



CCNP BCMSN Quick Reference Sheets

Exam 642-812

Brent Stewart
Denise Donohue

The Evolving Network Model

VLAN Implementation

Spanning Tree

InterVLAN Routing

Layer 3 Redundancy

Using Wireless LANs

VoIP in a Campus Network

Campus Network Security



About the Authors

Brent Stewart, CCNP, CCDP, MCSE, Certified Cisco Systems Instructor, is a network administrator for CommScope. He participated in the development of BSCI, and has separately developed training material for ICND, BSCI, BCMSN, BCRAN, and CIT. Brent lives in Hickory, NC, with his wife, Karen, and children, Benjamin, Kaitlyn, Madelyn, and William.

Denise Donohue, CCIE No. 9566, is a Design Engineer with AT&T. She is responsible for designing and implementing data and VoIP networks for SBC and AT&T customers. Prior to that, she was a Cisco instructor and course director for Global Knowledge. Her CCIE is in Routing and Switching.

Icons Used in This Book



Router

7507
RouterMultilayer Switch
with TextMultilayer
SwitchCommunication
Server

Switch



Internal Firewall



IDS

Web
Browser

Database



App Server

The Evolving Network Model


Cisco has developed specific architecture recommendations for Campus, Data Center, WAN, branches, and telecommuting. These recommendations add specific ideas about how current technologies and capabilities match the network roles within an enterprise.

Each of these designs builds on a traditional hierarchical design and adds features such as security, Quality of Service (QoS), caching, and convergence.

The Hierarchical Design Model

Cisco has used the three level *Hierarchical Design Model* for years. This older model provided a high-level idea of how a reliable network might be conceived, but it was largely conceptual because it did not provide specific guidance. Figure 1-1 is a simple drawing of how the three-layer model might have been built out. A distribution layer-3 switch would be used for each building on campus, tying together the access-switches on the floors. The core switches would links the various buildings together.

FIGURE 1-1 THE HIERARCHICAL DESIGN MODEL



The hierarchical design model divides a network into three layers:

- Access—End stations attach to VLANs.
 - Clients attach to switch ports.
 - VLAN assigned/broadcast domains established.
 - Built using low-cost ports.
- Distribution—Intermediate devices route and apply policies.
 - VLANs terminated, routing between.
 - Policies applied, such as route selection.
 - Access-lists.
 - Quality of Service (QoS).

THE EVOLVING NETWORK MODEL

- **Core**—The backbone that provides a high-speed path between distribution elements.
 - Distribution devices are interconnected.
 - High speed (there is a lot of traffic).
 - No policies (it is tough enough to keep up).

Later versions of this model include redundant distribution and core devices, and connections that make the model more fault-tolerant. A set of distribution devices and their accompanying access layer switches are called a switch block.

Problems with the Hierarchical Design Model

This early model was a good starting point, but it failed to address key issues, such as:

- Where do wireless devices fit in?
- How should Internet access and security be provisioned?
- How to account for remote-access, such as dial-up or virtual private network (VPN)?
- Where should workgroup and enterprise services be located?

Enterprise Composite Network Model

The newer Cisco model—the Enterprise Composite Model—is significantly more complex and attempts to address the major shortcoming of the Hierarchical Design Model by expanding the older version and making specific recommendations about how and where certain network functions should be implemented. This model is based on the principles described in the Cisco Architecture for Voice, Video, and Integrated Data (AVVID).

The Enterprise Composite Model is broken up into three large sections:

- **Enterprise Campus**—The portion of the design that is like the old hierarchical model.
- **Enterprise Edge**—The connections to the public network.
- **Service Provider Edge**—The different public networks that are attached.

The first section, the Enterprise Campus, looks like the old Hierarchical model with some added details. The Enterprise Campus is shown in Figure 1-2. It features six sections:

- **Campus Backbone**—The center of the network, like the old “core”.
- **Building Distribution**—Intermediate devices that route from the core to access devices.

THE EVOLVING NETWORK MODEL

- Building Access—Connections for end systems.
- Management—Command, control, and auditing features.
- Edge Distribution—A distribution layer out to the WAN.
- Server Farm—For Enterprise services.

The Enterprise Edge (shown in Figure 1-3) details the connections from the campus to the Wide Area Network and includes:

- E-Commerce—Externally accessible services that have ties to internal data stores.
- Internet Connectivity—Connectivity to outside services.
- Remote Access—Dial and VPN.
- WAN—Internal links.

