CHAPTER 3

Switching in an Enterprise Network

Objectives

Upon completion of this chapter, you should be able to answer the following questions:

- What types of switches are found in an enterprise network?
- How does Spanning Tree Protocol prevent switching loops?
- What is a VLAN and what purpose does it serve?
- How is a VLAN configured on a Cisco switch?
- What is inter-VLAN routing and how is it configured?
- What is VLAN Trunking Protocol and how does it help maintain VLANs in an enterprise network?

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary.

wire speed	*page 67*
content addressable memory (CAM)	*page 67*
aging time	*page 67*
microsegmentation	*page 67*
virtual circuit	*page 68*
symmetric switching	*page 68*
asymmetric switching	*page 68*
uplink ports	*page 68*
multilayer switching	*page 69*
application-specific integrated circuit (ASIC)	*page 69*
store-and-forward switching	*page 70*
cut-through switching	*page 70*
fast-forward switching	*page 70*
fragment-free switching	*page 70*
runts	*page 70*
adaptive cut-through switching	*page 70*
Redundancy	*page 72*
congestion	*page 72*
availability	*page 72*
load balancing	*page 72*
switching loops	*page 72*
broadcast storm	*page 72*
Spanning Tree Protocol (STP)	*page 75*
bridge protocol data units (BPDU)	*page 76*
blocking state	*page 77*
listening state	*page 77*
learning state	*page 77*
forwarding state	*page 77*
disabled	*page 77*
Access ports	*page 77*
Trunking ports	*page 77*
root bridge	*page 78*
bridge	*page 78*
ID (BID)	*page 78*
Root ports	*page 79*
least-cost path	*page 79*
Designated ports	*page 79*

continues

continued

Blocked ports page 79

PortFast page 80

UplinkFast page 80

BackboneFast page 81

Rapid Spanning Tree Protocol (RSTP) page 85

discarding state page 85

active topology page 85

virtual local-area network (VLAN) page 85

VLAN management policy server (VMPS) page 87

management VLAN page 87

VLAN ID (VID) page 91

frame tagging page 91

IEEE 802.1Q page 91

dot1q page 91

Inter-Switch Link (ISL) page 93

native VLAN page 94

subinterfaces page 95

router-on-a-stick page 95

VLAN Trunking Protocol (VTP) page 97

VTP server page 98

VTP client page 98

VTP transparent page 98

summary advertisements page 99

subset advertisement page 99

advertisement request page 99

Enterprise networks rely on switches in the access, distribution, and core layers to provide network segmentation and high-speed connectivity between users and networks. Spanning Tree Protocol (STP) is used in a hierarchical network to prevent switching loops that can seriously degrade network performance. Virtual LANs logically segment networks and contain broadcasts to improve network performance and security. Switches configured with trunking enable VLANs to span multiple geographic locations. VLAN Trunking Protocol (VTP) is used to simplify the configuration and management of VLANs in a complex, enterprise-level switched network.

Part II of this book includes the corresponding labs for this chapter.

Describing Enterprise-Level Switching

Although you can create an enterprise network with both routers and switches, the network design of most enterprises relies heavily on switches. Switches are cheaper per port than routers and provide fast forwarding of frames at ***wire speed***. Transmission at wire speed indicates that little overhead is associated with the transmission and that it occurs at the maximum speed of the hardware.

Switching and Network Segmentation

A switch is a very adaptable Layer 2 device. In its simplest role, it is used to replace a hub as the central point of connection for multiple hosts. In a more complex role, a switch connects to one or more other switches to create, manage, and maintain redundant links and VLAN connectivity. Regardless of the role a switch plays in a network, it processes all types of traffic in the same way.

A switch moves traffic based on MAC addresses. Each switch maintains a MAC address table in high-speed memory, called ***content addressable memory (CAM)*****.** The switch re-creates this table every time it is activated, using the source MAC addresses of incoming frames and the port number through which the frame entered the switch. The switch deletes entries from the MAC address table if they are not used within a certain period of time. The name given to this period is the ***aging time***; removal of an entry is called aging out.

As a unicast frame enters a port, the switch finds the source MAC address in the frame. It then searches the MAC table, looking for an entry that matches the address. If the source MAC address is not in the table, the switch adds a MAC address and port number entry and sets the aging timer. If the source MAC address already exists, the switch resets the aging timer associated with that entry. Next, the switch checks the table for the destination MAC address. If an entry exists, the switch forwards the frame out the appropriate port number. If the entry does not exist, the switch floods the frame out every active port except the port on which the frame was received.

High availability, speed, and throughput of the network are critical in an enterprise environment. These variables are affected by the size of the broadcast domain and the collision domain. In general, larger broadcast and collision domains negatively impact these mission-critical variables.

If a switch receives a broadcast frame, the switch floods it out every active interface, just as it does for an unknown destination MAC address. All devices that receive this broadcast make up the broadcast domain. As more switches are connected together, the size of the broadcast domain increases.

Collision domains create a similar problem. The more devices participating in a collision domain, the more collisions occur and the slower the throughput. Hubs create large collision domains. Switches, however, use a feature called ***microsegmentation*** to reduce the size of collision domains to a single switch port. When a host connects to a switch port, the switch creates a dedicated connection. When two connected hosts communicate with each other, the switch consults the switching table and establishes a virtual circuit, or microsegment, between the ports. The switch maintains the ***virtual circuit*** until the session terminates. Multiple virtual circuits can be active at the same time.

This process improves bandwidth utilization by reducing collisions and by allowing multiple simultaneous connections. Figure 3-1 shows the difference between a network that uses a hub versus one that uses a switch to connect hosts.

Figure 3-1 Connecting Hosts Using a Hub or a Switch

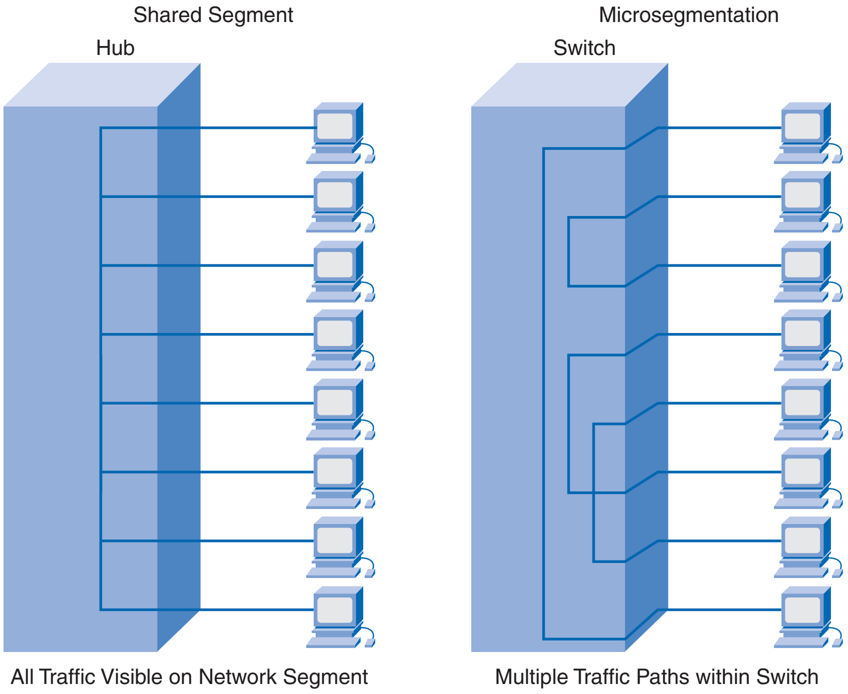
Shared Segment
Hub
Microsegmentation
Switch
All Traffic Visible on Network Segment
Multiple Traffic Paths within Switch

Many switches can support both ***symmetric switching*** and ***asymmetric switching***. Switches that have ports of all the same speeds are termed symmetric. Many switches, however, have two or more high-speed ports. These high-speed, or ***uplink ports***, connect to areas that have a higher demand for bandwidth. Typically, these areas include server farms or other networks. Connections between ports of different speeds use asymmetric switching. If necessary, a switch stores information in memory to provide a buffer between ports of different speeds. Asymmetric switches are common in the enterprise environment.

Interactive Activity 3-1: Switch Frame Forwarding (3.1.1)

In this activity, you determine to which ports a frame will be forwarded based on the information in the switch MAC table. Use file d3ia-311 on the CD-ROM that accompanies this book to perform this interactive activity.

Multilayer Switching

Traditionally, networks have been composed of separate Layer 2 and Layer 3 devices. Each device uses a different technique for processing and forwarding traffic and has a very specific role in the network design and functionality. Figure 3-2 compares Layer 2 switching and Layer 3 routing.

Figure 3-2 Layer 2 Switching and Layer 3 Routing

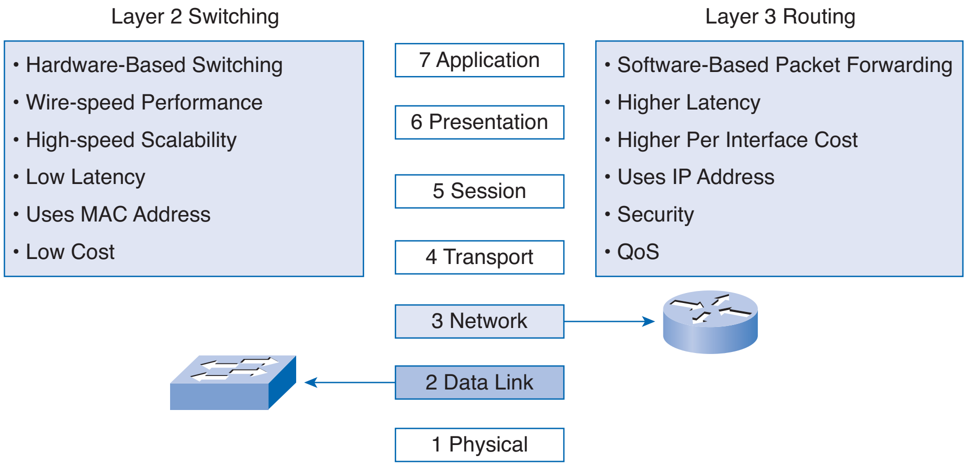

Layer 2

Layer 2 switches are hardware based. They forward traffic at wire speeds, using the internal circuits that physically connect each incoming port to every other port. The forwarding process uses the MAC address and relies on the existence of the destination MAC address in the MAC address table. A Layer 2 switch limits the forwarding of traffic to within a single network segment or subnet. Traffic that must pass from one network segment to another must pass through a Layer 3 device.

Layer 3

Routers are software based and use microprocessors to execute routing based on IP addresses. Layer 3 routing allows traffic to be forwarded between different networks and subnets. As a packet enters a router interface, the router uses software to find the destination IP address and select the best path toward the destination network. The router then switches the packet to the correct output interface.

Layer 3 switching, or ***multilayer switching***, combines hardware-based switching and hardware-based routing in the same device. A multilayer switch combines the features of a Layer 2 switch and a Layer 3 router. Layer 3 switching occurs in special ***application-specific integrated circuit (ASIC)*** hardware. The frame- and packet-forwarding functions use the same ASIC circuitry.

Multilayer switches often save, or cache, source and destination routing information from the first packet of a conversation. Subsequent packets do not have to execute a routing lookup, because they find the routing information in memory. This caching feature adds to the high performance of these devices.

Types of Switching

When switching was first introduced, a switch could support one of two major methods to forward a frame from one port to another. The two methods are store-and-forward and cut-through switching. Each of these methods has distinct advantages as well as some disadvantages. With recent advances in the speed of switching hardware, store-and-forward techniques have become the standard in many networked environments.

Store-and-Forward

In ***store-and-forward switching***, the entire frame is read and stored in memory before being sent to the destination device. The switch checks the integrity of the bits in the frame by recalculating the cyclic redundancy check (CRC) value. If the calculated CRC value is the same as the CRC field value in the frame, the switch forwards the frame out the destination port. The switch does not forward frames if the CRC values do not match. The CRC value is located within the frame check sequence (FCS) field of an Ethernet frame.

Although this method keeps damaged frames from being switched to other network segments, it introduces the highest amount of latency of any of the switching technologies. Because of the latency incurred by the store-and-forward method, it is typically only used in environments where errors are likely to occur. For example, an environment that has a high probability of electromagnetic interference (EMI) would create a large number of defective frames and would be appropriate for store-and-forward switching.

Cut-Through Switching

The other major method of switching is ***cut-through switching***. Cut-through switching subdivides into two other methods: ***fast-forward switching*** and ***fragment-free switching***. In both of these methods the switch forwards the frame before all of it is received. Because the switch does not calculate or check the CRC value before forwarding the frame, damaged frames can be switched.

Fast-forward is the fastest method of switching. The switch forwards the frames out the destination port as soon as it reads the destination MAC address. This method has the lowest latency but also forwards collision fragments and damaged frames. This method of switching works best in a stable network with few errors.

In fragment-free switching, the switch reads the first 64 bytes of the frame before it begins to forward it out the destination port. The shortest valid Ethernet frame is 64 bytes. Smaller frames are usually the result of a collision and are called ***runts***. Checking the first 64 bytes ensures that the switch does not forward collision fragments.

