VLAN design.

Refer to
Lab Activity
for this chapter

3.0.1.2 Class Activity Vacation Station

Vacation Station

Scenario

You have purchased a vacation home at the beach for rental purposes. There are three identical floors on each level of the home. Each floor offers one digital television for renters to use.

According to the local Internet service provider, only three stations may be offered within a television package. It is your job to decide which television packages you offer your guests.

- Divide the class into groups of three students per group.
- Choose three different stations to make one subscription package for each floor of your rental home.
- Complete the PDF for this activity.
- Share your completed group-reflection answers with the class.

Refer to
Online Course
for Illustration

3.1 VLAN Segmentation

3.1.1 Overview of VLANs

3.1.1.1 VLAN Definitions

Within a switched internetwork, VLANs provide segmentation and organizational flexibility. VLANs provide a way to group devices within a LAN. A group of devices within a VLAN communicate as if they were attached to the same wire. VLANs are based on logical connections, instead of physical connections.

VLANs allow an administrator to segment networks based on factors such as function, project team, or application, without regard for the physical location of the user or device. Devices within a VLAN act as if they are in their own independent network, even if they share a common infrastructure with other VLANs. Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations within the VLAN where the packets are sourced. Each VLAN is considered a separate logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a device that supports routing.

A VLAN creates a logical broadcast domain that can span multiple physical LAN segments. VLANs improve network performance by separating large broadcast domains into smaller ones. If a device in one VLAN sends a broadcast Ethernet frame, all devices in the VLAN receive the frame, but devices in other VLANs do not.

VLANs enable the implementation of access and security policies according to specific groupings of users. Each switch port can be assigned to only one VLAN (with the exception of a port connected to an IP phone or to another switch).

Refer to
Online Course
for Illustration

3.1.1.2 Benefits of VLANs

User productivity and network adaptability are important for business growth and success. VLANs make it easier to design a network to support the goals of an organization. The primary benefits of using VLANs are as follows:

- **Security** - Groups that have sensitive data are separated from the rest of the network, decreasing the chances of confidential information breaches. As shown in the figure, faculty computers are on VLAN 10 and completely separated from student and guest data traffic.
- **Cost reduction** - Cost savings result from reduced need for expensive network upgrades and more efficient use of existing bandwidth and uplinks.
- **Better performance** - Dividing flat Layer 2 networks into multiple logical workgroups (broadcast domains) reduces unnecessary traffic on the network and boosts performance.
- **Shrink broadcast domains** - Dividing a network into VLANs reduces the number of devices in the broadcast domain. As shown in the figure, there are six computers on this network but there are three broadcast domains: Faculty, Student, and Guest.
- **Improved IT staff efficiency** - VLANs make it easier to manage the network because users with similar network requirements share the same VLAN. When a new switch is provisioned, all the policies and procedures already configured for the particular

VLAN are implemented when the ports are assigned. It is also easy for the IT staff to identify the function of a VLAN by giving it an appropriate name. In the figure, for easy identification VLAN 10 has been named “Faculty”, VLAN 20 is named “Student”, and VLAN 30 “Guest.”

- **Simpler project and application management** - VLANs aggregate users and network devices to support business or geographic requirements. Having separate functions makes managing a project or working with a specialized application easier; an example of such an application is an e-learning development platform for faculty.

Each VLAN in a switched network corresponds to an IP network; therefore, VLAN design must take into consideration the implementation of a hierarchical network-addressing scheme. Hierarchical network addressing means that IP network numbers are applied to network segments or VLANs in an orderly fashion that takes the network as a whole into consideration. Blocks of contiguous network addresses are reserved for and configured on devices in a specific area of the network, as shown in the figure.

Refer to
Online Course
for Illustration

3.1.1.3 Types of VLANs

There are a number of distinct types of VLANs used in modern networks. Some VLAN types are defined by traffic classes. Other types of VLANs are defined by the specific function that they serve.

Data VLAN

A data VLAN is a VLAN that is configured to carry user-generated traffic. A VLAN carrying voice or management traffic would not be part of a data VLAN. It is common practice to separate voice and management traffic from data traffic. A data VLAN, is sometimes referred to as a user VLAN. Data VLANs are used to separate the network into groups of users or devices.

Default VLAN

All switch ports become a part of the default VLAN after the initial boot up of a switch loading the default configuration. Switch ports that participate in the default VLAN are part of the same broadcast domain. This allows any device connected to any switch port to communicate with other devices on other switch ports. The default VLAN for Cisco switches is VLAN 1. In the figure, the `show vlan brief` command was issued on a switch running the default configuration. Notice that all ports are assigned to VLAN 1 by default.

VLAN 1 has all the features of any VLAN, except it cannot be renamed or deleted. By default, all Layer 2 control traffic is associated with VLAN 1.

Native VLAN

A native VLAN is assigned to an 802.1Q trunk port. Trunk ports are the links between switches that support the transmission of traffic associated with more than one VLAN. An 802.1Q trunk port supports traffic coming from many VLANs (tagged traffic), as well as traffic that does not come from a VLAN (untagged traffic). The 802.1Q trunk port places untagged traffic on the native VLAN, which by default is VLAN 1.

Native VLANs are defined in the IEEE 802.1Q specification to maintain backward compatibility with untagged traffic common to legacy LAN scenarios. A native VLAN serves as a common identifier on opposite ends of a trunk link.

It is a best practice to configure the native VLAN as an unused VLAN, distinct from VLAN 1 and other VLANs. In fact, it is not unusual to dedicate a fixed VLAN to serve the role of the native VLAN for all trunk ports in the switched domain.

Management VLAN

A management VLAN is any VLAN configured to access the management capabilities of a switch. VLAN 1 is the management VLAN by default. To create the management VLAN, the switch virtual interface (SVI) of that VLAN is assigned an IP address and subnet mask, allowing the switch to be managed via HTTP, Telnet, SSH, or SNMP. Because the out-of-the-box configuration of a Cisco switch has VLAN 1 as the default VLAN, VLAN 1 would be a bad choice for the management VLAN.

In the past, the management VLAN for a 2960 switch was the only active SVI. On 15.x versions of the Cisco IOS for Catalyst 2960 Series switches, it is possible to have more than one active SVI. With Cisco IOS 15.x, the particular active SVI assigned for remote management must be documented. While theoretically a switch can have more than one management VLAN, having more than one increases exposure to network attacks.