FIGURE 1-2 THE ENTERPRISE CAMPUS

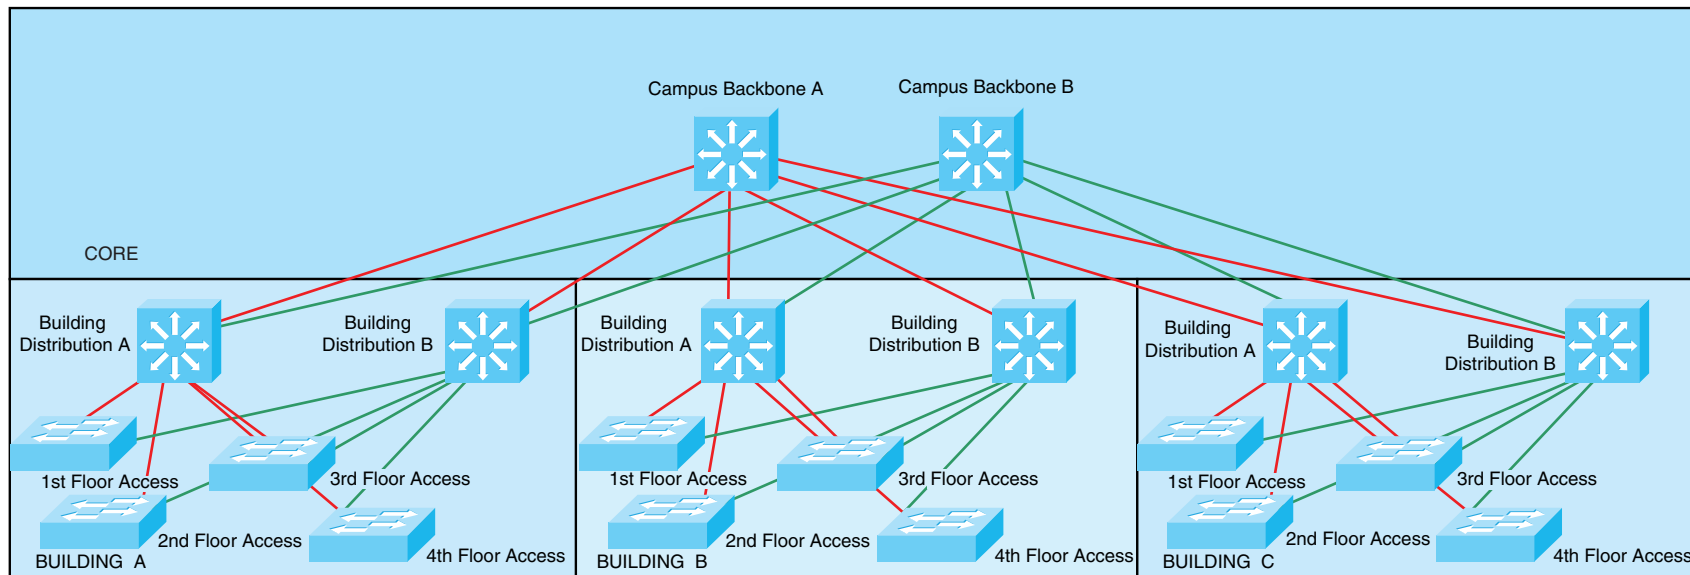