Store-and-forward has the highest latency and fast-forward has the lowest. The latency introduced by fragment-free switching is in the middle of these other methods. The fragment-free switching method works best in an environment where many collisions occur. In a properly constructed switched network, collisions are not a problem; therefore, fast-forward switching would be the preferred method.

Some newer Layer 2 and Layer 3 switches can adapt their switching method to changing network conditions. This is known as ***adaptive cut-through switching***. The switches begin by forwarding traffic using the fast-forward method to achieve the lowest latency possible. Even though the switch does not check for errors before forwarding the frame, it recognizes errors in the frames as they pass through the switch. The switch stores this value in an error counter in memory. It compares the number of errors found to a predefined threshold value. If the number of errors exceeds the threshold value, the switch has forwarded an unacceptable number of errors. In this situation, the switch modifies itself to perform store-and-forward switching. If the number of errors drops back below the threshold, the switch reverts back to fast-forward mode.

Switch Security

It is important to keep your network secure, regardless of the switching method used. Network security often focuses on routers and blocking traffic from the outside. Switches are internal to the organization and designed to allow ease of connectivity. For this reason, only limited or no security measures are applied in many switched environments. The following list includes some of the security measures that you should take to ensure that only authorized people have access to the switches in a network:

- **Physically secure the device.** Switches are a critical link in the network. Secure them physically, by mounting them in a rack and installing the rack in a secure room. Limit access to authorized network staff.
- **Use secure passwords.** Configure all passwords (user mode, privilege mode, and vty access) with a minimum of six nonrepeating characters. Change passwords on a regular basis. Never use words found in a dictionary.

  Use the **enable secret command** for privileged-level password protection, because it uses advanced encryption techniques. Encrypt all passwords in the display of the running configuration file using the IOS command **service password-encryption**.
- **Enable SSH access.** Secure Shell (SSH) is a client-server protocol used to log in to another device over a network. It provides strong authentication and secure communication over insecure channels. SSH encrypts the entire login session, including password transmission.
- **Monitor access and traffic.** Monitor all traffic passing through a switch to ensure that it complies with company policies. Additionally, record the MAC address of all devices connecting to a specific switch port and all login attempts on the switch. If the switch detects malicious traffic or unauthorized access, take action according to the security policy of the organization.
- **Disable HTTP access.** Disable HTTP access so that no one modifies the switch configuration through the web. The command to disable HTTP access is **no ip http server.**
- **Disable unused ports.** Disable all unused ports on the switch to prevent unknown PCs or wireless access points from connecting to an available port on the switch. Accomplish this by issuing a **shutdown** command on the interface.
- **Enable port security.** Port security restricts access to a switch port to a specific list of MAC addresses. Enter the MAC addresses manually or have the switch learn them dynamically. The specific switch port associates with the MAC addresses, allowing only traffic from those devices. If a device with a different MAC address plugs into the port, the switch automatically disables the port.
- **Disable Telnet.** A Telnet connection sends data over the public network in clear text. This includes usernames, passwords, and data. Disable Telnet access to all networking devices by not configuring a password for any vty sessions at login.

Lab 3-1: Applying Basic Switch Security (3.1.4)

In this lab, you configure and test basic switch security. Refer to the lab in Part II of this *Learning Guide*. You can perform this lab now or wait until the end of the chapter.

Preventing Switching Loops

Modern enterprises rely more and more on their networks for their very existence. The network is the lifeline of many organizations. Network downtime translates into potentially disastrous loss of business, income, and customer confidence. The failure of a single network link, a single device, or a critical port on a switch causes network downtime.

Redundancy in a Switched Network

Redundancy is required in the network design to maintain a high degree of reliability and eliminate any single point of failure. Redundancy is accomplished by installing duplicate equipment and network links for critical areas. Figure 3-3 shows a network that incorporates redundancy.

Figure 3-3 Redundancy in a Network

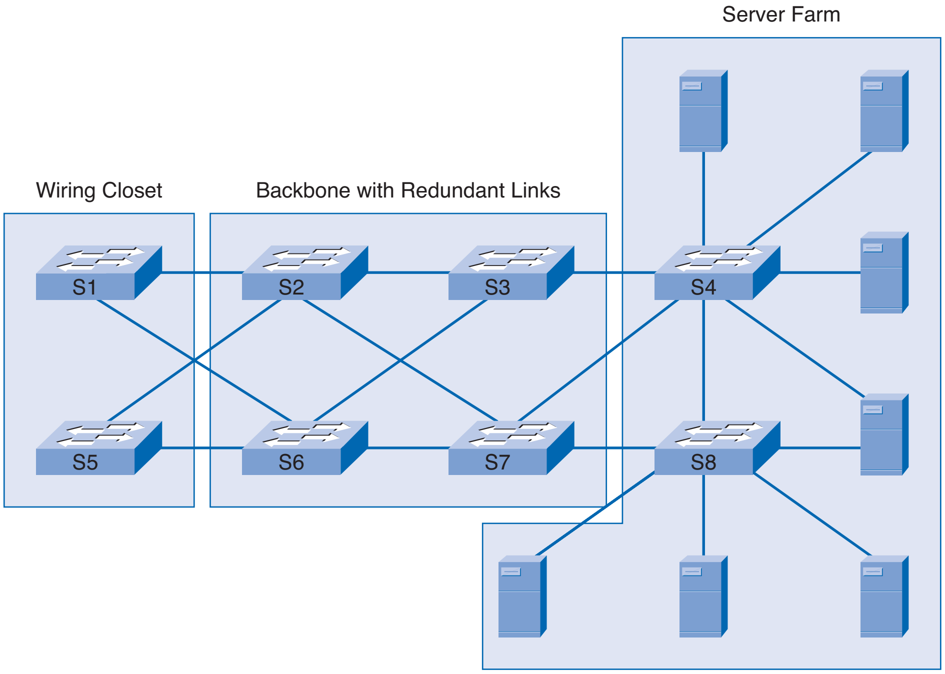

Sometimes, providing complete redundancy of all links and devices in a network becomes very expensive. Network engineers are often required to balance the cost of redundancy with the need for network availability. Network downtime translates into potential loss of business, income, and customer confidence. This loss must be assessed in the context of the business environment. Each company or organization can tolerate different levels of network outage. For most enterprise environments a 99.999 percent uptime is expected, and the network must be designed to provide this level of reliability.

Redundancy refers to having two different pathways to a particular destination. Examples of redundancy in nonnetworking environments include two roads into a town, two bridges to cross a river, or two doors to exit a building. If one way is blocked, another is still available. Redundancy in a switched network is achieved by connecting switches with multiple links. Redundant links in a switched network reduce ***congestion*** and support high ***availability*** and ***load balancing***.

Connecting switches together, however, can cause problems. For example, the broadcast nature of Ethernet traffic creates ***switching loops***. The broadcast frames go around and around in all directions, causing a ***broadcast storm***, as shown in Figure 3-4.

In this example, a host sends a broadcast into a switched network. Any switch that receives the broadcast sends it out all ports except the one on which it was originally received. In this case, the first switch forwards the broadcast message across multiple links to a second switch. The second switch repeats the process and forwards the broadcast message back to the first switch. The first switch then forwards the message to the second switch and the process continues to repeat itself, consuming large amounts of bandwidth and creating a broadcast storm.

Figure 3-4 Broadcast Storm

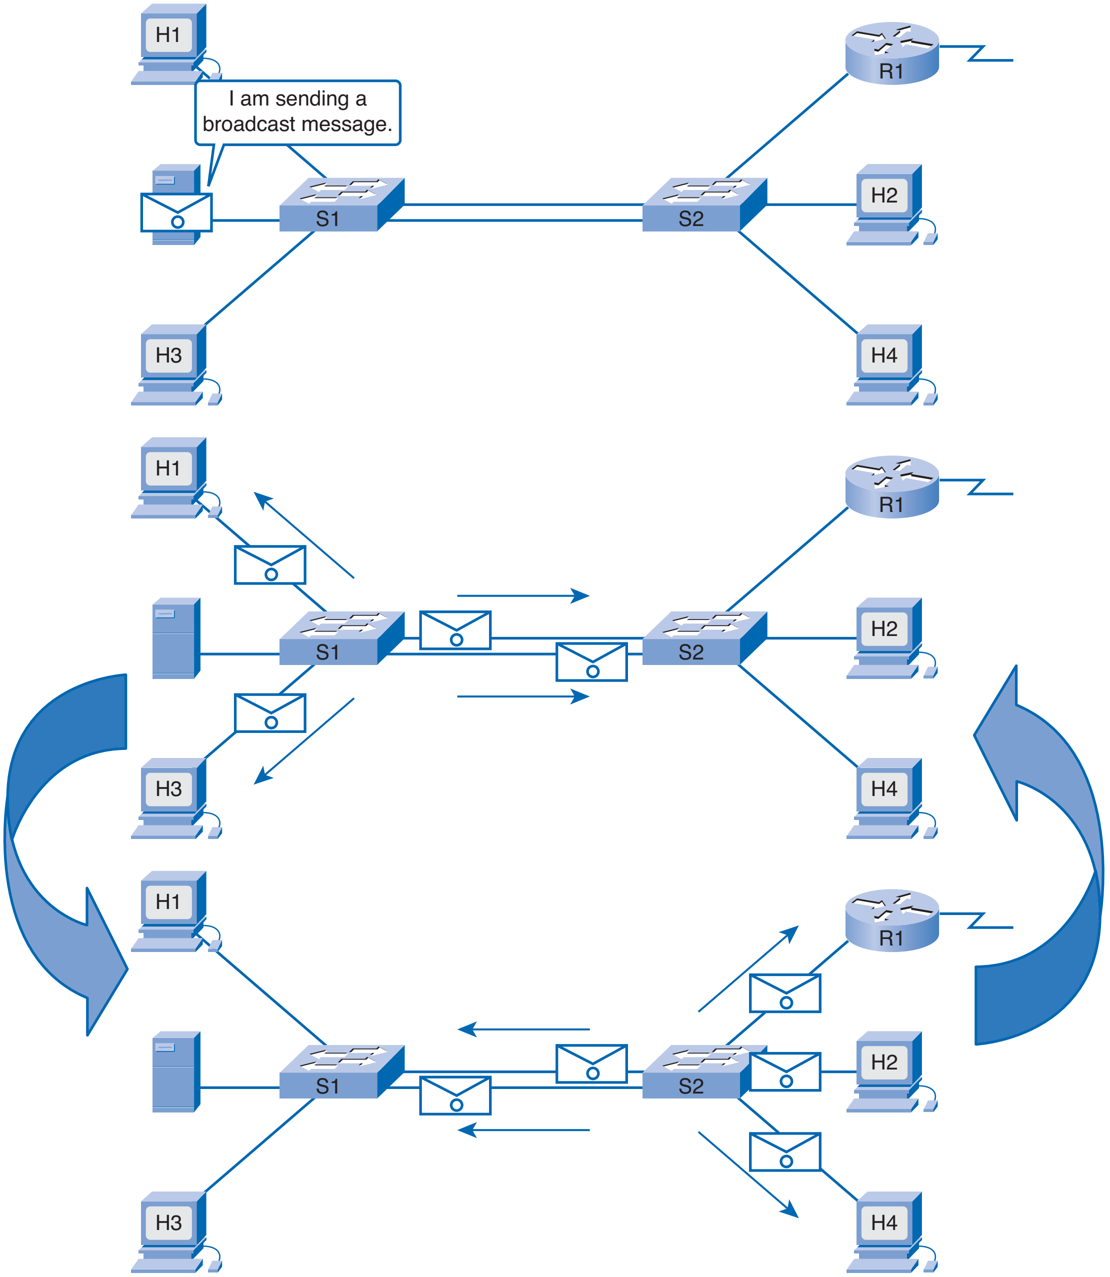

Broadcast storms use up all the available bandwidth, can prevent network connections from being established, and can cause existing network connections to be dropped.

Broadcast storms are not the only problem created by redundant links in a switched network. Unicast frames sometimes produce problems, such as multiple frame transmissions and MAC database instability.

Multiple Frame Transmissions

If a host sends a unicast frame to a destination host and the destination MAC address is not included in any of the connected switch MAC tables, every switch floods the frame out all ports. The process repeats, creating multiple copies of the frame on the network. Eventually the destination host receives multiple copies of the frame, as shown in Figure 3-5. In this example, a copy of the frame is sent across each link to the destination switch. The destination switch receives two copies of the frame addresses to the destination host and forwards both. This causes three problems: wasted bandwidth, wasted CPU time, and potential duplication of transaction traffic. Imagine the problems that could be caused if two invoices were issued or two buy requests placed in the stock market because of multiple frame transmissions.