In the figure, all ports are currently assigned to the default VLAN 1. No native VLAN is explicitly assigned and no other VLANs are active; therefore the network is designed with the native VLAN the same as the management VLAN. This is considered a security risk.

Refer to
Online Course
for Illustration

3.1.1.4 Voice VLANs

A separate VLAN is needed to support Voice over IP (VoIP). VoIP traffic requires:

- Assured bandwidth to ensure voice quality
- Transmission priority over other types of network traffic
- Ability to be routed around congested areas on the network
- Delay of less than 150 ms across the network

To meet these requirements, the entire network has to be designed to support VoIP. The details of how to configure a network to support VoIP are beyond the scope of this course, but it is useful to summarize how a voice VLAN works between a switch, a Cisco IP phone, and a computer.

In the figure, VLAN 150 is designed to carry voice traffic. The student computer PC5 is attached to the Cisco IP phone, and the phone is attached to switch S3. PC5 is in VLAN 20, which is used for student data.

Refer to **Packet Tracer Activity**
for this chapter

3.1.1.5 Packet Tracer - Who Hears the Broadcast?

In this activity, a 24-port Catalyst 2960 switch is fully populated. All ports are in use. You will observe broadcast traffic in a VLAN implementation and answer some reflection questions.

Refer to
Online Course
for Illustration

3.1.2 VLANs in a Multi-Switched Environment

3.1.2.1 VLAN Trunks

A VLAN trunk, or trunk, is a point-to-point link between two network devices that carries more than one VLAN. A VLAN trunk extends VLANs across an entire network. Cisco

supports IEEE 802.1Q for coordinating trunks on Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces.

VLANs would not be very useful without VLAN trunks. VLAN trunks allow all VLAN traffic to propagate between switches, so that devices which are in the same VLAN, but connected to different switches, can communicate without the intervention of a router.

A VLAN trunk does not belong to a specific VLAN; rather, it is a conduit for multiple VLANs between switches and routers. A trunk could also be used between a network device and server or other device that is equipped with an appropriate 802.1Q-capable NIC. By default, on a Cisco Catalyst switch, all VLANs are supported on a trunk port.

In the figure, the links between switches S1 and S2, and S1 and S3 are configured to transmit traffic coming from VLANs 10, 20, 30, and 99 across the network. This network could not function without VLAN trunks.

Refer to
Online Course
for Illustration

3.1.2.2 Controlling Broadcast Domains with VLANs

Network without VLANs

In normal operation, when a switch receives a broadcast frame on one of its ports, it forwards the frame out all other ports except the port where the broadcast was received. In the animation in Figure 1, the entire network is configured in the same subnet (172.17.40.0/24) and no VLANs are configured. As a result, when the faculty computer (PC1) sends out a broadcast frame, switch S2 sends that broadcast frame out all of its ports. Eventually the entire network receives the broadcast because the network is one broadcast domain.

Network with VLANs

As shown in the animation in Figure 2, the network has been segmented using two VLANs: Faculty devices are assigned to VLAN 10 and Student devices are assigned to VLAN 20. When a broadcast frame is sent from the faculty computer, PC1, to switch S2, the switch forwards that broadcast frame only to those switch ports configured to support VLAN 10.

The ports that comprise the connection between switches S2 and S1 (ports F0/1), and between S1 and S3 (ports F0/3) are trunks and have been configured to support all the VLANs in the network.

When S1 receives the broadcast frame on port F0/1, S1 forwards that broadcast frame out of the only other port configured to support VLAN 10, which is port F0/3. When S3 receives the broadcast frame on port F0/3, it forwards that broadcast frame out of the only other port configured to support VLAN 10, which is port F0/11. The broadcast frame arrives at the only other computer in the network configured in VLAN 10, which is faculty computer PC4.

When VLANs are implemented on a switch, the transmission of unicast, multicast, and broadcast traffic from a host in a particular VLAN are restricted to the devices that are in that VLAN.

Refer to
Online Course
for Illustration

3.1.2.3 Tagging Ethernet Frames for VLAN Identification

Catalyst 2960 Series switches are Layer 2 devices. They use the Ethernet frame header information to forward packets. They do not have routing tables. The standard Ethernet frame header does not contain information about the VLAN to which the frame belongs;

thus, when Ethernet frames are placed on a trunk, information about the VLANs to which they belong must be added. This process, called tagging, is accomplished by using the IEEE 802.1Q header, specified in the IEEE 802.1Q standard. The 802.1Q header includes a 4-byte tag inserted within the original Ethernet frame header, specifying the VLAN to which the frame belongs.

When the switch receives a frame on a port configured in access mode and assigned a VLAN, the switch inserts a VLAN tag in the frame header, recalculates the FCS, and sends the tagged frame out of a trunk port.

VLAN Tag Field Details

The VLAN tag field consists of a Type field, a tag control information field, and the FCS field:

- **Type** - A 2-byte value called the tag protocol ID (TPID) value. For Ethernet, it is set to hexadecimal 0x8100.
- **User priority** - A 3-bit value that supports level or service implementation.
- **Canonical Format Identifier (CFI)** - A 1-bit identifier that enables Token Ring frames to be carried across Ethernet links.
- **VLAN ID (VID)** - A 12-bit VLAN identification number that supports up to 4096 VLAN IDs.

After the switch inserts the Type and tag control information fields, it recalculates the FCS values and inserts the new FCS into the frame.

Refer to
Online Course
for Illustration

3.1.2.4 Native VLANs and 802.1Q Tagging

Tagged Frames on the Native VLAN

Some devices that support trunking, add a VLAN tag to native VLAN traffic. Control traffic sent on the native VLAN should not be tagged. If an 802.1Q trunk port receives a tagged frame with the VLAN ID the same as the native VLAN, it drops the frame. Consequently, when configuring a switch port on a Cisco switch, configure devices so that they do not send tagged frames on the native VLAN. Devices from other vendors that support tagged frames on the native VLAN include IP phones, servers, routers, and non-Cisco switches.