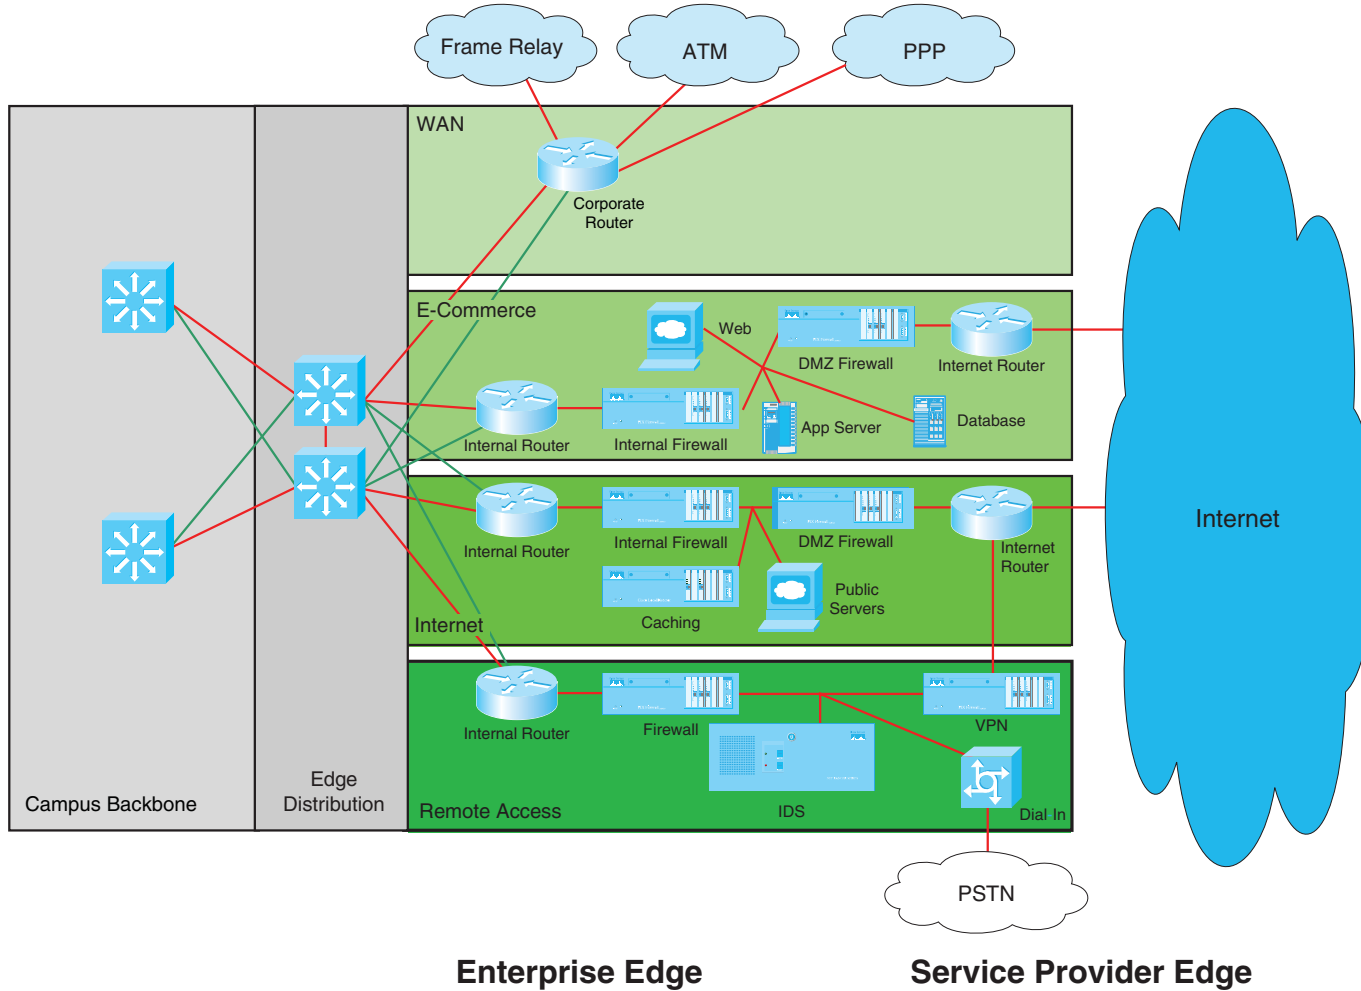