Figure 3-5 Multiple Frame Transmission on a Looped Network

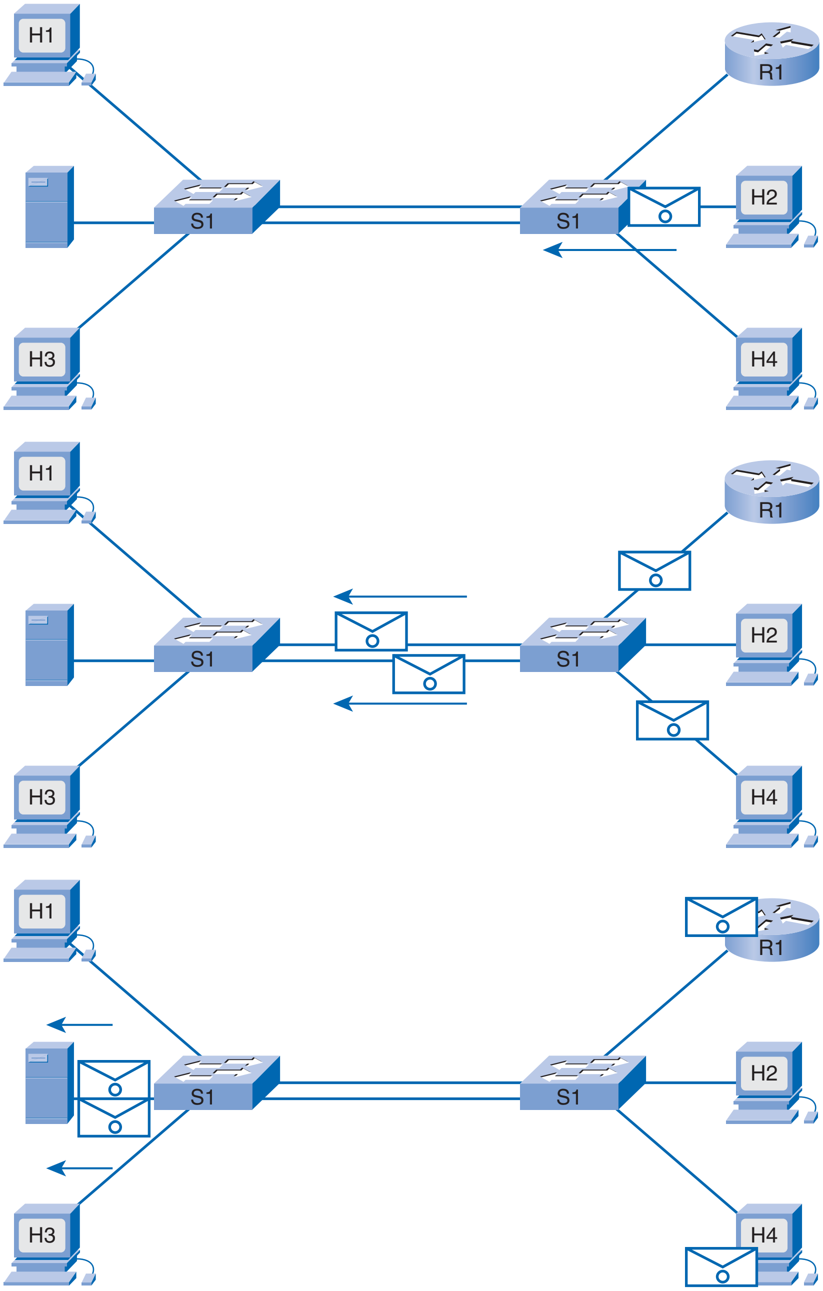

MAC Database Instability

Switches in a redundant network can learn the wrong information about the location of a host. If a loop exists, one switch might associate the destination MAC address with two separate ports. This is because the switch receives information from the same source on two different ports, causing the switch to continually update its MAC address table. This causes suboptimal forwarding of frames.

Packet Tracer
□ **Activity**

Disabling Redundant Links to Avoid Switching Loops (3.2.1)

In this activity, you disable redundant links in a network to avoid switching loops. Use file d3-321.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

Spanning Tree Protocol (STP)

Spanning Tree Protocol (STP) provides a mechanism for disabling redundant links in a switched network. STP provides the redundancy required for reliability without creating switching loops. STP is an open standard protocol, used in a switched environment to create a loop-free logical topology.

STP is relatively self-sufficient and requires little configuration. When switches are first powered up with STP enabled, they check the switched network for the existence of loops. Switches detecting a potential loop block some of the connecting ports, while leaving other ports active to forward frames, as shown in Figure 3-6.

Figure 3-6 Spanning Tree Protocol Preventing a Switching Loop

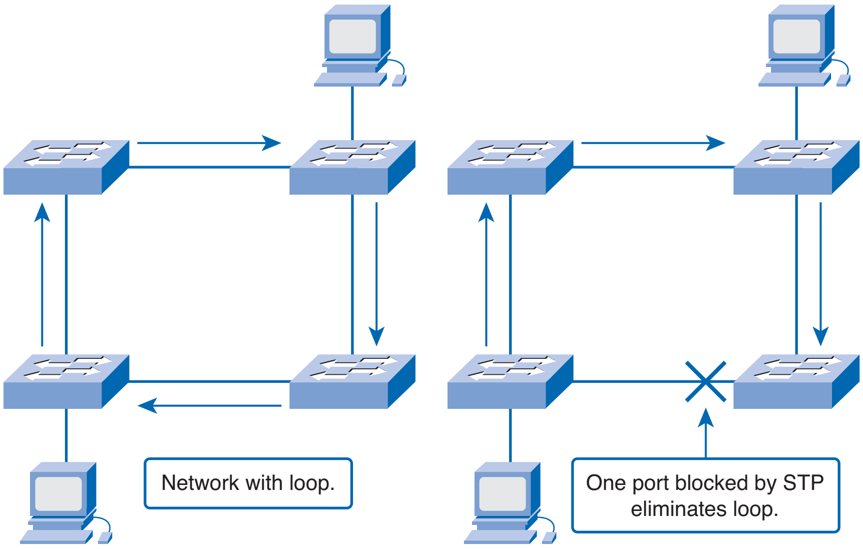

STP defines a tree that spans all the switches in an extended star switched network. Switches are constantly checking the network to ensure that no loops exist and that all ports function as required. To prevent switching loops, STP does the following:

- Forces certain interfaces into a blocked state
- Leaves other interfaces in a forwarding state
- Reconfigures the network by activating the appropriate path, if the forwarding path becomes unavailable

In STP terminology, the term *bridge* is frequently used to refer to a switch. For example, the root bridge is the primary switch or focal point in the STP topology. The root bridge communicates with the other switches using ***bridge protocol data units (BPDU)***. BPDUs are frames that multicast every 2 seconds to all other switches. BPDUs contain information such as

- Identity of the source switch
- Identity of the source port
- Cumulative cost to the root bridge
- Value of aging timers
- Value of the hello timer

The structure of a BPDU is shown in Figure 3-7, and the individual fields are described in Table 3-1.

Figure 3-7 Bridge Protocol Data Unit Structure

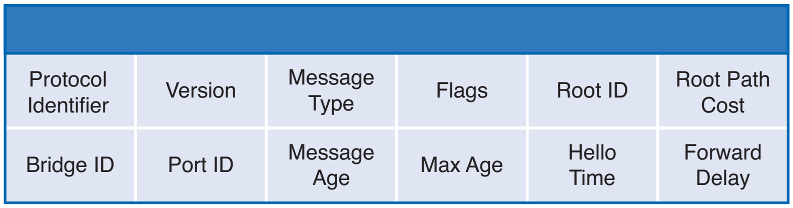

Protocol Identifier	Version	Message Type	Flags	Root ID	Root Path Cost
Bridge ID	Port ID	Message Age	Max Age	Hello Time	Forward Delay

Table 3-1 BPDU Fields

Field	Octets	Description
Protocol Identifier	1–2	Always 0.
Version	3	Always 0.
Message Type	4	Specifies the type of BPDU (Configuration or Topology Change Notification) that the frame contains.
Flags	5	Used to handle changes in active topology.
Root ID	6–13	Contains the bridge ID of the root bridge. This value is the same for all BPDUs in a bridged network after the network has converged.
Root Path Cost	14–17	Reflects the cumulative cost of all links leading to the root bridge.
Bridge ID	18–25	Contains the BID of the bridge that created the current BPDU.
Port ID	26–27	Contains a unique value for every port. For example, this field contains the value 0x8001 for port 1/1, whereas port 1/2 contains 0x8002.
Message Age	28–29	Records the time since the root bridge originally generated the information from which the current BPDU is derived.
Max Age	30–31	Shows the maximum time that a BPDU is saved. This influences the bridge table aging timer during the Topology Change Notification process.

Field	Octets	Description
Hello Time	32–33	Shows the time between periodic configuration BPDUs.
Forward Delay	34–35	Shows the time spent in the listening and learning states. This influences timers during the Topology Change Notification process.

As a switch powers on, each port cycles through a series of four states: blocking, listening, learning, and forwarding. A fifth state, disabled, indicates that the administrator has shut down the switch port. As the port cycles through these states, the LEDs on the switch change from flashing orange to steady green. It can take as long as 50 seconds for a port to cycle through all of these states and be ready to forward frames.

Blocking

When a switch powers on, it first goes into a ***blocking state*** to immediately prevent the formation of a loop. In this state, the bridge receives BPDUs but discards all data frames. In this state, the bridge does not learn new addresses and the port status light is steady amber. This phase can last up to 20 seconds before the port transitions to the listening state.

Listening

In the ***listening state***, the switch continues to listen to BPDUs but does not forward data frames or learn addresses. During this state, the switch determines which ports can forward frames without creating a loop. If enabling the port would create a loop, the switch returns the port to the blocking state. If no loop would be created, the switch transitions the port to the learning state. It takes 15 seconds to transition to the learning state, during which time the port indicator flashes amber.

Learning

In the ***learning state***, the switch receives and processes both BPDUs and data frames. It does not forward data frames while in this state but does learn MAC addresses from the data received. The port LED continues to flash amber during this state, which takes 15 seconds to complete before transitioning to the forwarding state.

Forwarding

In the ***forwarding state***, the switch continues to process both BPDUs and learn MAC addresses. It also now forwards data frames on the network. While in this state, the port status LED blinks green.

Disabled

If the administrator shuts down a port, it is considered to be ***disabled***. The port status indicator on a disabled port is off.

Access ports are ports that connect to an end host and carry the data for only a single VLAN. Because they do not normally connect to other switches, they do not create loops in a switched network. These ports always transition to forwarding if they have a host attached. ***Trunking ports*** can connect to other switches and normally carry data for multiple VLANs. These ports can potentially create a looped network and transition to either a forwarding or blocking state.

Interactive Activity 3-2: Spanning Tree (3.2.2)

In this activity, you associate the process with the correct spanning-tree state. Use file d3ia-322 on the CD-ROM that accompanies this book to perform this interactive activity.

Root Bridges

For STP to function, the switches in the network determine a switch that is the focal point in that network. STP uses this focal point, called a ***root bridge*** or root switch, to determine which ports to block and which ports to put into the forwarding state. The root bridge sends out BPDUs containing network topology information to all other switches. This information allows the network to reconfigure itself in the event of a failure.

Only one root bridge exists on each switched network, and it is elected based on the ***bridge-ID (BID)***. The bridge priority value plus the MAC address create the BID, as shown in Figure 3-8. Bridge priority has a default value of 32,768. If a switch has a MAC address of AA-11-BB-22-CC-33, the BID for that switch would be 32768: AA-11-BB-22-CC-33. The root bridge is based on the lowest BID value. Because switches typically use the same default priority value, the switch with the lowest MAC address becomes the root bridge.

Figure 3-8 Bridge ID

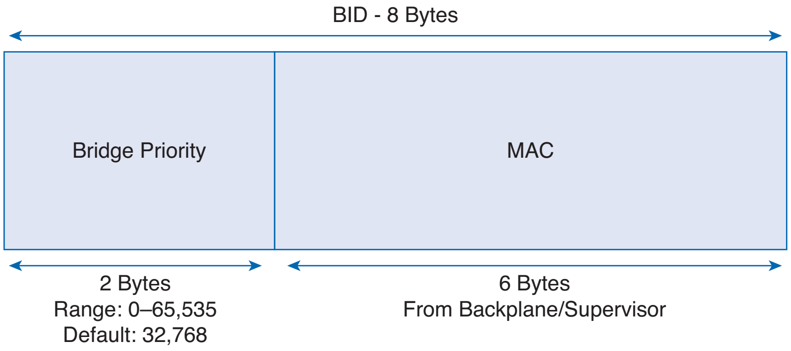

As each switch powers on, it assumes that it is the root bridge and sends out BPDUs containing its own BID. Consider the network shown in Figure 3-9. If S2 advertises a root ID that is a lower number than S1, S1 stops the advertisement of its root ID and accepts the root ID of S2. S2 is now the root bridge.