Untagged Frames on the Native VLAN

When a Cisco switch trunk port receives untagged frames (which are unusual in a well-designed network), it forwards those frames to the native VLAN. If there are no devices associated with the native VLAN (which is not unusual) and there are no other trunk ports (which is not unusual), then the frame is dropped. The default native VLAN is VLAN 1. When configuring an 802.1Q trunk port, a default Port VLAN ID (PVID) is assigned the value of the native VLAN ID. All untagged traffic coming in or out of the 802.1Q port is forwarded based on the PVID value. For example, if VLAN 99 is configured as the native VLAN, the PVID is 99 and all untagged traffic is forwarded to VLAN 99. If the native VLAN has not been reconfigured, the PVID value is set to VLAN 1.

In the figure, PC1 is connected by a hub to an 802.1Q trunk link. PC1 sends untagged traffic which the switches associate with the native VLAN configured on the trunk ports, and forward accordingly. Tagged traffic on the trunk received by PC1 is dropped. This scenario

reflects poor network design for several reasons: it uses a hub, it has a host connected to a trunk link, and it implies that the switches have access ports assigned to the native VLAN. But it illustrates the motivation for the IEEE 802.1Q specification for native VLANs as a means of handling legacy scenarios.

Refer to
Online Course
for Illustration

3.1.2.5 Voice VLAN Tagging

As shown in Figure 1, the F0/18 port on S3 is configured to be in voice mode so that voice frames will be tagged with VLAN 150. Data frames coming through the Cisco IP phone from PC5 are left untagged. Data frames destined for PC5 coming from port F0/18 are tagged with VLAN 20 on the way to the phone. The phone strips the VLAN tag before the data is forwarded to PC5.

The Cisco IP Phone contains an integrated three-port 10/100 switch. The ports provide dedicated connections to these devices:

- Port 1 connects to the switch or other VoIP device.
- Port 2 is an internal 10/100 interface that carries the IP phone traffic.
- Port 3 (access port) connects to a PC or other device.

When the switch port has been configured with a voice VLAN, the link between the switch and the IP phone acts as a trunk to carry both the tagged voice traffic and untagged data traffic. Communication between the switch and IP phone is facilitated by the Cisco Discovery Protocol (CDP).

Sample Configuration

Figure 2 shows a sample output. A discussion of voice Cisco IOS commands are beyond the scope of this course but the highlighted areas in the sample output show the F0/18 interface configured with a VLAN configured for data (VLAN 20) and a VLAN configured for voice (VLAN 150).

Refer to
Interactive Graphic
in online course.

3.1.2.6 Activity - Predict Switch Behavior

Refer to Packet
Tracer Activity
for this chapter

3.1.2.7 Packet Tracer - Investigating a VLAN Implementation

In this activity, you will observe how broadcast traffic is forwarded by the switches when VLANs are configured and when VLANs are not configured.

Refer to
Online Course
for Illustration

3.2 VLAN Implementations

3.2.1 VLAN Assignment

3.2.1.1 VLAN Ranges on Catalyst Switches

Different Cisco Catalyst switches support various numbers of VLANs. The number of supported VLANs is large enough to accommodate the needs of most organizations. For example, the Catalyst 2960 and 3560 Series switches support over 4,000 VLANs. Normal range VLANs on these switches are numbered 1 to 1,005 and extended range VLANs are

numbered 1,006 to 4,094. The figure illustrates the available VLANs on a Catalyst 2960 switch running Cisco IOS Release 15.x.

Normal Range VLANs

- Used in small- and medium-sized business and enterprise networks.
- Identified by a VLAN ID between 1 and 1005.
- IDs 1002 through 1005 are reserved for Token Ring and FDDI VLANs.
- IDs 1 and 1002 to 1005 are automatically created and cannot be removed.
- Configurations are stored within a VLAN database file, called `vlan.dat`. The `vlan.dat` file is located in the flash memory of the switch.
- The VLAN Trunking Protocol (VTP), which helps manage VLAN configurations between switches, can only learn and store normal range VLANs.

Extended Range VLANs

- Enable service providers to extend their infrastructure to a greater number of customers. Some global enterprises could be large enough to need extended range VLAN IDs.
- Are identified by a VLAN ID between 1006 and 4094.
- Configurations are not written to the `vlan.dat` file.
- Support fewer VLAN features than normal range VLANs.
- Are, by default, saved in the running configuration file.
- VTP does not learn extended range VLANs.

Note 4096 is the upper boundary for the number of VLANs available on Catalyst switches, because there are 12 bits in the VLAN ID field of the IEEE 802.1Q header.

Refer to
Online Course
for Illustration

3.2.1.2 Creating a VLAN

When configuring normal range VLANs, the configuration details are stored in flash memory on the switch in a file called `vlan.dat`. Flash memory is persistent and does not require the `copy running-config startup-config` command. However, because other details are often configured on a Cisco switch at the same time that VLANs are created, it is good practice to save running configuration changes to the startup configuration.

Figure 1 displays the Cisco IOS command syntax used to add a VLAN to a switch and give it a name. Naming each VLAN is considered a best practice in switch configuration.

Figure 2 shows how the student VLAN (VLAN 20) is configured on switch S1. In the topology example, the student computer (PC1) has not been associated with a VLAN yet, but it does have an IP address of 172.17.20.22.

Use the Syntax Checker in Figure 3 to create a VLAN and use the `show vlan brief` command to display the contents of the `vlan.dat` file.

In addition to entering a single VLAN ID, a series of VLAN IDs can be entered separated by commas, or a range of VLAN IDs separated by hyphens using the `vlan vlan-id` command. For example, use the following command to create VLANs 100, 102, 105, 106, and 107:

```
S1(config)# vlan 100,102,105-107
```

Refer to
Online Course
for Illustration

3.2.1.3 Assigning Ports to VLANs

After creating a VLAN, the next step is to assign ports to the VLAN. An access port can belong to only one VLAN at a time; one exception to this rule is that of a port connected to an IP phone, in which case, there are two VLANs associated with the port: one for voice and one for data.

Figure 1 displays the syntax for defining a port to be an access port and assigning it to a VLAN. The `switchport mode access` command is optional, but strongly recommended as a security best practice. With this command, the interface changes to permanent access mode.

Note Use the `interface range` command to simultaneously configure multiple interfaces.

In the example in Figure 2, VLAN 20 is assigned to port F0/18 on switch S1; therefore, the student computer (PC2) is in VLAN 20. When VLAN 20 is configured on other switches, the network administrator knows to configure the other student computers to be in the same subnet as PC2 (172.17.20.0/24).