FIGURE 1-3 THE ENTERPRISE EDGE



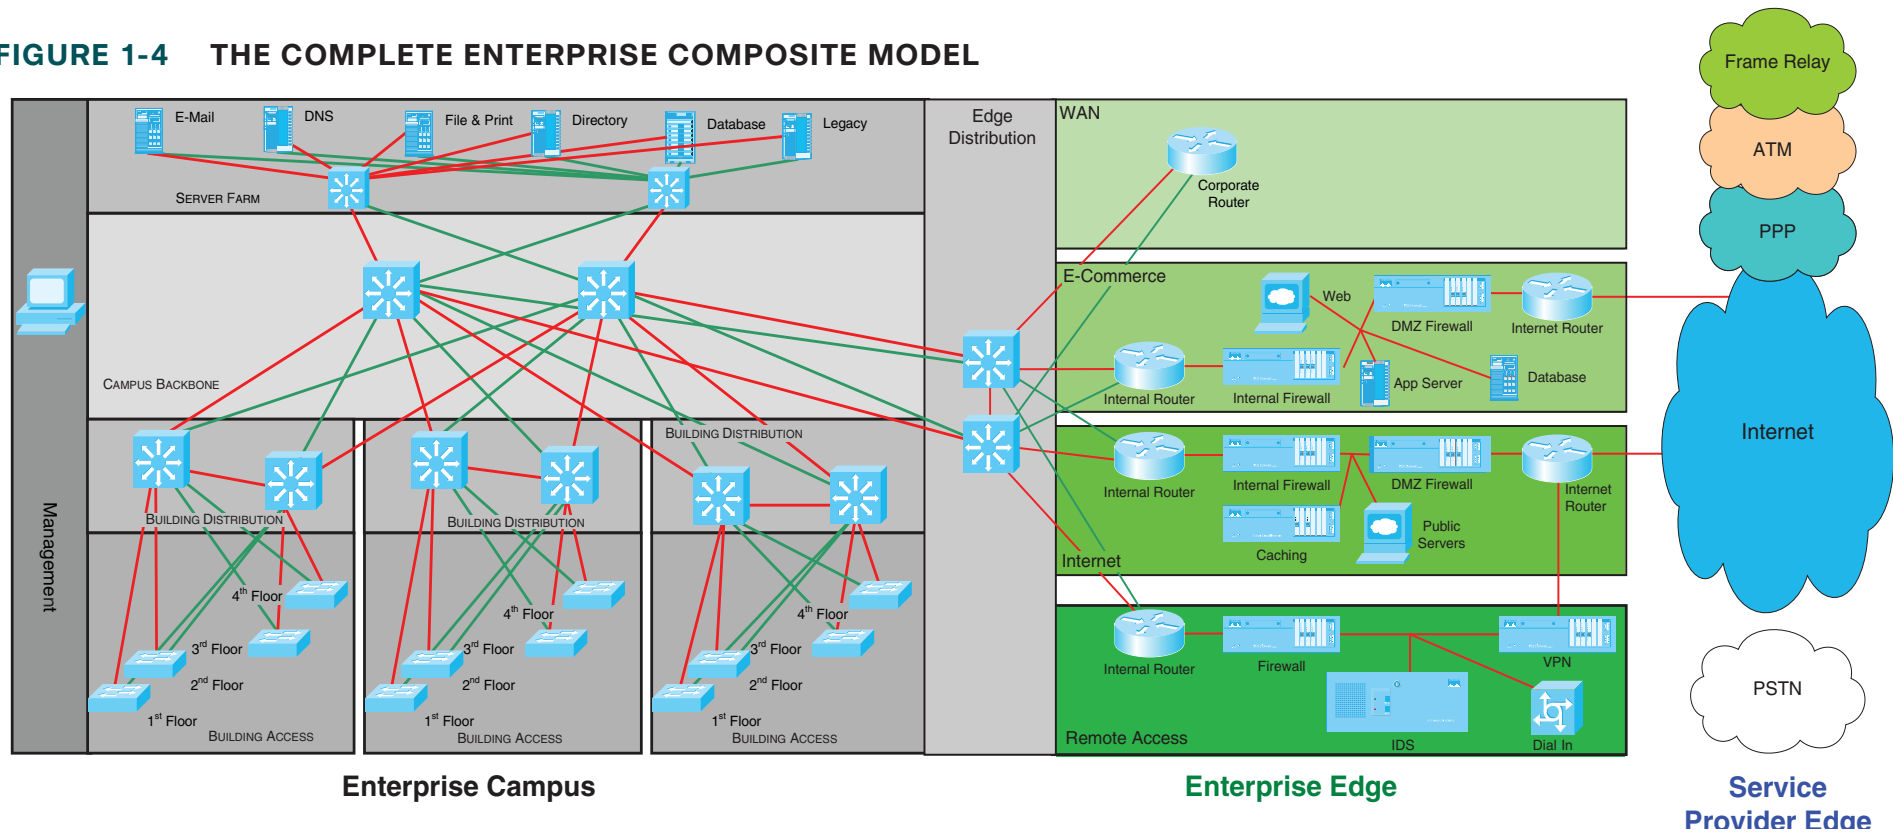
THE EVOLVING NETWORK MODEL

The Service Provider Edge consists of the public networks that facilitate wide-area network connectivity:

- Internet Service Provider (ISP)—Public connectivity
- Public Switched Telephone Network (PSTN)—Dial up
- Frame Relay, ATM, and PPP—Private connectivity

Figure 1-4 puts together the various pieces: Campus, Enterprise Edge, and Service Provider Edge. Security implemented on this model is described in the Cisco SAFE (Security Architecture for Enterprise) blueprint.

FIGURE 1-4 THE COMPLETE ENTERPRISE COMPOSITE MODEL



CHAPTER 1

THE EVOLVING NETWORK MODEL

SONA and IIN

Modern converged networks include different traffic types, each with unique requirements for security, QoS, transmission capacity, and delay. These include:

- Voice signaling and bearer
- Core Application traffic, such as Enterprise Resource Programming (ERP) or Customer Relationship Management (CRM)
- Database Transactions
- Multicast multimedia
- Network management
- “Other” traffic, such as web pages, e-mail, and file transfer

Cisco routers are able to implement filtering, compression, prioritization, and policing (dedicating network capacity). Except for filtering, these capabilities are referred to collectively as QoS.