Figure 3-9 Port Designations

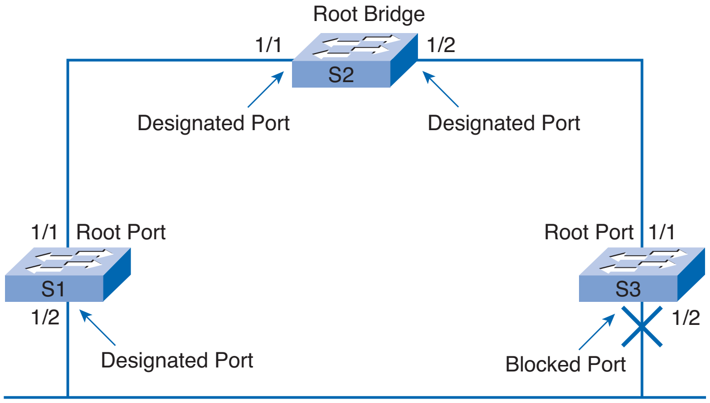

STP designates three types of ports, as shown in Figure 3-9:

- ***Root ports***: A port that provides the least-cost path back to the root bridge becomes the root port. Switches calculate the ***least-cost path*** using the bandwidth cost of each link required to reach the root bridge.
- ***Designated ports***: A designated port is a port that forwards traffic toward the root bridge but does not connect to the least-cost path. One designated port is elected per link (segment). The designated port is the port closest to the root bridge.
- ***Blocked ports***: A blocked port is a port that does not forward traffic.

Before configuring STP, the network technician plans and evaluates the network to select the best switch to become the root of the spanning tree. If the root switch selection is allowed to default to the one with the lowest MAC address, forwarding might not be optimal. A centrally located switch works best as the root bridge. A blocked port situated at the extreme edge of the network might cause traffic to take a longer route to get to the destination than if the switch is centrally located.

To specify the root bridge, the BID of the chosen switch is configured with the lowest priority value. The bridge priority command is used to configure the bridge priority. The range for the priority is from 0 to 65535, but values are in increments of 4096. The default value is 32768. The global **spanning-tree vlan VLAN-ID priority <0-61440>** command is used to set the priority of a switch.

Lab 3-2: Building a Switched Network with Redundant Links (3.2.3)

In this lab, you configure the BID on a switch to control which one becomes the root bridge. You also observe the spanning tree and traffic flow patterns as different switches are configured as root. Refer to the lab in Part II of this *Learning Guide*. You can perform this lab now or wait until the end of the chapter.

Spanning Tree in a Hierarchical Network

After establishing the root bridge, root ports, designated ports, and blocked ports, STP sends BPDUs throughout the switched network at 2-second intervals. STP continues to listen to these BPDUs to ensure that no links fail and no new loops appear. If a link failure occurs, STP recalculates as follows:

- Changing some blocked ports to forwarding ports
- Changing some forwarding ports to blocked ports
- Forming a new STP tree to maintain the loop-free integrity of the network

STP is not instantaneous. When a link goes down, STP detects the failure and recalculates the best paths across the network. This calculation and transition period takes about 30 to 50 seconds on each switch. During this recalculation, no user data passes through the recalculating ports. Figure 3-10 shows how the port designations are altered after a link failure.

Some user applications time out during the recalculation period, which can result in lost productivity and revenue. Frequent STP recalculations negatively impact uptime. A high-volume, enterprise server is normally connected to a switch port. If that port recalculates because of STP, the server is down for 50 seconds. It would be difficult to imagine the number of transactions lost during that time frame.

Figure 3-10 STP Recalculation

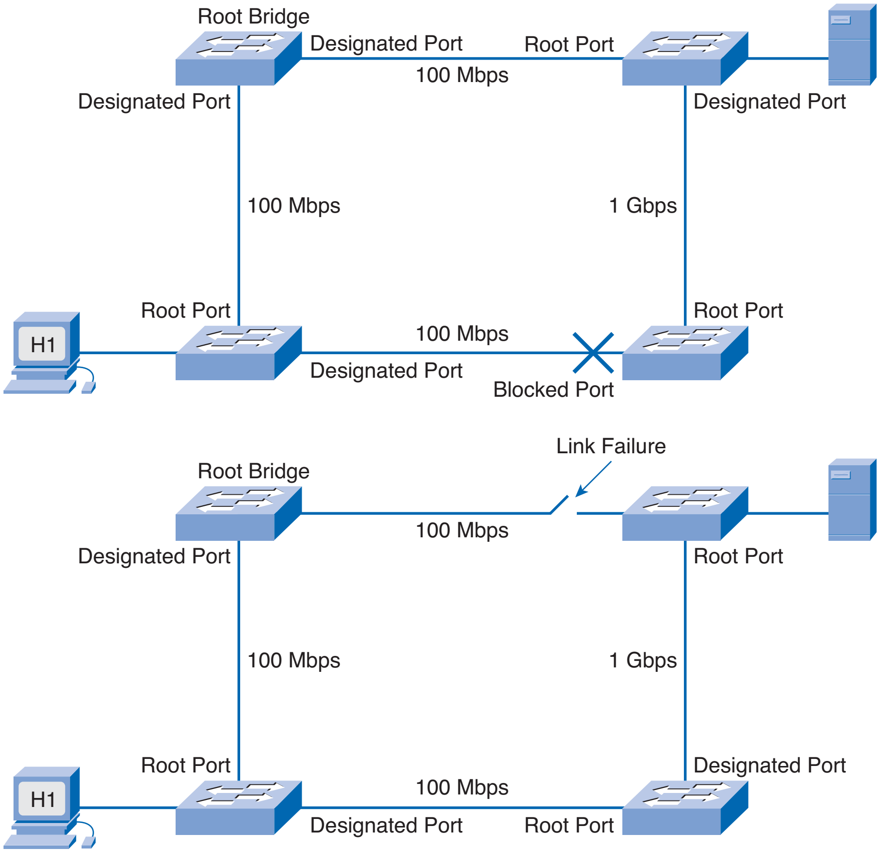

In a stable network, STP recalculations are infrequent. In an unstable network, it is important to check the switches for stability and configuration changes. One of the most common causes of frequent STP recalculations is a faulty power supply or power feed to a switch. A faulty power supply causes the device to reboot unexpectedly.

Several proprietary enhancements to STP exist to minimize the downtime incurred during an STP recalculation. These include PortFast, UplinkFast, and BackboneFast. These enhancements are Cisco proprietary; therefore, they cannot be used if the network includes switches from other vendors. In addition, all of these features require configuration. The following sections describe each enhancement and then conclude with **show** commands that provide information about STP on a network.

PortFast

STP ***PortFast*** causes an access port to enter the forwarding state immediately, bypassing the listening and learning states. Using PortFast on access ports that are connected to a single workstation or server allows those devices to connect to the network immediately, instead of waiting for STP to converge.

UplinkFast

STP ***UplinkFast*** accelerates the choice of a new root port when a link or switch fails or when STP reconfigures itself. The root port transitions to the forwarding state immediately without going through the listening and learning states, as it would do with normal STP procedures.

BackboneFast

BackboneFast provides fast convergence after a spanning-tree topology change occurs. It quickly restores backbone connectivity. BackboneFast is used at the distribution and core layers, where multiple switches connect.

STP Diagnostic show Commands

You can use a number of **show** commands to obtain information about the functionality of STP on a network. Sample output from some of the more useful commands is provided in Examples 3-1 to 3-6.

The **show spanning-tree** command displays root ID, bridge ID, and port states. A sample output is shown in Example 3-1.

Example 3-1 STP show spanning tree Sample Output

```
milfeet3548# show spanning-tree

VLAN0140
  Spanning tree enabled protocol ieee
  Root ID    Priority    32768
             Address     0004.4d3f.02c0
             Cost        8
             Port        49 (GigabitEthernet0/1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32908  (priority 32768 sys-id-ext 140)
             Address     000a.8a53.2200
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Gi0/1            Root FWD 4         128.49   P2p


VLAN0145
  Spanning tree enabled protocol ieee
  Root ID    Priority    32768
             Address     0004.4d3f.02c1
             Cost        8
             Port        49 (GigabitEthernet0/1)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32913  (priority 32768 sys-id-ext 145)
             Address     000a.8a53.2200
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

Interface        Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
Gi0/1            Root FWD 4         128.49   P2p
<output omitted>
```

The **show spanning-tree summary** command displays a summary of the port states. Example 3-2 shows sample output from this command.

Example 3-2 STP show spanning-tree summary Sample Output

```
milfeet3548# show spanning-tree summary

Switch is in pvst mode
Root bridge for: none
EtherChannel misconfig guard is enabled
Extended system ID           is enabled
Portfast Default             is disabled
PortFast BPDU Guard Default  is disabled
Portfast BPDU Filter Default is disabled
Loopguard Default            is disabled
UplinkFast                   is disabled
BackboneFast                 is disabled
Configured Pathcost method used is short

Name                   Blocking Listening Learning Forwarding STP Active
---------------------- -------- --------- -------- ---------- ----------
VLAN0140                     0         0        0          1          1
VLAN0145                     0         0        0          1          1
VLAN0151                     0         0        0          1          1
VLAN0152                     0         0        0          1          1
VLAN0153                     0         0        0          1          1
VLAN0155                     0         0        0          1          1
VLAN0157                     0         0        0          1          1
VLAN0231                     0         0        0          1          1
VLAN0262                     0         0        0          3          3
VLAN0420                     0         0        0          1          1
VLAN0430                     0         0        0          1          1
VLAN0900                     0         0        0          1          1
VLAN0901                     0         0        0          1          1
---------------------- -------- --------- -------- ---------- ----------
13 vlans                     0         0        0         15         15
```

Example 3-3 shows sample output from the **show spanning-tree root** command. Use this command to obtain information on the status and configuration of the root bridge.

Example 3-3 STP show spanning-tree root Sample Output

```
milfeet3548# show spanning-tree root

                                        Root Hello Max Fwd
Vlan                   Root ID          Cost  Time Age Dly  Root Port
---------------- -------------------- ------ ----- --- ---  ----------------
VLAN0140         32768 0004.4d3f.02c0         8    2   20  15  Gi0/1
VLAN0145         32768 0004.4d3f.02c1         8    2   20  15  Gi0/1
VLAN0151         32768 0004.4d3f.02c2         8    2   20  15  Gi0/1
VLAN0152         32768 0004.4d3f.02c3         8    2   20  15  Gi0/1
VLAN0153         32768 0004.4d3f.02d1         8    2   20  15  Gi0/1
VLAN0155         32768 0004.4d3f.02c4         8    2   20  15  Gi0/1
```

```
VLAN0157          32768 0004.4d3f.02c5         8    2   20  15  Gi0/1
VLAN0231          32768 0004.4d3f.02c6         8    2   20  15  Gi0/1
VLAN0262          32768 0004.4d3f.02c7         8    2   20  15  Gi0/1
VLAN0420          32768 0004.4d3f.02cc         8    2   20  15  Gi0/1
VLAN0430          32768 0004.4d3f.02cd         8    2   20  15  Gi0/1
VLAN0900          32768 0004.4d3f.02cf         8    2   20  15  Gi0/1
VLAN0901          32768 0004.4d3f.02d0         8    2   20  15  Gi0/1
```

To obtain detailed information on the spanning-tree ports, use the **show spanning-tree detail** command, as shown in Example 3-4.

Example 3-4 STP show spanning-tree detail Sample Output

```
milfeet3548# show spanning-tree detail

 VLAN0140 is executing the ieee compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, sysid 140, address 000a.8a53.2200
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 32768, address 0004.4d3f.02c0
  Root port is 49 (GigabitEthernet0/1), cost of root path is 8
  Topology change flag not set, detected flag not set
  Number of topology changes 15 last change occurred 4w5d ago
  Times:  hold 1, topology change 35, notification 2
          hello 2, max age 20, forward delay 15
  Timers: hello 0, topology change 0, notification 0, aging 300

 Port 49 (GigabitEthernet0/1) of VLAN0140 is forwarding
   Port path cost 4, Port priority 128, Port Identifier 128.49.
   Designated root has priority 32768, address 0004.4d3f.02c0
   Designated bridge has priority 32768, address 0015.c79e.0d8c
   Designated port id is 128.165, designated path cost 4
   Timers: message age 2, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default
   BPDU: sent 16, received 3186595


 VLAN0145 is executing the ieee compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, sysid 145, address 000a.8a53.2200
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 32768, address 0004.4d3f.02c1
  Root port is 49 (GigabitEthernet0/1), cost of root path is 8
  Topology change flag not set, detected flag not set
  Number of topology changes 12 last change occurred 4w5d ago
  Times:  hold 1, topology change 35, notification 2
          hello 2, max age 20, forward delay 15
  Timers: hello 0, topology change 0, notification 0, aging 300

 Port 49 (GigabitEthernet0/1) of VLAN0145 is forwarding
   Port path cost 4, Port priority 128, Port Identifier 128.49.
   Designated root has priority 32768, address 0004.4d3f.02c1
   Designated bridge has priority 32768, address 0015.c79e.0d91
   Designated port id is 128.165, designated path cost 4
   Timers: message age 2, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default
   BPDU: sent 12, received 3186349
<output omitted>
```

To view STP status and configuration information on a specific interface, use the **show spanning-tree interface** command, as illustrated in Example 3-5.