Use the Syntax Checker in Figure 3 to assign a VLAN and use the `show vlan brief` command to display the contents of the `vlan.dat` file.

The `switchport access vlan` command forces the creation of a VLAN if it does not already exist on the switch. For example, VLAN 30 is not present in the `show vlan brief` output of the switch. If the `switchport access vlan 30` command is entered on any interface with no previous configuration, then the switch displays the following:

```
% Access VLAN does not exist. Creating vlan 30
```

Refer to
Online Course
for Illustration

3.2.1.4 Changing VLAN Port Membership

There are a number of ways to change VLAN port membership. Figure 1 shows the syntax for changing a switch port to VLAN 1 membership with the `no switchport access vlan` interface configuration mode command.

Interface F0/18 was previously assigned to VLAN 20. The `no switchport access vlan` command is entered for interface F0/18. Examine the output in the `show vlan brief` command that immediately follows as shown in Figure 2. The `show vlan brief` command displays the VLAN assignment and membership type for all switch ports. The `show vlan brief` command displays one line for each VLAN. The output for each VLAN includes the VLAN name, status, and switch ports.

VLAN 20 is still active, even though no ports are assigned to it. In Figure 3, the `show interfaces f0/18 switchport` output verifies that the access VLAN for interface F0/18 has been reset to VLAN 1.

A port can easily have its VLAN membership changed. It is not necessary to first remove a port from a VLAN to change its VLAN membership. When an access port has its VLAN membership reassigned to another existing VLAN, the new VLAN membership simply replaces the previous VLAN membership. In Figure 4, port F0/11 is assigned to VLAN 20.

Use the Syntax Checker in Figure 5 to change VLAN port membership.

Refer to
Online Course
for Illustration

3.2.1.5 Deleting VLANs

In the figure, the `no vlan vlan-id` global configuration mode command is used to remove VLAN 20 from the switch. Switch S1 had a minimal configuration with all ports in VLAN 1 and an unused VLAN 20 in the VLAN database. The `show vlan brief` command verifies that VLAN 20 is no longer present in the `vlan.dat` file after using the `no vlan 20` command.

Caution Before deleting a VLAN, be sure to first reassign all member ports to a different VLAN. Any ports that are not moved to an active VLAN are unable to communicate with other hosts after the VLAN is deleted and until they are assigned to an active VLAN.

Alternatively, the entire `vlan.dat` file can be deleted using the `delete flash:vlan.dat` privileged EXEC mode command. The abbreviated command version (`delete vlan.dat`) can be used if the `vlan.dat` file has not been moved from its default location. After issuing this command and reloading the switch, the previously configured VLANs are no longer present. This effectively places the switch into its factory default condition concerning VLAN configurations.

Note For a Catalyst switch, the `erase startup-config` command must accompany the `delete vlan.dat` command prior to reload to restore the switch to its factory default condition.

Refer to
Online Course
for Illustration

3.2.1.6 Verifying VLAN Information

After a VLAN is configured, VLAN configurations can be validated using Cisco IOS `show` commands.

Figure 1 displays examples of common `show vlan` and `show interfaces` command options.

In the example in Figure 2, the `show vlan name student` command produces output that is not easily interpreted. The preferable option is to use the `show vlan brief` command. The `show vlan summary` command displays the count of all configured VLANs. The output in Figure 2 shows seven VLANs.

The `show interfaces vlan vlan-id` command displays details that are beyond the scope of this course. The important information appears on the second line in Figure 3, indicating that VLAN 20 is up.

Use the Syntax Checker in Figure 4 to display the switch port information using the `show interfaces interface-id switchport` command. This command can be used to verify VLAN assignments and mode.

Refer to Packet Tracer Activity for this chapter

3.2.1.7 Packet Tracer - Configuring VLANs

VLANs are helpful in the administration of logical groups, allowing members of a group to be easily moved, changed, or added. This activity focuses on creating and naming VLANs, and assigning access ports to specific VLANs.

Refer to Online Course for Illustration

3.2.2 VLAN Trunks

3.2.2.1 Configuring IEEE 802.1Q Trunk Links

A VLAN trunk is an OSI Layer 2 link between two switches that carries traffic for all VLANs (unless the allowed VLAN list is restricted manually or dynamically). To enable trunk links, configure the ports on either end of the physical link with parallel sets of commands.

To configure a switch port on one end of a trunk link, use the `switchport mode trunk` command. With this command, the interface changes to permanent trunking mode. The port enters into a Dynamic Trunking Protocol (DTP) negotiation to convert the link into a trunk link even if the interface connecting to it does not agree to the change. DTP is described in the next topic. In this course, the `switchport mode trunk` command is the only method implemented for trunk configuration.

The Cisco IOS command syntax to specify a native VLAN (other than VLAN 1) is shown in Figure 1. In the example, VLAN 99 is configured as the native VLAN using the `switchport trunk native vlan 99` command.

Use the Cisco IOS `switchport trunk allowed vlan vlan-list` command to specify the list of VLANs to be allowed on the trunk link.

In Figure 2, VLANs 10, 20, and 30 support the Faculty, Student, and Guest computers (PC1, PC2, and PC3). The F0/1 port on switch S1 is configured as a trunk port and forwards traffic for VLANs 10, 20, and 30. VLAN 99 is configured as the native VLAN.

Figure 3 displays the configuration of port F0/1 on switch S1 as a trunk port. The native VLAN is changed to VLAN 99 and the allowed VLAN list is restricted to 10, 20, and 30.

If the native VLAN is not allowed on the trunk link, the trunk will not allow any data traffic for the native VLAN.

Note This configuration assumes the use of Cisco Catalyst 2960 switches which automatically use 802.1Q encapsulation on trunk links. Other switches may require manual configuration of the encapsulation. Always configure both ends of a trunk link with the same native VLAN. If 802.1Q trunk configuration is not the same on both ends, Cisco IOS Software reports errors.

Refer to
Online Course
for Illustration

3.2.2.2 Resetting the Trunk to Default State

Figure 1 shows the commands to remove the allowed VLANs and reset the native VLAN of the trunk. When reset to the default state, the trunk allows all VLANs and uses VLAN 1 as the native VLAN.

The command to reset the switch port to an access port and, in effect, delete the trunk port is also shown.