Note

The best way to meet capacity requirements is to have twice as much bandwidth as needed. Financial reality, however, usually requires QoS instead.

Although QoS is wonderful, it is not the only way to address bandwidth shortage. Cisco espouses an ideal called the Intelligent Information Network (IIN).

IIN describes an evolutionary vision of a network that integrates network and application functionality cooperatively and allows the network to be smart about how it handles traffic to minimize the footprint of applications. IIN is built on top of the Enterprise Composite Model and describes structures overlaid on to the Composite design as needed in three phases.

Phase 1, “Integrated Transport,” describes a converged network, which is built along the lines of the Composite model and based on open standards. This is the phase that the industry has been transitioning to for the last few years, and the Cisco Integrated Services Routers (ISR) are an example of this trend.

Phase 2, “Integrated Services,” attempts to virtualize resources, such as servers, storage, and network access and move to an “on-demand” model.

By “virtualize” Cisco means that the services are not associated with a particular device or location. Instead, many services can reside in one device to ease management, or many devices can provide one service that is more reliable.

An ISR brings together routing, switching, voice, security, and wireless. It is an example of many services existing on one device. A load balancer, which makes many servers look like one, is a second example.

CHAPTER 1

THE EVOLVING NETWORK MODEL

VRFs are an example of taking one resource and making it look like many. Some versions of IOS are capable of having a router present itself as many virtual router forwarding (VRF) instances, allowing your company to deliver different logical topologies on the same physical infrastructure. Server virtualization is another example. The classic example of taking one resource and making it appear to be many resources is the use of a virtual LAN (VLAN) and a virtual storage area network (VSAN).

Virtualization provides flexibility in configuration and management.

Phase 3, “Integrated Applications,” uses application-oriented networking (AON) to make the network application-aware and to allow the network to actively participate in service delivery.

An example of this phase 3 IIN systems approach to service delivery is Network Admission Control (NAC). Before NAC, authentication, VLAN assignment, and anti-virus updates were separately managed. With NAC in place, the network is able to check the policy stance of a client and admit, deny, or remediate based on policies.

IIN allows the network to deconstruct packets, parse fields, and take actions based on the values it finds. An ISR equipped with an AON blade might be configured to route traffic from a business partner. The AON blade can examine traffic, recognize the application, and rebuild XML files in memory. Corrupted XML fields might represent an attack (called *schema poisoning*), so the AON blade could react by blocking

that source from further communication. In this example, routing, an awareness of the application data flow, and security are combined to allow the network to contribute to the success of the application.

Services-Oriented Network Architecture (SONA) applies the IIN ideals to Enterprise networks. Figure 1-5 shows how SONA breaks down the IIN functions into three layers:

- Network Infrastructure—Hierarchical converged network and attached end systems.
- Interactive Services—Resources allocated to applications.
- Applications—Includes business policy and logic.

THE EVOLVING NETWORK MODEL

FIGURE 1-5 IIN AND SONA COMPARED




IIN Phases

Phase 3 – Integrated Applications
 (“application aware”)

Phase 2 – Integrated Services (virtualized resources)

Phase 1 – Integrated Transport (converged network)

SONA Framework Layers



VLAN Implementation

VLANs are used to break large campus networks into smaller pieces. The benefit of this is to minimize the amount of broadcast traffic on a logical segment.

What Is a VLAN?

A virtual LAN (VLAN) is a logical LAN, or a logical subnet. It defines a broadcast domain. A physical subnet is a group of devices that shares the same physical wire. A logical subnet is a group of switch ports assigned to the same VLAN, regardless of their physical location in a switched network.

Two types of VLANs are:

- **End-to-end VLAN**—VLAN members are assigned by function and can reside on different switches. They are used when hosts are assigned to VLANs based on functions or workgroups, rather than physical location. VLANs should not extend past the Building Distribution submodule. Figure 2-1 shows end-to-end VLANs.
- **Local VLAN**—Hosts are assigned to VLANs based on their location, such as a floor in a building. A router accomplishes sharing of resources between VLANs. This type is typically found in the Building Access submodule. Figure 2-2 shows an example of local VLANs.

FIGURE 2-1 END-TO-END VLANS





FIGURE 2-2 LOCAL VLANS



VLAN membership can be assigned either statically by port or dynamically by MAC address using a VLAN Membership Policy Server (VMPS).

Best Practices

VLAN networks need many of the same considerations that normal Ethernet lines demand. For instance, VLANs should have one IP subnet. By supplying consecutive subnets to VLANs, the routing advertisements can be summarized (which has many benefits to convergence).