Example 3-5 STP show spanning-tree interface Sample Output

```
milfeet3548# show spanning-tree interface gigabitEthernet 0/1

Vlan             Role Sts Cost      Prio.Nbr Type
---------------- ---- --- --------- -------- --------------------------------
VLAN0140         Root FWD 4         128.49   P2p
VLAN0145         Root FWD 4         128.49   P2p
VLAN0151         Root FWD 4         128.49   P2p
VLAN0152         Root FWD 4         128.49   P2p
VLAN0153         Root FWD 4         128.49   P2p
VLAN0155         Root FWD 4         128.49   P2p
VLAN0157         Root FWD 4         128.49   P2p
VLAN0231         Root FWD 4         128.49   P2p
VLAN0262         Root FWD 4         128.49   P2p
VLAN0420         Root FWD 4         128.49   P2p
VLAN0430         Root FWD 4         128.49   P2p
VLAN0900         Root FWD 4         128.49   P2p
VLAN0901         Root FWD 4         128.49   P2p
```

The **show spanning-tree blockedports** command is used to view any ports that are currently blocked by STP. This is shown in Example 3-6.

Example 3-6 STP show spanning-tree blockedports Sample Output

```
milfeet3548# show spanning-tree blockedports

Name                 Blocked Interfaces List
-------------------- ------------------------------------


Number of blocked ports (segments) in the system : 0

milfeet3548#
```

Lab 3-3: Verifying STP with show Commands (3.2.4)

In this lab, you use various **show** commands to verify STP operation. Refer to the lab in Part II of this *Learning Guide*. You can perform this lab now or wait until the end of the chapter.

Rapid Spanning Tree Protocol (RSTP)

When the IEEE developed the original 802.1D Spanning Tree Protocol (STP), recovery time of 1 to 2 minutes was acceptable. Today, Layer 3 switching and advanced routing protocols provide a faster alternative path to the destination. The need to carry delay-sensitive traffic, such as voice and video, requires that switched networks converge quickly to keep up with the new technology. ***Rapid Spanning Tree Protocol (RSTP)***, defined in IEEE 802.1w, significantly speeds the recalculation of the spanning tree. Unlike PortFast, UplinkFast, and BackboneFast, RSTP is not proprietary.

RSTP requires a full-duplex, point-to-point connection between switches to achieve the highest reconfiguration speed. Reconfiguration of the spanning tree by RSTP occurs in less than 1 second, as compared to 50 seconds in STP. RSTP eliminates the requirements for features such as PortFast and UplinkFast. RSTP can revert to STP to provide services for legacy equipment.

To speed the recalculation process, RSTP reduces the number of port states to three: discarding, learning, and forwarding. The ***discarding state*** is similar to three of the original STP states: blocking, listening, and disabled. RSTP also introduces the concept of ***active topology***. All ports that are not discarding, or are blocked, are considered to be part of the active topology and will immediately transition to the forwarding state.

Configuring VLANs

Hosts and servers that are connected to Layer 2 switches are part of the same network segment. This arrangement poses two significant problems:

- Switches flood broadcasts out all ports, which consumes unnecessary bandwidth. As the number of devices connected to a switch increases, more broadcast traffic is generated and more bandwidth is wasted.
- Every device that is attached to a switch can forward and receive frames from every other device on that switch.

As a network design best practice, broadcast traffic is contained to the area of the network in which it is required. There are business reasons why certain hosts access each other while others do not. As an example, members of the accounting department might be the only users who need to access the accounting server. In a switched network, virtual local-area networks (VLAN) are created to contain broadcasts and group hosts together in communities of interest.

Virtual LAN

A ***virtual local-area network (VLAN)*** is a logical broadcast domain that can span multiple physical LAN segments. It allows an administrator to group stations by logical function, by project teams, or by applications, without regard to physical location of the users. This is shown in Figure 3-11.

Figure 3-11 VLANs

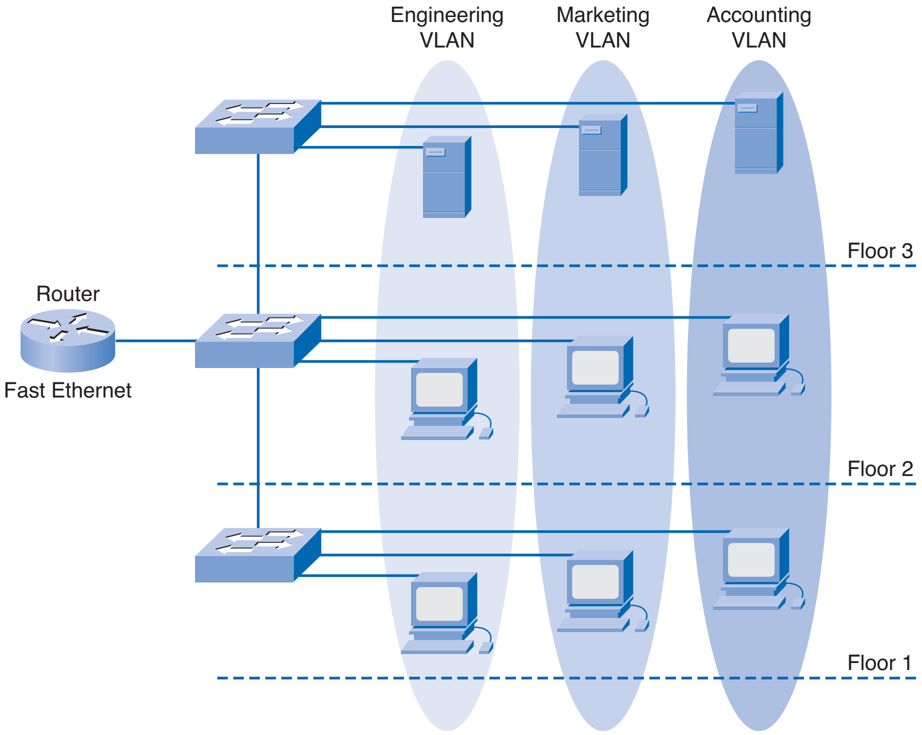

To understand the difference between a physical and a logical network, consider the following example. The students in a school are divided into two groups. In the first group, each student is given a red card, for identification. In the second group, each student is given a blue card. The principal announces that students with red cards can only speak to other students with red cards and that students with blue cards can only speak to other students with blue cards. The students are now logically separated into two virtual groups, which function as VLANs do in a network. Using this logical grouping, a broadcast goes out only to the red card group, even though both the red card group and the blue card group are physically located within the same school.

This example also shows another feature of VLANs. Broadcasts do not forward between VLANs; they are contained within the VLAN. Each VLAN functions as a separate LAN. A VLAN spans one or more switches, which allows host devices to behave as if they were on the same network segment. A VLAN has two major functions: It contains broadcasts and groups devices. Devices located on one VLAN are not visible to devices located on another VLAN. To move traffic between different VLANs requires the use of a Layer 3 device.

In a switched network, a device can be assigned to a VLAN based on its location, MAC address, IP address, or the applications that the device most frequently uses. Administrators assign membership in a VLAN either statically or dynamically.

Static VLANs

Static VLAN membership requires an administrator to manually assign each switch port to a specific VLAN. As an example, port Fa0/3 might be assigned to VLAN 20. Any device that plugs into port Fa0/3 automatically becomes a member of VLAN 20. This type of VLAN membership is the easiest to configure and is also the most popular; however, it requires the most administrative support for adds, moves, and changes. For example, moving a host from one VLAN to another requires either the switch port to be manually reconfigured to the new VLAN or the workstation cable to be plugged into a different switch port on the new VLAN. Membership in a specific VLAN is transparent to the users. Users working on a device plugged into a switch port have no knowledge that they are members of a VLAN.

Dynamic VLANs

Dynamic VLAN membership requires a ***VLAN management policy server (VMPS)***. The VMPS contains a database that maps MAC addresses to VLAN assignments. When a device plugs into a switch port, the VMPS searches the database for a match of the MAC address and temporarily assigns that port to the appropriate VLAN. Dynamic VLAN membership requires more organization and configuration but creates a structure with much more flexibility than static VLAN membership. In dynamic VLAN, moves, adds, and changes are automated and do not require intervention from the administrator.

Note

Not all Catalyst switches support the use of VMPSs.

Interactive Activity 3-3: Implementing VLANs (3.3.1)

In this activity, you decide whether VLANs can solve the stated problem. Use file d3ia-331 on the CD-ROM that accompanies this book to perform this interactive activity.

Configuring a Virtual LAN

Whether VLANs are created statically or dynamically, the maximum number of VLANs depends on the type of switch and the IOS. By default, VLAN1 is the ***management VLAN***. An administrator will use the IP address of the management VLAN to configure the switch remotely. When accessing the switch remotely, the network administrator can configure and maintain all VLAN configurations. Additionally, the management VLAN is used to exchange information, such as Cisco Discovery Protocol (CDP) traffic and VLAN Trunking Protocol (VTP) traffic, with other networking devices.

When a VLAN is created, it is assigned a number and a name. The VLAN number is any number from the range available on the switch, except for VLAN1. Some switches support approximately 1000 VLANs; others support more than 4000. Naming a VLAN is considered a network management best practice. To create a VLAN on a switch and give it a name, issue the following commands in global configuration mode:

```
Switch(config)# vlan vlan_number
Switch(config-vlan)# name vlan_name
Switch(config-vlan)# exit
```

After the VLAN is created, ports can be assigned either individually or as a range. By default, all ports are initially members of VLAN1. Use the following commands to assign individual ports to VLANs:

```
Switch(config)# interface fa#/#
Switch(config-if)# switchport access vlan vlan_number
Switch(config-if)# exit
```

Use the following commands to assign a range of ports to a VLAN:

```
Switch(config)# interface range fa#/start_of_range - end_of_range
Switch(config-if)# switchport access vlan vlan_number
Switch(config-if)# exit
```

Example 3-7 shows the creation of an accounting and production VLAN on a switch and the assignment of ports Fa0/3 and Fa0/5 to the accounting VLAN and ports Fa0/6 to Fa0/11 to the production VLAN.

Example 3-7 Creating VLANs and Assigning Ports

```
Switch(config)# vlan 101
Switch(config-vlan)# name accounting
Switch(config-vlan)# exit
Switch(config)# vlan 102
Switch(config-vlan)# name production
Switch(config-vlan)# exit
Switch(config)# interface fa0/3
Switch(config-if)# switchport access vlan 101
Switch(config)# interface fa0/5
Switch(config-if)# switchport access vlan10
Switch(config)# interface range fa0/6 - 11
Switch(config-if)# switchport access vlan 102
Switch(config-if)# exit
```

When working with VLANs, it is important to understand the key **show** commands that are available in the Cisco IOS. Sample output from some of the more important **show** commands used to verify, maintain, and troubleshoot VLANs is given in Examples 3-8 to 3-11.

The **show vlan** command displays a detailed list of all the VLAN numbers and names currently active on the switch, along with the ports associated with each one. This command also displays STP statistics if configured on a per-VLAN basis. Example 3-8 provides sample output from this command.

Example 3-8 show vlan Sample Output

```
Switch# show vlan

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/14, Fa0/15, Fa0/16
                                                Fa0/17, Fa0/18, Fa0/19, Fa0/20
                                                Fa0/21, Fa0/22, Fa0/23, Fa0/24
                                                Gi0/1, Gi0/2
27   accounting                       active    Fa0/13
28   engineering                      active    Fa0/6, Fa0/7, Fa0/8, Fa0/9
                                                Fa0/10, Fa0/11, Fa0/12
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
```

```
VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
1    enet  100001     1500  -      -      -        -    -        0      0
27   enet  100027     1500  -      -      -        -    -        0      0
28   enet  100028     1500  -      -      -        -    -        0      0
1002 fddi  101002     1500  -      -      -        -    -        0      0
1003 tr    101003     1500  -      -      -        -    -        0      0

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
1004 fdnet 101004     1500  -      -      -        ieee -        0      0
1005 trnet 101005     1500  -      -      -        ibm  -        0      0

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



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

Switch#
```

Sometimes detailed VLAN information is not required. In cases such as this, the **show vlan brief** command might be more appropriate. As shown in Example 3-9, this command displays a summarized list showing only the active VLANs and the ports associated with each one.

Example 3-9 `show vlan brief` Sample Output

```
Switch# show vlan brief

VLAN    Name                            Status     Ports
------- ----------------------------------------------- ------- ----------------
1      default                          active     Fa0/1, Fa0/4, Fa0/5,
                                                   Fa0/14, Fa0/15, Fa0/16, Fa0/17
                                                   Fa0/18, Fa0/19, Fa0/20, Fa0/21
                                                   Fa0/22, Fa0/23, Fa0/24
27    accounting                        active     Fa0/13
28    engineering                       active     Fa0/6, Fa0/7, Fa0/8, Fa0/9
                                                   Fa0/10, Fa0/11, Fa0/12
1002  fddi-default                      active
1003  token-ring-default                active
1004  fddinet-default                   active
1005  trnet-default
```

If information is required on only a single VLAN, the **show vlan id** or **show vlan name** command can be used to display information on the VLAN by VLAN ID number or VLAN name, respectively. Examples 3-10 and 3-11 provide sample output from these two commands.