Figure 2 shows the commands used to reset all trunking characteristics of a trunking interface to the default settings. The `show interfaces f0/1 switchport` command reveals that the trunk has been reconfigured to a default state.

In Figure 3, the sample output shows the commands used to remove the trunk feature from the F0/1 switch port on switch S1. The `show interfaces f0/1 switchport` command reveals that the F0/1 interface is now in static access mode.

Refer to
Online Course
for Illustration

3.2.2.3 Verifying Trunk Configuration

Figure 1 displays the configuration of switch port F0/1 on switch S1. The configuration is verified with the `show interfaces interface-ID switchport` command.

The top highlighted area shows that port F0/1 has its administrative mode set to `trunk`. The port is in trunking mode. The next highlighted area verifies that the native VLAN is VLAN 99. Further down in the output, the bottom highlighted area shows that all VLANs are enabled on the trunk.

Use the Syntax Checker in Figure 2 to configure a trunk supporting all VLANs on interface F0/1, with native VLAN 99. Verify the trunk configuration with the `show interfaces f0/1 switchport` command.

Refer to Packet
Tracer Activity
for this chapter

3.2.2.4 Packet Tracer - Configuring Trunks

Trunks are required to pass VLAN information between switches. A port on a switch is either an access port or a trunk port. Access ports carry traffic from a specific VLAN assigned to the port. A trunk port by default is a member of all VLANs; therefore, it carries traffic for all VLANs. This activity focuses on creating trunk ports, and assigning them to a native VLAN other than the default.

Refer to
Lab Activity
for this chapter

3.2.2.5 Lab - Configuring VLANs and Trunking

In this lab, you will complete the following objectives:

- Part 1: Build the Network and Configure Basic Device Settings
- Part 2: Create VLANs and Assign Switch Ports
- Part 3: Maintain VLAN Port Assignments and the VLAN Database
- Part 4: Configure an 802.1Q Trunk between the Switches
- Part 5: Delete the VLAN Database

Refer to
Online Course
for Illustration

3.2.3 Dynamic Trunking Protocol

3.2.3.1 Introduction to DTP

Ethernet trunk interfaces support different trunking modes. An interface can be set to trunking or nontrunking, or to negotiate trunking with the neighbor interface. Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which operates on a point-to-point basis only, between network devices.

DTP is a Cisco proprietary protocol that is automatically enabled on Catalyst 2960 and Catalyst 3560 Series switches. Switches from other vendors do not support DTP. DTP manages trunk negotiation only if the port on the neighbor switch is configured in a trunk mode that supports DTP.

Caution Some internetworking devices might forward DTP frames improperly, which can cause misconfigurations. To avoid this, turn off DTP on interfaces on a Cisco switch connected to devices that do not support DTP.

The default DTP configuration for Cisco Catalyst 2960 and 3560 switches is dynamic auto as shown in Figure 1 on interface F0/3 of switches S1 and S3.

To enable trunking from a Cisco switch to a device that does not support DTP, use the `switchport mode trunk` and `switchport nonegotiate` interface configuration mode commands. This causes the interface to become a trunk, but not generate DTP frames.

In Figure 2, the link between switches S1 and S2 becomes a trunk because the F0/1 ports on switches S1 and S2 are configured to ignore all DTP advertisements, and to come up in and stay in trunk port mode. The F0/3 ports on switches S1 and S3 are set to dynamic auto, so the negotiation results in the access mode state. This creates an inactive trunk link. When configuring a port to be in trunk mode, there is no ambiguity about which state the trunk is in; it is always on. With this configuration, it is easy to remember which state the trunk ports are in; if the port is supposed to be a trunk, the mode is set to trunk.

Refer to
Online Course
for Illustration

3.2.3.2 Negotiated Interface Modes

Ethernet interfaces on Catalyst 2960 and Catalyst 3560 Series switches support different trunking modes with the help of DTP:

- **switchport mode access**- Puts the interface (access port) into permanent nontrunking mode and negotiates to convert the link into a nontrunk link. The interface becomes a nontrunk interface, regardless of whether the neighboring interface is a trunk interface.
- **switchport mode dynamic auto**- Makes the interface able to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to trunk or desirable mode. The default switchport mode for all Ethernet interfaces is `dynamic auto`.
- **switchport mode dynamic desirable**- Makes the interface actively attempt to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to trunk, desirable, or auto mode. This is the default switchport mode on older switches, such as the Catalyst 2950 and 3550 Series switches.

- **switchport mode trunk**- Puts the interface into permanent trunking mode and negotiates to convert the neighboring link into a trunk link. The interface becomes a trunk interface even if the neighboring interface is not a trunk interface.
- **switchport nonegotiate**- Prevents the interface from generating DTP frames. You can use this command only when the interface switchport mode is **access** or **trunk**. You must manually configure the neighboring interface as a trunk interface to establish a trunk link.

Figure 1 illustrates the results of the DTP configuration options on opposite ends of a trunk link connected to Catalyst 2960 switch ports.

Configure trunk links statically whenever possible. The default DTP mode is dependent on the Cisco IOS Software version and on the platform. To determine the current DTP mode, issue the **show dtp interface** command as shown in Figure 2.

Use the Syntax Checker in Figure 3 to determine the DTP mode on interface F0/1.

Note A general best practice is to set the interface to **trunk** and **nonegotiate** when a trunk link is required. On links where trunking is not intended, DTP should be turned off.

Refer to
Interactive Graphic
in online course.

3.2.3.3 Activity - Predict DTP Behavior

Refer to
Online Course
for Illustration

3.2.4 Troubleshoot VLANs and Trunks

3.2.4.1 IP Addressing Issues with VLAN

Each VLAN must correspond to a unique IP subnet. If two devices in the same VLAN have different subnet addresses, they cannot communicate. This is a common problem, and it is easy to solve by identifying the incorrect configuration and changing the subnet address to the correct one.

In Figure 1, PC1 cannot connect to the Web/TFTP server shown.

A check of the IP configuration settings of PC1 shown in Figure 2 reveals the most common error in configuring VLANs: an incorrectly configured IP address. PC1 is configured with an IP address of 172.172.10.21, but it should have been configured with 172.17.10.21.

The PC1 Fast Ethernet configuration dialog box shows the updated IP address of 172.17.10.21. In Figure 3, the output on the bottom reveals that PC1 has regained connectivity to the Web/TFTP server found at IP address 172.17.10.30.