A stereotypical description of capacity requirements is possible. Access ports are assigned to a single VLAN and should be Fast Ethernet or faster. Ports to the distribution layer should be Gigabit Ethernet or better. Core ports are Gigabit Etherchannel or 10-Gig Ethernet. Remember that uplink ports need to be able to handle all hosts communicating concurrently, and remember that although VLANs logically separate traffic, traffic in different VLANs can still experience contention with other VLANs when both VLANs travel over the same trunk line.

Take into account the entire traffic pattern of applications found in your network. For instance, Voice VLANs pass traffic to a remote Call Manager. Multicast traffic has to communicate back to the routing process and possibly call upon a Rendezvous Point.

Creating a VLAN in Global Config Mode

VLANs must be created before they may be used. VLANs may be created in global configuration mode or in VLAN database mode. Creating VLANs in global configuration is easy—just identify the VLAN number and name it!

```
(config)#vlan 12
(config-vlan)#name MYVLAN
```

Creating a VLAN in Database Mode

Creating a VLAN in VLAN database mode is very similar to global configuration. There are no advantages to either method. Either method creates an entry in a VLAN.DAT file. Remember that copying the configuration, by itself, does not move the VLAN information! To do that you must move the VLAN.DAT file.

```
#vlan database
(vlan)#vlan 12 name MYVLAN
```

Delete a VLAN by using the same command with **no** in front of it. There is no need to include the name when deleting.

VLAN IMPLEMENTATION

Assigning Ports to VLANs

When statically assigning ports to VLANs, first make it an access port, and then assign the port to a VLAN. At the interface configuration prompt:

```
(config-if)#switchport mode access
(config-if)#switchport access vlan 12
```

The commands are similar when using dynamic VLAN assignment. At interface configuration mode:

```
(config-if)#switchport mode access
(config-if)#switchport access vlan dynamic
```

If you use dynamic, you must also enter the IP address of the VMPS server at global configuration mode:

```
(config-if)#vmps server ip address
```

Verifying VLAN Configuration

To see a list of all the VLANs and the ports assigned to them, use the command **show vlan**. To narrow down the information displayed, you can use these keywords after the command: **brief**, **id**, *vlan-number*, or **name** *vlan-name*:

```
ASW# show vlan brief
VLAN Name      Status      Ports
-----
1      default     active      Fa0/1, Fa0/2, Fa0/3,
          Fa0/10,Fa0/11,Fa0/12
20     VLAN0020    active      Fa0/5,Fa0/6,Fa0/7
21     VLAN0021    active      Fa0/8,Fa0/9
```

```
1002 fddi-default    active
1003 trcrf-default  active
1004 fddinet-default active
1005 trbrf-default  active
```

Other verification commands include:

- **show running-config interface** *interface no.*—Use the following to verify the VLAN membership of the port:

```
ASW# show run interface fa0/5
Building configuration...
Current configuration 64 bytes
interface FastEthernet 0/5
  switchport access vlan 20
  switchport mode access
```

- **show mac address-table interface** *interface no. vlan vlan no.*—Use the following to view MAC addresses learned through that port for the specified VLAN:

```
ASW# show mac address-table interface fa0/1
Mac Address Table
-----
Vlan      Mac Address      Type              Ports
---      -
1         0030.b656.7c3d   DYNAMIC           Fa0/1
Total Mac Addresses for this criterion: 1
```

- **show interfaces** *interface no. switchport*—Use the following to see detailed information about the port configuration, such as entries in the Administrative Mode and Access Mode VLAN fields:

```
ASW# show interfaces fa0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: dynamic desirable
```

VLAN IMPLEMENTATION

```

Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Protected: false
Unknown unicast blocked: false
Unknown multicast blocked: false
Broadcast Suppression Level: 100
Multicast Suppression Level: 100
Unicast Suppression Level: 100

```

Troubleshooting VLAN Issues

The following are three steps in troubleshooting VLAN problems:

- Check the physical connectivity—Make sure the cable, the network adapter, and switch port are good. Check the port's link LED.
- Check the switch configuration—If you see FCS errors or late collisions, suspect a duplex mismatch. Also check configured speed on both ends of the link. Increasing collisions can mean an overloaded link, such as with a broadcast storm.
- Check the VLAN configuration—If two hosts cannot communicate, make sure they are both in the same VLAN. If a host cannot connect to a switch, make sure the host and the switch are in the same VLAN.

VLAN Trunking

A *trunk* is a link that carries traffic for more than one VLAN. Trunks multiplex traffic from multiple VLANs. Trunks connect switches and allow ports on multiple switches to be assigned to the same VLAN.