Example 3-10 show vlan id Sample Output

```
Switch# show vlan id 28

VLAN  Name                             Status     Ports
------- --------------------------------------------- ------- ----------------
28    engineering                      active     Fa0/6, Fa0/7, Fa0/8, Fa0/9



VLAN  Type   SAID    MTU  Parent  RingNo  BridgeNo  Stp  BrdgMode  Tran1  Trans2
------- ------- ------- ------- ------- ------- -------------- ------- ----------
28   enet    100028  1500  -        -       -        -      -        0       0


Remote  SPAN VLANS
------------------
Disabled


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

Example 3-11 show vlan name Sample Output

```
Switch# show vlan name engineering

VLAN  Name                                  Status    Ports
------- ----------------------------------------- ------- -------------------------
28    engineering                           active    Fa0/6, Fa0/7, Fa0/8, Fa0/9



VLAN  Type  SAID   MTU  Parent  RingNo  BridgeNo  Stp  BrdgMode  Tran1  Trans2
------- ----- ------ ----- ------- ------- -------- ------ -------- ------- -----
28    enet  100028 1500   -      -       -        -     -         0      0


Remote  SPAN VLANS
------------------
Disabled


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

In an organization, employees are frequently added, removed, or moved to a different department or project. This constant movement requires VLAN maintenance, including removal or reassignment to different VLANs. The removal of VLANs and the reassignment of ports to different VLANs are two separate and distinct functions. When a port is disassociated from a specific VLAN, it returns to VLAN1. If a VLAN is deleted, the interfaces assigned to that VLAN will become inactive until assigned to another VLAN.

To delete a VLAN use the following command:

```
Switch(config)# no vlan vlan_number
```

To disassociate a port from a specific VLAN use these commands:

```
Switch(config)# interface fa#/#
Switch(config-if)# no switchport access vlan vlan_number
```

Lab 3-4: Configuring, Verifying, and Troubleshooting VLANs (3.3.2)

In this lab, you configure, verify, and troubleshoot VLANs on a switch. Refer to the lab in Part II of this *Learning Guide*. You can perform this lab now or wait until the end of the chapter.

Identifying VLANs

Devices connected to a VLAN only communicate with other devices in the same VLAN, regardless of whether those devices are on the same switch or different switches. A switch associates each port with a specific VLAN number. If VLANs are to span multiple switches or traffic must be routed between VLANs, trunking is used to carry the traffic from multiple VLANs over a single physical link. The trunking device inserts a tag into the original frame before sending the frame over the trunk link. The tag contains the ***VLAN ID (VID)***, which identifies the VLAN to which the traffic belongs. At the receiving end, the tag is removed and the frame is forwarded to the assigned VLAN. The addition of the VLAN ID number into the Ethernet frame is called ***frame tagging***.

The most commonly used frame-tagging standard is ***IEEE 802.1Q***. The 802.1Q standard, sometimes abbreviated to ***dot1q***, inserts a 4-byte tag field into the Ethernet frame. This tag sits between the source address and the type/length field. Untagged Ethernet frames have a minimum size of 64 bytes and a maximum size of 1518 bytes. This tag field increases the minimum Ethernet frame from 64 to 68 bytes. The maximum size increases from 1518 to 1522 bytes. The switch recalculates the FCS because the number of bits in the frame has been modified. The FCS field provides error checking to ensure the integrity of all the bits within the frame.

An 802.1Q tagged Ethernet frame is shown in Figure 3-12. Table 3-2 describes the fields in the 801.1Q tag.

Figure 3-12 802.1Q Tagged Ethernet Frame

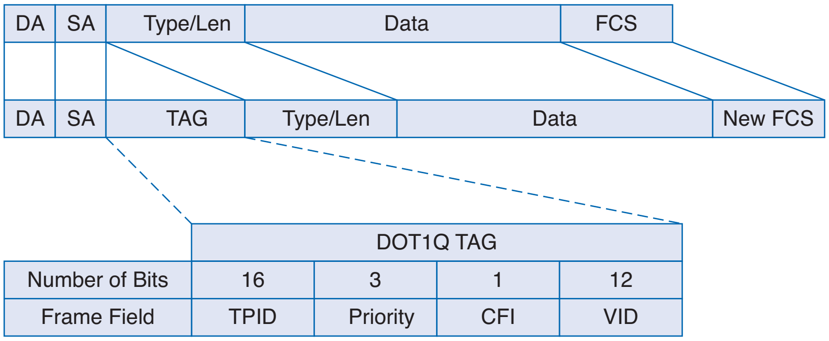

	DOT1Q TAG			
Number of Bits	16	3	1	12
Frame Field	TPID	Priority	CFI	VID

Table 3-2 802.1Q Tag Fields

Field	Description
Tag Protocol Identifier (TPID)	Set to a value of 0x8100 to identify the frame as an IEEE 802.1Q tagged frame.
Priority	Known as user priority. This 3-bit field refers to the IEEE 802.1Q priority. The field indicates the frame priority level used for the prioritization of traffic. The field can represent 8 levels (0 through 7).
Canonical Format Identifier (CFI)	A 1-bit indicator used for compatibility between Ethernet and Token Ring networks. Always set to 0 for Ethernet switches.
VLAN Identifier (VID)	Uniquely identifies the VLAN to which the frame belongs. The field has a value between 0 and 4095.

If an 802.1Q-compliant port is connected to another 802.1Q-compliant port, the VLAN tagging information passes between them. If a non-802.1Q-enabled device or an access port receives an 802.1Q frame, the tag data is ignored, and the packet is switched at Layer 2 as a standard Ethernet frame. This allows the placement of Layer 2 intermediate devices, such as other switches or bridges, along the 802.1Q trunk path. To process an 802.1Q tagged frame, a device must allow an MTU of 1522 or higher. Some older devices and network cards that are not 802.1Q compatible view a tagged Ethernet frame as too large. These devices drop the frame and log it as an error, called a baby giant.

Interactive Activity 3-4: Frame Delivery (3.3.3)

In this activity, you decide whether a frame can be delivered based on the VLAN and port configurations. Use file d3ia-333 on the CD-ROM that accompanies this book to perform this interactive activity.

Trunking and Inter-VLAN Routing

A VLAN limits the size of broadcast domains. This has the effect that there is less traffic on an area of the network, which improves network performance. It also provides a level of security by containing traffic within a certain area of the network. A Layer 3 device is required to move traffic between VLANs. This Layer 3 device can filter traffic as it passes between VLANs, thus allowing the network administrator to have complete control over which traffic is allowed to move between VLANs. To take full advantage of the benefits of VLANs, they are extended across multiple switches in the enterprise network.

Trunk Ports

Switch ports can be configured for two different roles. A port is classified as either an access port or a trunk port, as shown in Figure 3-13.

Figure 3-13 Access and Trunk Ports

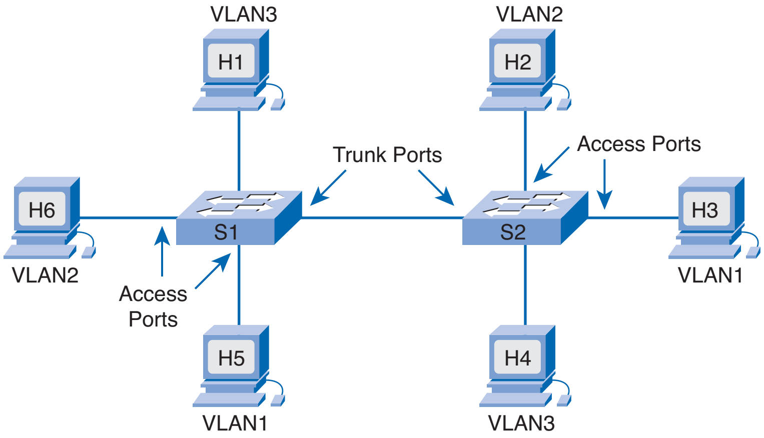

Access Port

An access port belongs to only one VLAN. Typically, single devices such as PCs or servers connect to this type of port. If a hub connects multiple PCs to the single access port, each device connected to the hub is a member of the same VLAN.

Trunk Port

A trunk port is a point-to-point link between the switch and another networking device. Trunks carry the traffic of multiple VLANs over a single link and allow VLANs to reach across an entire network. Trunk ports are necessary to carry the traffic from multiple VLANs between devices when connecting either two switches together, a switch to a router, or a host NIC that supports 802.1Q trunking.

Without trunk ports, each VLAN requires a separate connection between switches. For example, an enterprise with 100 VLANs requires 100 connecting links. This type of arrangement does not scale well and is very expensive. Trunk links provide a solution to this problem by transporting traffic from multiple VLANs on the same physical link. When multiple VLANs travel on the same link, they need VLAN identification. A trunk port supports frame tagging. Frame tagging adds VLAN information to the frame.

IEEE 802.1Q is the standardized and approved method of frame tagging. Cisco developed a proprietary frame-tagging protocol called ***Inter-Switch Link (ISL)*****.** Higher-end switches, such as the Catalyst 6500 series, still support both tagging protocols; however, most LAN switches, such as the 2960, support only 802.1Q.

Switch ports are access ports by default. To configure a switch port as a trunk port, use the following commands:

```
Switch(config)# interface fa(controller # / port #)
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport trunk encapsulation {dot1q ¦ isl ¦ negotiate}
```

Switches that support both 802.1Q and ISL require the last configuration statement. The 2960 does not require that statement because it only supports 802.1Q. The negotiate parameter is the default mode on many Cisco switches. This parameter automatically detects the encapsulation type of the neighbor switch.

Newer switches can detect the type of link configured at the other end. Based on the attached device, the link configures itself as either a trunk port or an access port. To turn on this feature use the following command:

```
Switch(config-if)# switchport mode dynamic {desirable ¦ auto}
```

In desirable mode, the port becomes a trunk port if the other end is set to either trunk, desirable, or auto. In auto mode, the port becomes a trunk port if the other end is set to either trunk or desirable. To return a trunk port to an access port, issue either of the following commands:

```
Switch(config-if)# no switchport mode trunk
Switch(config-if)# switchport mode access
```

Lab 3-5: Creating VLANs and Assigning Ports (3.4.1)

In this lab, you create VLANs on a switch and assign ports. Refer to the lab in Part II of this *Learning Guide*. You can perform this lab now or wait until the end of the chapter.

Extending VLANs Across Switches

Trunking enables VLANs to forward traffic between switches using only a single port. A trunk link configured with 802.1Q on both ends allows traffic that has a 4-byte tag field added to the frame. This frame tag contains the VLAN ID. When a switch receives a tagged frame on a trunk port, it removes the tag before sending it out an access port. This is shown in Figure 3-14. The switch forwards the frame only if the access port is a member of the same VLAN as the tagged frame.

Figure 3-14 VLAN Tags

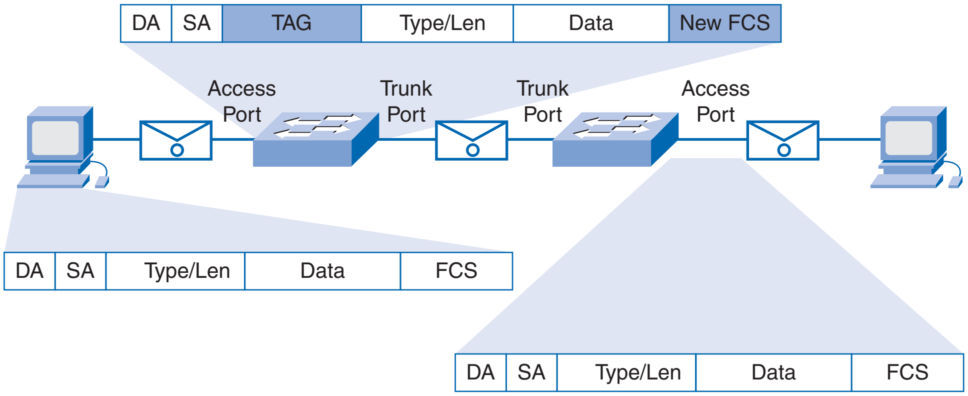
DA | SA | TAG | Type/Len | Data | New FCS
Access Port
Trunk Port
Trunk Port
Access Port
DA | SA | Type/Len | Data | FCS
DA | SA | Type/Len | Data | FCS

Some traffic however, needs to cross the 802.1Q configured link without a VLAN ID. Traffic with no VLAN ID is called untagged. Examples of untagged traffic are Cisco Discovery Protocol (CDP), VTP, and certain types of voice traffic. Untagged traffic minimizes the delays associated with inspection of the VLAN ID tag. To accommodate untagged traffic, a special VLAN called a ***native VLAN*** is available. Untagged frames received on the 802.1Q trunk port will become members of the native VLAN. On Cisco Catalyst switches, VLAN 1 is the native VLAN by default.

Any VLAN can be configured as the native VLAN. Ensure that the native VLAN for an 802.1Q trunk is the same on both ends of the trunk line. If they are different, switching loops might result. On an 802.1Q trunk, use the following command to assign the native VLAN ID on a physical interface:

```
Switch(config-if)# switchport trunk native vlan vlan-id
```

Lab 3-6: Configuring Trunk Ports to Connect Switches (3.4.2)

In this lab, you configure trunk ports to connect two switches and verify connectivity across the trunk link. Refer to the lab in Part II of this *Learning Guide*. You can perform this lab now or wait until the end of the chapter.