Refer to
Online Course
for Illustration

3.2.4.2 Missing VLANs

If there is still no connection between devices in a VLAN, but IP addressing issues have been ruled out, refer to the flowchart in Figure 1 to troubleshoot:

- Use the **show vlan** command to check whether the port belongs to the expected VLAN. If the port is assigned to the wrong VLAN, use the **switchport access vlan** command to correct the VLAN membership. Use the **show mac address-table** command to check which addresses were learned on a particular port of the switch and to which VLAN that port is assigned.

- If the VLAN to which the port is assigned is deleted, the port becomes inactive. Use the `show vlan` or `show interfaces switchport` command.

To display the MAC address table, use the `show mac-address-table` command. The example in Figure 2 shows MAC addresses that were learned on the F0/1 interface. It can be seen that MAC address `000c.296a.a21c` was learned on interface F0/1 in VLAN 10. If this number is not the expected VLAN number, change the port VLAN membership using the `switchport access vlan` command.

Each port in a switch belongs to a VLAN. If the VLAN to which the port belongs is deleted, the port becomes inactive. All ports belonging to the VLAN that was deleted are unable to communicate with the rest of the network. Use the `show interface f0/1 switchport` command to check whether the port is inactive. If the port is inactive, it is not functional until the missing VLAN is created using the `vlan vlan_id` command.

Refer to
Online Course
for Illustration

3.2.4.3 Introduction to Troubleshooting Trunks

A common task of a network administrator is to troubleshoot trunk link formation or links incorrectly behaving as trunk links. Sometimes a switch port may behave like a trunk port even if it is not configured as a trunk port. For example, an access port might accept frames from VLANs different from the VLAN to which it is assigned. This is called VLAN leaking.

Figure 1 displays a flowchart of general trunk troubleshooting guidelines.

To troubleshoot issues when a trunk is not forming or when VLAN leaking is occurring, proceed as follows:

- Use the `show interfaces trunk` command to check whether the local and peer native VLANs match. If the native VLAN does not match on both sides, VLAN leaking occurs.
- Use the `show interfaces trunk` command to check whether a trunk has been established between switches. Statically configure trunk links whenever possible. Cisco Catalyst switch ports use DTP by default and attempt to negotiate a trunk link.

To display the status of the trunk, the native VLAN used on that trunk link, and verify trunk establishment, use the `show interfaces trunk` command. The example in Figure 2 shows that the native VLAN on one side of the trunk link was changed to VLAN 2. If one end of the trunk is configured as native VLAN 99 and the other end is configured as native VLAN 2, a frame sent from VLAN 99 on one side is received on VLAN 2 on the other side. VLAN 99 leaks into the VLAN 2 segment.

CDP displays a notification of a native VLAN mismatch on a trunk link with this message:

```
*Mar 1 06:45:26.232: %CDP-4-NATIVE_VLAN_MISMATCH: Native VLAN mismatch discovered
on FastEthernet0/1 (2), with S2 FastEthernet0/1 (99).
```

Connectivity issues occur in the network if a native VLAN mismatch exists. Data traffic for VLANs, other than the two native VLANs configured, successfully propagates across the trunk link, but data associated with either of the native VLANs does not successfully propagate across the trunk link.

Note The output in Figure 2 indicates that there is an active trunk despite the native VLAN mismatch. Configure the native VLAN to be the same VLAN on both sides of the link to correct this behavior.

Refer to
Online Course
for Illustration

3.2.4.4 Common Problems with Trunks

Trunking issues are usually associated with incorrect configurations. When configuring VLANs and trunks on a switched infrastructure, the following types of configuration errors are the most common:

- **Native VLAN mismatches** - Trunk ports are configured with different native VLANs. This configuration error generates console notifications, and causes control and management traffic to be misdirected. This poses a security risk.
- **Trunk mode mismatches** - One trunk port is configured with trunk mode off and the other with trunk mode on. This configuration error causes the trunk link to stop working.
- **Allowed VLANs on trunks** - The list of allowed VLANs on a trunk has not been updated with the current VLAN trunking requirements. In this situation, unexpected traffic or no traffic is being sent over the trunk.

If an issue with a trunk is discovered and if the cause is unknown, start troubleshooting by examining the trunks for a native VLAN mismatch. If that is not the cause, check for trunk mode mismatches, and finally check for the allowed VLAN list on the trunk. The next two pages examine how to fix the common problems with trunks.

Refer to
Online Course
for Illustration

3.2.4.5 Trunk Mode Mismatches

Trunk links are normally configured statically with the `switchport mode trunk` command. Cisco Catalyst switch trunk ports use DTP to negotiate the state of the link. When a port on a trunk link is configured with a trunk mode that is incompatible with the neighboring trunk port, a trunk link fails to form between the two switches.

In the scenario illustrated in Figure 1, PC4 cannot connect to the internal web server. The topology indicates a valid configuration. Why is there a problem?

Check the status of the trunk ports on switch S1 using the `show interfaces trunk` command. The output shown in Figure 2 reveals that interface Fa0/3 on switch S1 is not currently a trunk link. Examining the F0/3 interface reveals that the switch port is actually in dynamic auto mode. An examination of the trunks on switch S3 reveals that there are no active trunk ports. Further checking reveals that the Fa0/3 interface is also in dynamic auto mode. This explains why the trunk is down.

To resolve the issue, reconfigure the trunk mode of the F0/3 ports on switches S1 and S3, as shown in Figure 3. After the configuration change, the output of the `show interfaces` command indicates that the port on switch S1 is now in trunking mode. The output from PC4 indicates that it has regained connectivity to the Web/TFTP server found at IP address 172.17.10.30.

Refer to
Online Course
for Illustration

3.2.4.6 Incorrect VLAN List

For traffic from a VLAN to be transmitted across a trunk, it must be allowed on the trunk. To do so, use the `switchport trunk allowed vlan vlan-id` command.

In Figure 1, VLAN 20 (Student) and PC5 have been added to the network. The documentation has been updated to show that the VLANs allowed on the trunk are 10, 20, and 99. In this scenario, PC5 cannot connect to the student email server.