Two methods of identifying VLANs over trunk links are:

- Inter-Switch Link (ISL)—A Cisco proprietary method that encapsulates the original frame in a header, which contains VLAN information. It is protocol-independent and can identify Cisco Discovery Protocol (CDP) and bridge protocol data unit (BPDU) frames.
- 802.1Q—Standards-based, tags the frames (inserts a field into the original frame immediately after the source MAC address field), and supports Ethernet and Token Ring networks.

When a frame comes into a switch port, the frame is tagged internally within the switch with the VLAN number of the port. When it reaches the outgoing port, the internal tag is removed. If the exit port is a trunk port, then its VLAN is identified in either the ISL encapsulation or the 802.1Q tag. The switch on the other end of the trunk removes the ISL or 802.1Q information, checks the VLAN of the frame, and adds the internal tag. If the exit port is a user port, then the original frame is sent out unchanged, making the use of VLANs transparent to the user.

If a nontrunking port receives an ISL-encapsulated frame, the frame is dropped. If the ISL header and footer cause the MTU size to be

VLAN IMPLEMENTATION

exceeded, it might be counted as an error.

If a nontrunking port receives an 802.1Q frame, the source and destination MAC addresses are read, the tag field is ignored, and the frame is switched normally at Layer 2.

Configuring a Trunk Link

Ports can become trunk ports either by static configuration or dynamic negotiation using Dynamic Trunking Protocol (DTP). A switch port can be in one of five DTP modes:

- Access—The port is a user port in a single VLAN.
- Trunk—The port negotiates trunking with the port on the other end of the link.
- Non-negotiate—The port is a trunk and does not do DTP negotiation with the other side of the link.
- Dynamic Desirable—Actively negotiates trunking with the other side of the link. It becomes a trunk if the port on the other switch is set to **trunk**, **dynamic desirable**, or **dynamic auto** mode.
- Dynamic Auto—Passively waits to be contacted by the other switch. It becomes a trunk if the other end is set to **trunk** or **dynamic desirable** mode.

Configure a port for trunking at the interface configuration mode:

```
(config-if)#switchport mode {dynamic {auto | desirable} | trunk}
```

If dynamic mode is used, DTP negotiates the trunking state and encapsulation. If trunk mode is used, you must specify encapsulation:

```
(config-if)#switchport trunk encapsulation {isl | dot1q | negotiate}
```

Native VLAN with 802.1Q

If you are using 802.1Q, specify a native VLAN for the trunk link with the command:

```
(config-if)#switchport trunk native vlan vlan no
```

Frames from the native VLAN are sent over the trunk link untagged. Native VLAN is the VLAN the port would be in if it were not a trunk, and it must match on both sides of the trunk link. VLAN 1 is the default native VLAN for all ports.

VLAN Mapping

ISL trunking recognizes only VLANs numbered 1–1001, but 802.1Q can use VLANs 0–4094. If you are using both ISL and 802.1Q in your network and have VLANs numbered above 1001, you have to map the 802.1Q VLANs to ISL numbers. Some rules about mapping VLANs include:

- You can configure only eight mappings.
- Mappings are local to the switch; the same mappings must be configured on all switches in the network.

VLAN IMPLEMENTATION

- You can map only to Ethernet ISL VLANs.
- The 802.1Q VLANs with the same number as mapped ISL VLANs are blocked. (For example, you map 802.1Q VLAN 1500 to ISL VLAN 150, then 802.1Q VLAN 150 is blocked on that switch.)
- You should not map the 802.1Q native VLAN.

VLANs Allowed on the Trunk

By default, a trunk carries traffic for all VLANs. You can change that behavior for a particular trunk link by giving the following command at the interface config mode:

```
switchport trunk allowed vlan vians
```

Make sure that both sides of a trunk link allow the same VLANs.

Verifying a Trunk Link

Two commands you can use to verify your trunk configuration are:

```
#show running-config
#show interfaces [interface no.] switchport | trunk
```