Inter-VLAN Switching

Although VLANs extend to span multiple switches, only members of the same VLAN can communicate without the assistance of a Layer 3 device. This arrangement enables the network administrator to strictly control the type of traffic that flows from one VLAN to another. One method of accomplishing the inter-VLAN routing requires a separate interface connection to the Layer 3 device for each VLAN, as shown in Figure 3-15.

Figure 3-15 Inter-VLAN Routing

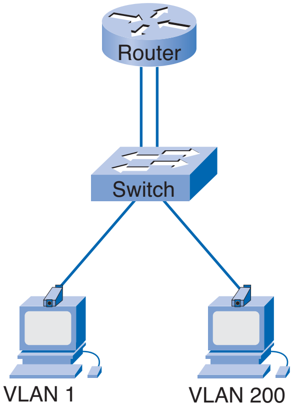

Another method for providing connectivity between different VLANs requires a feature called ***subinterfaces***, as shown in Figure 3-16. Subinterfaces logically divide one physical interface into multiple logical pathways. Configure one pathway or subinterface for each VLAN. Supporting inter-VLAN communication using subinterfaces requires configuration on both the switch and the router.

The switch interface connecting to the router must be configured as an 802.1Q trunk link. The router interfaces must be a minimum of 100-Mbps Fast Ethernet and support 802.1Q encapsulation. A separate subinterface must be configured for each VLAN. These subinterfaces allow each VLAN to have its own logical pathway and default gateway into the router.

The host from the sending VLAN forwards traffic to the router using the default gateway. The subinterface for the VLAN specifies the default gateway for all hosts in that VLAN. The router locates the destination IP address and does a routing table lookup. If the destination VLAN is on the same switch as the source VLAN, the router forwards the traffic back down to the source switch using the subinterface parameters of the destination VLAN ID. This type of configuration is often referred to as a ***router-on-a-stick***.

Figure 3-16 Inter-VLAN Routing Using Subinterfaces

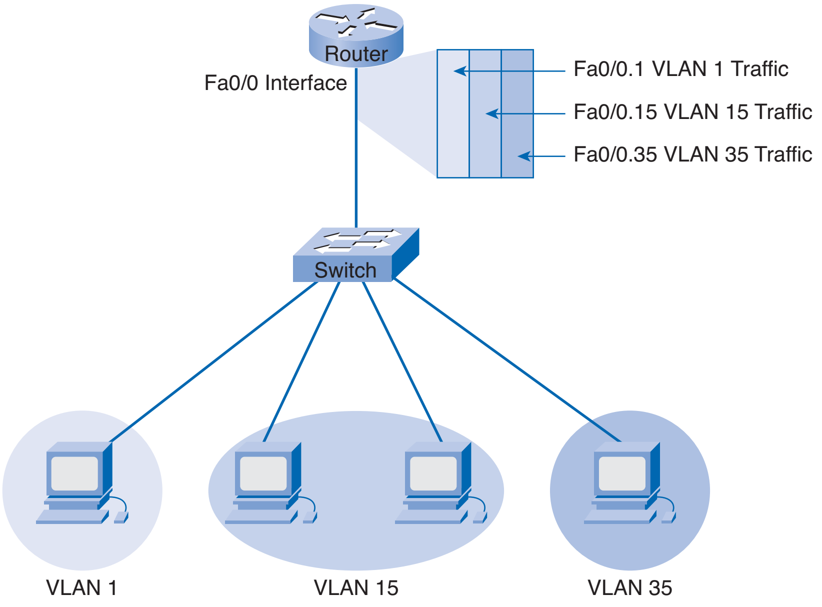

If the exit interface of the router is 802.1Q compatible, the frame retains its 4-byte VLAN tag. If the outbound interface is not 802.1Q compatible, the router strips the tag from the frame and returns the frame to its original Ethernet format.

Consider an example where port Fa0/1 on a router is connected to Fa0/2 on a switch. To configure inter-VLAN routing, use the following steps:

Step 1. Configure a trunk port on the switch.

```
Switch(config)# interface fa0/2
Switch(config-if)# switchport mode trunk
```

Step 2. On the router, configure a Fast Ethernet interface with no IP address or subnet mask.

```
Router(config)# interface fa0/1
Router(config-if)# no ip address
Router(config-if)# no shutdown
```

Step 3. On the router, configure one subinterface with an IP address and subnet mask for each VLAN. Each subinterface has an 802.1Q encapsulation. The number following the dot1q statement is the VLAN ID. Each subinterface must be assigned an IP address because it will act as the default gateway for that subnet.

```
Router(config)# interface fa0/1.10
Router(config-subif)# encapsulation dot1q 10
Router(config-subif)# ip address 192.168.10.1 255.255.255.0
```

Step 4. Use the following commands to verify the inter-VLAN routing configuration and functionality:

```
Switch# show trunk


Router# show ip interfaces

Router# show ip interfaces brief

Router# show ip route
```

Note

If inter-VLAN routing fails to function as expected, make certain that no IP address is assigned to the physical Ethernet interface and that the physical interface has been activated.

Lab 3-7: Part A: Configuring Inter-VLAN Routing (3.4.3)

In this lab, you configure inter-VLAN routing using separate interfaces for each VLAN. Refer to the lab in Part II of this *Learning Guide*. You can perform this lab now or wait until the end of the chapter.

Lab 3-7: Part B: Configuring Inter-VLAN Routing (3.4.3)

In this lab, you configure inter-VLAN routing using a router on-a-stick configuration. Refer to the lab in Part II of this *Learning Guide*. You can perform this lab now or wait until the end of the chapter.

Maintaining VLANs on an Enterprise LAN

As networks grow in size and complexity, centralized management of the VLAN structure becomes crucial. If there is no automated way to manage an enterprise network with hundreds of VLANs, manual configuration of each VLAN on each switch is necessary. Any change to the VLAN structure requires further manual configuration. One incorrectly keyed number causes inconsistencies in connectivity throughout the entire network. To resolve this issue, Cisco created VTP to automate many of the VLAN configuration functions.

VLAN Trunking Protocol (VTP)

VLAN Trunking Protocol (VTP) is a Layer 2 messaging protocol that provides a method for the distribution and management of the VLAN database from a centralized server in a network segment. Routers do not forward VTP updates. VTP ensures that VLAN configuration is consistently maintained across the network and reduces the task of VLAN management and monitoring.

VTP is a client/server messaging protocol that adds, deletes, and renames VLANs in a single VTP domain. All switches under a common administration are part of a domain. Each domain has a unique name. VTP switches only share VTP messages with other switches in the same domain. Two different versions of VTP exist: Version 1 and Version 2. Version 1 is the default and it is not compatible with Version 2. All switches must be configured with the same version.

With VTP, each switch advertises messages on its trunk ports. Messages include the management domain, configuration revision number, known VLANs, and parameters for each VLAN. These advertisement frames are sent to a multicast address so that all neighbor devices receive the frames.

VTP Modes

VTP has three modes: server, client, and transparent. The following sections describe each.

VTP Server Mode

By default, all Cisco switches are in ***VTP server*** mode. In this mode, the administrator can create, modify, and delete VLANs and VLAN configuration parameters for the entire domain. A VTP server saves VLAN configuration information in the switch NVRAM and sends out VTP messages on all trunk ports. It is good practice to have at least two switches configured as servers on a network to provide backup and redundancy.

VTP Client Mode

A switch in ***VTP client*** mode does not create, modify, or delete VLAN information for the VTP domain. It accepts VTP messages from the VTP server and modifies its own database with this information. A VTP client sends VTP messages out all trunk ports.

VTP Transparent Mode

A switch in ***VTP transparent*** mode ignores information in the VTP messages. It does not modify its database with information received from the VTP server but does forward VTP advertisements. A switch in VTP transparent mode will not send out an update that indicates a change in its own VLAN database. Therefore, all VLAN created on a switch in this mode remain local to the switch.

Challenge Lab 3-8: VTP Modes

In this lab, you configure VTP and propagate VLAN information through a network. Refer to the lab in Part II of this *Learning Guide*. You can perform this lab now or wait until the end of the chapter.

VTP Revision Numbers

Each VTP switch saves a VLAN database in NVRAM that contains a revision number. If a VTP receives an update message that has a higher revision number than the one stored in the database, the switch updates its VLAN database with this new information. The VTP configuration revision number begins at 0. As changes occur, the configuration revision number increases by 1. The revision number continues to increment until it reaches 2,147,483,648. When it reaches that point, the counter resets to 0.

There are two ways to reset the VTP revision number to 0. The first is to set the switch you are inserting into the network to VTP transparent mode and then set it back to either a VTP client or server. A second method is to change the VTP domain name to something else and then change it back again.

A problem situation can occur related to the revision number if someone inserts a switch with a higher revision number into the network without first resetting the VTP revision number. Because a switch is a server by default, this results in new, but incorrect, information overwriting the legitimate VLAN information on all the other switches. Another way to protect against this critical situation is to configure a VTP password to validate the switch. When adding a new switch to an existing network, always reset the revision number. In addition, when adding a switch and when a server switch already exists, make sure that the new switch is configured in client or transparent mode.

VTP Message Types

VTP messages come in three varieties: summary advertisements, subset advertisements, and advertisement requests.

Summary Advertisements

Catalyst switches issue ***summary advertisements*** every 5 minutes or whenever a change to the VLAN database occurs. Summary advertisements contain the current VTP domain name and the configuration revision number. If VLANs are added, deleted, or changed, the server increments the configuration revision number and issues a summary advertisement.

When a switch receives a summary advertisement packet, it compares the VTP domain name to its own VTP domain name. If the domain name is the same, the switch compares the configuration revision number to its own number. If it is lower or equal, the switch ignores the packet. If the revision number is higher, an advertisement request is sent.

Subset Advertisements

A ***subset advertisement*** follows the summary advertisement. A subset advertisement contains a list of VLAN information. The subset advertisement contains the new VLAN information based on the summary advertisement. If several VLANs exist, they require more than one subset advertisement.

Advertisement Requests

Catalyst switches use an ***advertisement request*** to ask for VLAN information. Advertisement requests are required if the switch has been reset or if the VTP domain name has been changed. The switch receives a VTP summary advertisement with a higher configuration revision number than its own.

Interactive Activity 3-5: VTP Mode Characteristics (3.5.1)

In this activity, you select the characteristics of VTP client, server, and transparent modes. Use file d3ia-351 on the CD-ROM that accompanies this book to perform this interactive activity.

Configuring VTP

Switches are servers by default. If a switch in server mode issues an update with a higher revision number than the number currently in place, all switches will modify their databases to match the new switch.

When adding a new switch to an existing VTP domain, use the following steps:

Step 1. Configure VTP off-line as follows:

```
Switch(config)# vtp domain domain_name
Switch(config)# vtp mode {server ¦ client ¦ transparent}
Switch(config)# vtp password password
Switch(config)# end
Switch# copy running-config startup-config
```

Step 2. Verify the VTP configuration using the following commands. Ensure that the revision number is not higher than the network the switch is joining.

```
Switch# show vtp status

VTP Version                     : 2
Configuration Revision          : 6
Maximum VLANs supported locally : 64
Number of existing VLANs        : 9
VTP Operating Mode              : Server
VTP Domain Name                 : headoffice
VTP Pruning Mode                : Disabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x24 0xF1 0xB2 0xF8 0xC9 0x0E 0x9F 0x96
Configuration last modified by 0.0.0.0 at 3-1-93 00:10:17
Local updater ID is 0.0.0.0 (no valid interface found)

Switch# show vlan

VLAN Name                             Status    Ports
---- -------------------------------- --------- ----------------------------
---
1    default                          active    Fa0/1, Fa0/13, Fa0/14,
Fa0/15
                                                Fa0/16, Fa0/17, Fa0/18,
Fa0/19
                                                Fa0/20, Fa0/21, Fa0/22,
Fa0/23
                                                Fa0/24
10   inventory                        active    Fa0/6
20   marketing                        active    Fa0/4, Fa0/5
30   sales                            active    Fa0/7, Fa0/8, Fa0/9, Fa0/10
                                                Fa0/11, Fa0/12
50   administrators                   active
1002 fddi-default                     act/unsup
1003 token-ring-default               act/unsup
1004 fddinet-default                  act/unsup
1005 trnet-default                    act/unsup
```

```
VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1
Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ----
--
1    enet  100001     1500  -      -      -        -    -        0      0
10   enet  100010     1500  -      -      -        -    -        0      0
20   enet  100020     1500  -      -      -        -    -        0      0
30   enet  100030     1500  -      -      -        -    -        0      0

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1
Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ----
--
50   enet  100050     1500  -      -      -        -    -        0      0
1002 fddi  101002     1500  -      -      -        -    -        0      0
1003 tr    101003     1500  -      -      -        -    -        0      0
1004 fdnet 101004     1500  -      -      -        ieee -        0      0
1005 trnet 101005     1500  -      -      -        ibm  -        0      0

Switch# show vtp password

VTP Password: itsasecret

Switch# show vtp counters

VTP statistics:
Summary advertisements received    : 4
Subset advertisements received     : 4
Request advertisements received    : 2
Summary advertisements transmitted : 6
Subset advertisements transmitted  : 6
Request advertisements transmitted : 0Number of config revision errors   : 0
Number of config digest errors     : 0
Number of V1 summary errors        : 0

VTP pruning statistics:

Trunk            Join Transmitted Join Received    Summary advts received
from
```

```
                                                   non-pruning-capable
device
---------------- ---------------- ---------------- -------------------------
--
Fa0/2               0                1                0
Fa0/3               0                1                0
```

Step 3. Reboot the switch:

```
Switch# reload
```

Packet Tracer
☐ Activity

Building and Testing a VTP Domain (3.5.2.2)

In this activity, you build and test a VTP domain. Use file d3-3522.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

Packet Tracer
☐ Activity

Adding a Switch to a VTP Domain (3.5.2.3)

In this activity, you add a new switch to an existing VTP domain. Use file d3-3523.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

VLAN Support for IP Telephony and Wireless

The main purpose of VLANs is to separate traffic into logical groups. Traffic from one VLAN will not impact traffic from another VLAN. A common network design is to separate voice and wireless traffic from the rest of the traffic. This is shown in Figure 3-17.