Check the trunk ports on switch S1 using the `show interfaces trunk` command as shown in Figure 2. The command reveals that the interface F0/3 on switch S3 is correctly configured to allow VLANs 10, 20, and 99. An examination of the F0/3 interface on switch S1 reveals that interfaces F0/1 and F0/3 only allow VLANs 10 and 99. Someone updated the documentation but forgot to reconfigure the ports on the S1 switch.

Reconfigure F0/1 and F0/3 on switch S1 using the `switchport trunk allowed vlan 10,20,99` command as shown in Figure 3. The output shows that VLANs 10, 20, and 99 are now added to the F0/1 and F0/3 ports on switch S1. The `show interfaces trunk` command is an excellent tool for revealing common trunking problems. PC5 has regained connectivity to the student email server found at IP address 172.17.20.10.

Refer to Packet
Tracer Activity
for this chapter

3.2.4.7 Packet Tracer - Troubleshooting a VLAN Implementation - Scenario 1

In this activity, you will troubleshoot connectivity problems between PCs on the same VLAN. The activity is complete when PCs on the same VLAN can ping each other. Any solution you implement must conform to the Addressing Table.

Refer to Packet
Tracer Activity
for this chapter

3.2.4.8 Packet Tracer - Troubleshooting a VLAN Implementation - Scenario 2

In this activity, you will troubleshoot a misconfigured VLAN environment. The initial network has errors. Your objective is to locate and correct the errors in the configurations and establish end-to-end connectivity. Your final configuration should match the Topology diagram and Addressing Table. The native VLAN for this topology is VLAN 56.

Refer to
Lab Activity
for this chapter

3.2.4.9 Lab - Troubleshooting VLAN Configurations

In this lab, you will complete the following objectives:

- Part 1: Build the Network and Configure Basic Device Settings
- Part 2: Troubleshoot VLAN 10
- Part 3: Troubleshoot VLAN 20

Refer to
Online Course
for Illustration

3.3 VLAN Security and Design

3.3.1 Attacks on VLANs

3.3.1.1 Switch Spoofing Attack

There are a number of different types of VLAN attacks in modern switched networks. The VLAN architecture simplifies network maintenance and improves performance, but it also opens the door to abuse. It is important to understand the general methodology behind these attacks and the primary approaches to mitigate them.

VLAN hopping enables traffic from one VLAN to be seen by another VLAN. Switch spoofing is a type of VLAN hopping attack that works by taking advantage of an incorrectly configured trunk port. By default, trunk ports have access to all VLANs and pass traffic for multiple VLANs across the same physical link, generally between switches.

Click the Play button in the figure to see an animation of a switch spoofing attack.

In a basic switch spoofing attack, the attacker takes advantage of the fact that the default configuration of the switch port is dynamic auto. The network attacker configures a system to spoof itself as a switch. This spoofing requires that the network attacker be capable of emulating 802.1Q and DTP messages. By tricking a switch into thinking that another switch is attempting to form a trunk, an attacker can gain access to all the VLANs allowed on the trunk port.

The best way to prevent a basic switch spoofing attack is to turn off trunking on all ports, except the ones that specifically require trunking. On the required trunking ports, disable DTP, and manually enable trunking.

Refer to
Online Course
for Illustration

3.3.1.2 Double-Tagging Attack

Another type of VLAN attack is a double-tagging (or double-encapsulated) VLAN hopping attack. This type of attack takes advantage of the way that hardware on most switches operates. Most switches perform only one level of 802.1Q de-encapsulation, which allows an attacker to embed a hidden 802.1Q tag inside the frame. This tag allows the frame to be forwarded to a VLAN that the original 802.1Q tag did not specify. An important characteristic of the double-encapsulated VLAN hopping attack is that it works even if trunk ports are disabled, because a host typically sends a frame on a segment that is not a trunk link.

A double-tagging VLAN hopping attack follows three steps:

1. The attacker sends a double-tagged 802.1Q frame to the switch. The outer header has the VLAN tag of the attacker, which is the same as the native VLAN of the trunk port. The assumption is that the switch processes the frame received from the attacker as if it were on a trunk port or a port with a voice VLAN (a switch should not receive a tagged Ethernet frame on an access port). For the purposes of this example, assume that the native VLAN is VLAN 10. The inner tag is the victim VLAN; in this case, VLAN 20.
2. The frame arrives on the switch, which looks at the first 4-byte 802.1Q tag. The switch sees that the frame is destined for VLAN 10, which is the native VLAN. The switch forwards the packet out on all VLAN 10 ports after stripping the VLAN 10 tag. On

the trunk port, the VLAN 10 tag is stripped, and the packet is not retagged because it is part of the native VLAN. At this point, the VLAN 20 tag is still intact and has not been inspected by the first switch.

3. The second switch looks only at the inner 802.1Q tag that the attacker sent and sees that the frame is destined for VLAN 20, the target VLAN. The second switch sends the frame on to the victim port or floods it, depending on whether there is an existing MAC address table entry for the victim host.

This type of attack is unidirectional and works only when the attacker is connected to a port residing in the same VLAN as the native VLAN of the trunk port. Thwarting this type of attack is not as easy as stopping basic VLAN hopping attacks.

The best approach to mitigating double-tagging attacks is to ensure that the native VLAN of the trunk ports is different from the VLAN of any user ports. In fact, it is considered a security best practice to use a fixed VLAN that is distinct from all user VLANs in the switched network as the native VLAN for all 802.1Q trunks.

Refer to
Online Course
for Illustration

3.3.1.3 PVLAN Edge

Some applications require that no traffic be forwarded at Layer 2 between ports on the same switch so that one neighbor does not see the traffic generated by another neighbor. In such an environment, the use of the Private VLAN (PVLAN) Edge feature, also known as protected ports, ensures that there is no exchange of unicast, broadcast, or multicast traffic between these ports on the switch (Figure 1).

The PVLAN Edge feature has the following characteristics:

- A protected port does not forward any traffic (unicast, multicast, or broadcast) to any other port that is also a protected port, except for control traffic. Data traffic cannot be forwarded between protected ports at Layer 2.
- Forwarding behavior between a protected port and a nonprotected port proceeds as usual.
- Protected ports must be manually configured.

To configure the PVLAN Edge feature, enter the `switchport protected` command in interface configuration mode (Figure 2). To disable protected port, use the `no switchport protected` interface configuration mode command. To verify the configuration of the PVLAN Edge feature, use the `show interfaces interface-id switchport` global configuration mode command.

Use the Syntax Checker in Figure 3 to configure the PVLAN Edge feature on interface G0/1 and verify the configuration.

Refer to
Interactive Graphic
in online course.

3.3.1.4 Activity - Identify the Type of VLAN Attacks

Refer to
Online Course
for Illustration

3.3.2 Design Best Practices for VLANs

3.3.2.1 VLAN Design Guidelines

Cisco switches have a factory configuration in which default VLANs are preconfigured to support various media and protocol types. The default Ethernet VLAN is VLAN 1. It

is a security best practice to configure all the ports on all switches to be associated with VLANs other than VLAN 1. This is usually done by configuring all unused ports to a black hole VLAN that is not used for anything on the network. All used ports are associated with VLANs distinct from VLAN 1 and distinct from the black hole VLAN. It is also a good practice to shut down unused switch ports to prevent unauthorized access.

A good security practice is to separate management and user data traffic. The management VLAN, which is VLAN 1 by default, should be changed to a separate, distinct VLAN. To communicate remotely with a Cisco switch for management purposes, the switch must have an IP address configured on the management VLAN. Users in other VLANs would not be able to establish remote access sessions to the switch unless they were routed into the management VLAN, providing an additional layer of security. Also, the switch should be configured to accept only encrypted SSH sessions for remote management.

All control traffic is sent on VLAN 1. Therefore, when the native VLAN is changed to something other than VLAN 1, all control traffic is tagged on IEEE 802.1Q VLAN trunks (tagged with VLAN ID 1). A recommended security practice is to change the native VLAN to a different VLAN than VLAN 1. The native VLAN should also be distinct from all user VLANs. Ensure that the native VLAN for an 802.1Q trunk is the same on both ends of the trunk link.

DTP offers four switch port modes: access, trunk, dynamic auto, and dynamic desirable. A general guideline is to disable autonegotiation. As a port security best practice, do not use the dynamic auto or dynamic desirable switch port modes.

Finally, voice traffic has stringent QoS requirements. If user PCs and IP phones are on the same VLAN, each tries to use the available bandwidth without considering the other device. To avoid this conflict, it is good practice to use separate VLANs for IP telephony and data traffic.

Refer to
Lab Activity
for this chapter

3.3.2.2 Lab - Implementing VLAN Security

In this lab, you will complete the following objectives:

- Part 1: Build the Network and Configure Basic Device Settings
- Part 2: Implement VLAN Security on the Switches

Refer to
Online Course
for Illustration

3.4 Summary

Refer to
Lab Activity
for this chapter

3.4.1.1 Class Activity VLAN Plan

VLAN Plan

Scenario

You are designing a VLAN switched network for your small- to medium- sized business.

Your business owns space on two floors of a high-rise building. The following elements need VLAN consideration and access for planning purposes:

- Management
- Finance
- Sales
- Human Resources
- Network administrator
- General visitors to your business location

You have two Cisco 3560-24PS switches.

Use a word processing software program to design your VLAN-switched network scheme.

Section 1 of your design should include the regular names of your departments, suggested VLAN names and numbers, and which switch ports would be assigned to each VLAN.

Section 2 of your design should list how security would be planned for this switched network.

When your VLAN plan is finished, complete the reflection questions from this activity's PDF.

Save your work. Be able to explain and discuss your VLAN design with another group or with the class.

Refer to **Packet Tracer Activity** for this chapter

3.4.1.2 Packet Tracer - Skills Integration Challenge

In this activity, two switches are completely configured. On a third switch, you are responsible for assigning IP addressing to the SVI, configuring VLANs, assigning VLANs to interfaces, configuring trunking, and performing basic switch security.

Refer to **Online Course** for illustration

3.4.1.3 Summary

This chapter introduced VLANs. VLANs are based on logical connections, instead of physical connections. VLANs are a mechanism to allow network administrators to create logical broadcast domains that can span across a single switch or multiple switches, regardless of physical proximity. This function is useful to reduce the size of broadcast domains or to allow groups or users to be logically grouped without the need to be physically located in the same place.

There are several types of VLANs:

- Default VLAN
- Management VLAN
- Native VLAN
- User/Data VLANs
- Black Hole VLAN
- Voice VLAN

On a Cisco switch, VLAN 1 is the default Ethernet VLAN, the default native VLAN, and the default management VLAN. Best practices suggest that the native and management VLANs be moved to another distinct VLAN and that unused switch ports be moved to a “black hole” VLAN for increased security.

The `switchport access vlan` command is used to create a VLAN on a switch. After creating a VLAN, the next step is to assign ports to the VLAN. The `show vlan brief` command displays the VLAN assignment and membership type for all switch ports. Each VLAN must correspond to a unique IP subnet.

Use the `show vlan` command to check whether the port belongs to the expected VLAN. If the port is assigned to the wrong VLAN, use the `switchport access vlan` command to correct the VLAN membership. Use the `show mac address-table` command to check which addresses were learned on a particular port of the switch and to which VLAN that port is assigned.

A port on a switch is either an access port or a trunk port. Access ports carry traffic from a specific VLAN assigned to the port. A trunk port by default is a member of all VLANs; therefore, it carries traffic for all VLANs.

VLAN trunks facilitate inter-switch communication by carrying traffic associated with multiple VLANs. IEEE 802.1Q frame tagging differentiates between Ethernet frames associated with distinct VLANs as they traverse common trunk links. To enable trunk links, use the `switchport mode trunk` command. Use the `show interfaces trunk` command to check whether a trunk has been established between switches.

Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which operates on a point-to-point basis only, between network devices. DTP is a Cisco proprietary protocol that is automatically enabled on Catalyst 2960 and Catalyst 3560 Series switches.

To place a switch into its factory default condition with 1 default VLAN, use the command `delete flash:vlan.dat and erase startup-config`.

This chapter also examined the configuration, verification, and troubleshooting of VLANs and trunks using the Cisco IOS CLI and explored basic security and design considerations in the context of VLANs.

Go to the online course to take the quiz and exam.

Chapter 3 Quiz

This quiz is designed to provide an additional opportunity to practice the skills and knowledge presented in the chapter and to prepare for the chapter exam. You will be allowed multiple attempts and the grade does not appear in the gradebook.

Chapter 3 Exam

The chapter exam assesses your knowledge of the chapter content.

Your Chapter Notes