A VLAN environment is ideal for traffic that is sensitive to time delays, such as voice. Voice traffic must be given priority over normal data traffic to avoid jerky or jittery conversations. Providing a dedicated VLAN for voice traffic prevents voice traffic from having to compete with data for available bandwidth.

An IP phone usually has two ports, one for voice and one for data. Packets traveling to and from the PC and the IP phone share the same physical link to the switch and the same switch port. To segment the voice traffic, enable a separate voice VLAN on the switch.

Wireless is another type of traffic that benefits from VLANs. Wireless is, by nature, very insecure and prone to attacks by hackers. VLANs created for wireless traffic isolate some of the problems that can occur. A compromise to the integrity of the wireless VLAN has no effect on any other VLAN within the organization. Most wireless deployments place the user in a VLAN on the outside of the firewall for added security. Users have to authenticate to gain entry into the internal network from the wireless network.

In addition, many organizations provide guest access to their wireless network. Guest accounts provide anyone, within a limited range, temporary wireless services such as web access, e-mail, ftp, and SSH. Guest accounts are either included in the wireless VLAN or reside in a VLAN of their own.

Figure 3-17 Separating Voice and Wireless Traffic

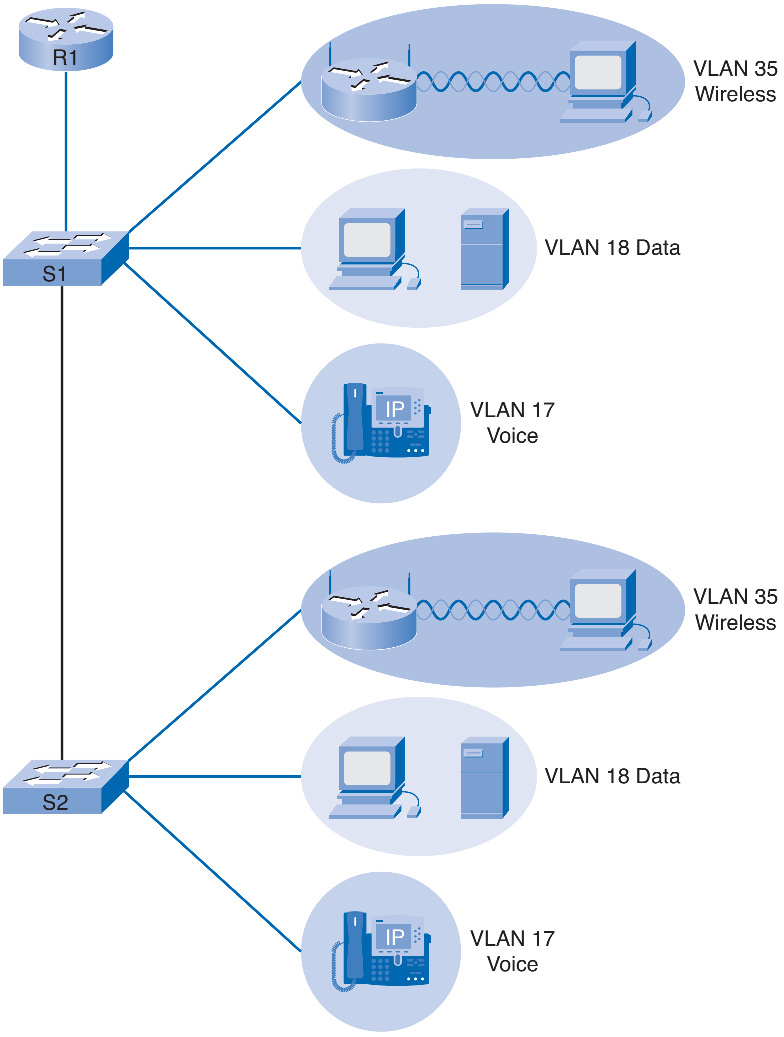

Packet Tracer
□ Activity

Configuring Wireless and Voice VLANs (3.5.3)

In this activity, you create separate VLANs for voice and wireless traffic. Use file d3-353.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

VLAN Best Practices

When carefully planned and designed, VLANs provide security, conserve bandwidth, and localize traffic on an enterprise network. All of these features combine to improve network performance. VLANs, however, are not the answer to every problem. If VLANs are not correctly implemented, they can overly complicate a network, resulting in inconsistent connectivity and slow network performance. VLANs isolate certain types of traffic for reasons of security. Moving traffic between VLANs requires a Layer 3 device, which increases the cost of implementation and introduces an increased level of latency into the network.

Table 3-3 presents some recommended best practices for configuring VLANs in an enterprise network.

Table 3-3 VLAN Best Practices

Best Practice	Description
Server placement	Ensure that all servers required by a particular group are members of the same VLAN.
Unused ports	Disable unused ports. Put unused ports in an unused VLAN. Stop unauthorized access by not granting connectivity or by placing a device into an unused VLAN.
Management VLAN	By default, the management VLAN and the native VLAN are VLAN1. Do not use VLAN1 for in-band management traffic. Select a different, dedicated VLAN to keep management traffic separate from user data and protocol traffic.
VLAN Trunking Protocol	Standardize the VLAN configuration across the enterprise. Provide easy VLAN management and maintenance. Reduce the time required for VLAN administration and maintenance.
VTP domains	Prevent the risk of an administrator error propagating to the entire network. Configure the VTP domains carefully and consistently. Turn off VTP when not required.
VTP revision number	Ensure that any new switch added to the network has a revision number of 0.

Packet Tracer
□ **Activity**

Planning and Building an Enterprise Network (3.5.4)

In this activity, you plan and build a switched network to meet client specifications. Use file d3-354.pka on the CD-ROM that accompanies this book to perform this activity using Packet Tracer.

Summary

Modern enterprise networks rely on switches for their basic functionality. Switches use microsegmentation to limit the size of the broadcast domain to a single switch port. Traditionally, switches used either store-and-forward or cut-through techniques. Improvements in technology allow some switches to automatically adapt their switching technique to match network conditions.

Redundancy is incorporated into network design to minimize downtime in the network. Spanning Tree Protocol is used to prevent the formation of switching loops in these redundant networks. Spanning Tree elects a root bridge based on the lowest bridge ID. Spanning Tree can take 50 seconds to reach the state where it is forwarding packets through the network. Rapid Spanning Tree Protocol has been developed to greatly shorten this convergence time.

A VLAN is a collection of hosts that are on the same local-area network, even though they might be physically separated from each other. A 4-byte 802.1Q header is inserted into a standard Ethernet frame to identify the VLAN to which the frame belongs. By default, VLAN 1 is the management VLAN.

An access port connects a device to a switch and is a member of a single VLAN. A trunk port carries traffic from multiple VLANs and normally connects a switch to another switch or a router. The native VLAN is the VLAN that uses untagged traffic. A Layer 3 device is required to move traffic between VLANs. Routers are usually configured with subinterfaces to handle multiple VLANs on a single physical link.

VTP provides a mechanism for the centralized control, distribution, and maintenance of the VLAN database. Switches can either be server, client, or transparent. VTP updates contain a revision number. A higher revision number update will overwrite information from a lower revision number.

VLANs are suitable for delay-sensitive traffic such as voice. They are also beneficial when it is important to keep different types of traffic off different areas of the network to provide a level of security and bandwidth management. When implementing VLANs, a number of best practices should be followed, including using a VTP password, consistent domain name, and revision control.

Activities and Labs

This summary outlines the activities and labs you can perform to help reinforce important concepts described in this chapter. You can find the activity and Packet Tracer files on the CD-ROM accompanying this book. The complete hands-on labs appear in Part II.

Interactive Activities on the CD-ROM:

Interactive Activity 3-1: Switch Frame Forwarding (3.1.1)

Interactive Activity 3-2: Spanning Tree (3.2.2)

Interactive Activity 3-3: Implementing VLANs (3.3.1)

Interactive Activity 3-4: Frame Delivery (3.3.3)

Interactive Activity 3-5: VTP Mode Characteristics (3.5.1)

Packet Tracer
□ **Activity**

Packet Tracer Activities on the CD-ROM:

Disabling Redundant Links to Avoid Switching Loops (3.2.1)

Building and Testing a VTP Domain (3.5.2.2)

Adding a Switch to a VTP Domain (3.5.2.3)

Configuring Wireless and Voice VLANs (3.5.3)

Planning and Building an Enterprise Network (3.5.4)

Hands-on Labs in Part II of this book:

Lab 3-1: Applying Basic Switch Security (3.1.4)

Lab 3-2: Building a Switched Network with Redundant Links (3.2.3)

Lab 3-3: Verifying STP with **show** Commands (3.2.4)

Lab 3-4: Configuring, Verifying, and Troubleshooting VLANs (3.3.2)

Lab 3-5: Creating VLANs and Assigning Ports (3.4.1)

Lab 3-6: Configuring a Trunk Port to Connect Switches (3.4.2)

Lab 3-7: Part A: Configuring Inter-VLAN Routing (3.4.3)

Lab 3-7: Part B: Configuring Inter-VLAN Routing (3.4.3)

Challenge Lab 3-8: VTP Modes

Check Your Understanding

Complete all the review questions listed here to check your understanding of the topics and concepts in this chapter. Appendix A, "Check Your Understanding and Challenge Questions Answer Key," lists the answers.

1. What does a switch do with a frame if the destination address is not in the MAC address database?

   A. It floods the frame out all ports except the one on which it was received.
   B. It forwards the frame out each port until it receives a receipt acknowledgment from the destination.
   C. It discards the frame.
   D. It buffers the frame until the destination address is learned.

2. What type of switching technology is most often used by modern LAN switches?

   A. Store-and-forward
   B. Fast-forward
   C. Fragment-free
   D. Adaptive

3. What type of switching is most appropriate for connecting multiple hosts to a server farm?

   A. Symmetric
   B. Asymmetric
   C. Cut-through
   D. Store-and-forward

4. Which type of switching would not switch runts? (Select all that apply.)

   A. Store-and-forward
   B. Fragment-free
   C. Fast-forward
   D. Adaptive

5. Which are security measures that should be implemented in a switched network? (Select all that apply.)

   A. Disable Telnet and HTTP access
   B. Disable unused ports
   C. Enable port security
   D. Restrict access to the physical switch

6. Which can be the result of redundancy in a switched network? (Select all that apply.)

   A. MAC database instability
   B. Broadcast storms
   C. Multiple frame transmission
   D. Increased availability

7. Which protocol is used to create a loop-free environment in a switched network that has redundant links?

   A BPDU
   B. STP
   C. VLAN
   D. VTP

8. Which open standard was developed to reduce the time required for a switched network to reach convergence?

   A. BackboneFast
   B. PortFast
   C. UplinkFast
   D. RSTP

9. When a new host is connected to a switch with STP enabled, what mode is the port placed into?

   A. Blocking
   B. Listening
   C. Learning
   D. Forwarding
   E. Disabled

10. In which STP mode does the switch listen for BPDUs? (Select all that apply.)

    A. Blocking

    B. Listening

    C. Learning

    D. Forwarding

11. In a switched network, what is used to contain broadcasts and group physically separated hosts together in communities of interest.

    A. Switch

    B. Router

    C. VLAN

    D. Active topology

12. On a Cisco switch, which VLAN is used for untagged traffic?

    A. VLAN 1

    B. Management VLAN

    C. Native VLAN

    D. Untagged VLAN

13. What are three types of VTP messages and what are they used for?

14. How does STP elect a root bridge?

15. What are the three modes in VTP and how does each function?

Challenge Questions and Activities

These questions require a deeper application of the concepts covered in this chapter. You can find the answers in Appendix A.

1. You have just been hired by AnyCompany as its network administrator. Your first task is to centralize all the internal company servers into a server farm and to set up a method by which certain hosts can be denied access to this server farm based on a statically assigned IP address. What equipment would you need to purchase and what high-level configuration would you implement to accomplish this?

2. The AnyCompany network has grown to over 200 switches and 43 different VLANs. To simplify the management of this network, you are using VTP. You have just received a call from the company president, who is complaining that she no longer has access to the company network. After a bit of investigating, you determine that the VLAN configuration on the network has been altered. At the same time you are informed by the junior network administrator that he replaced a switch in the engineering department. What could have caused the VLAN configuration change and how would you prevent it in the future?