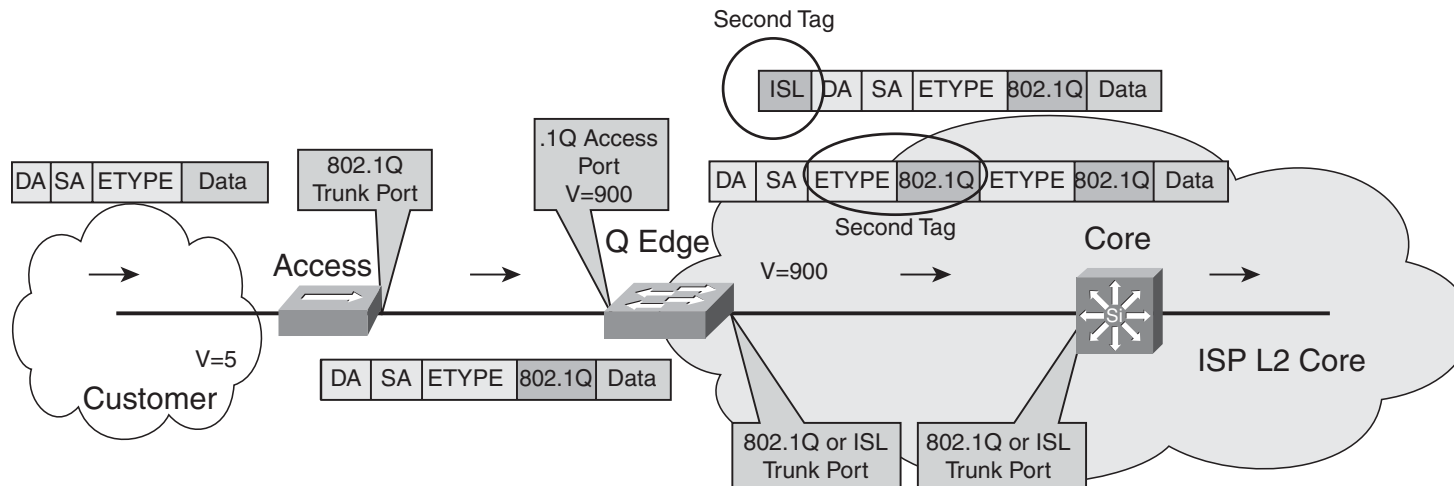
Using the **trunk** keyword with the **show interfaces** command gives information about the trunk link:

```
# show interfaces fastethernet 0/1 trunk
Port      Mode           Encapsulation  Status      Native
vlan
Fa0/1     desirable     n-802.1q       trunking    1
Port      Vlans allowed on trunk
Fa0/1     1-150
<further output omitted>
```

802.1Q Tunnels

Tunneling is a way to send 802.1Q-tagged frames across a foreign network (such as a Service Provider's network) and still preserve the original 802.1Q tag. The SP configures their end of the trunk link as a tunnel port and assigns a VLAN to carry your traffic within their network. The SP switch then adds a second 802.1Q tag to each frame that came in the tunnel port. Other switches in the SP network see only this second tag, and do not read the original tag. When the frame exits the SP network, the extra tag is removed, leaving the original 802.1Q tag to be read by the receiving switch in your network.

FIGURE 2-3 802.1Q



Layer 2 Protocol Tunneling (GBPT)

If a Service Provider separates sections of your network, you can use Layer 2 protocol tunneling to tunnel CDP, Spanning Tree Protocol (STP), and VLAN Trunking Protocol (VTP) frames across the SP's cloud. This is called Generic Bridge PDU Tunneling (GBPT). Frames from the above control protocols are encapsulated as they enter the SP's network on a tunnel port, and de-encapsulated when they exit that network.

Troubleshooting Trunking

Troubleshooting trunking links happens mostly at the physical and datalink layers. Start with the most basic assumptions and work your way "up" the OSI model. It is important to show that physical layer connectivity is present, before moving on to, for instance before trying to troubleshoot IP problems.

- Are both sides of the link in the correct trunking mode?
- Is the same trunk encapsulation on both sides?
- If 802.1Q, is the same native VLAN on both sides?
- Are the same VLANs permitted on both sides?

CHAPTER 2

VLAN IMPLEMENTATION

VLAN Trunking Protocol (VTP)

VTP is a protocol that runs over trunk links and synchronizes the VLAN databases of all switches in the VTP domain. A VTP domain is an administrative group—all switches within that group must have the same VTP domain name configured or they do not synchronize databases.

VTP works by using Configuration Revision numbers and VTP advertisements:

- All switches send out VTP advertisements every five minutes, or when there is a change to the VLAN database (when a VLAN is created, deleted, or renamed).
- VTP advertisements contain a Configuration Revision number. This number is increased by one for every VLAN change.
- When a switch receives a VTP advertisement, it compares the Configuration Revision number against the one in its VLAN database.
- If the new number is higher, the switch overwrites its database with the new VLAN information, and forwards the information to its neighbor switches.
- If the number is the same, the switch ignores the advertisement.
- If the new number is lower, the switch replies with the more up-to-date information contained in its own database.

VTP Switch Roles

A switch can be a VTP:

- **Server**—The default VTP role. Servers can create, delete, and rename VLANs. They originate both periodic and triggered VTP advertisements and synchronize their databases with other switches in the domain.
- **Client**—Clients cannot make VLAN changes. They originate periodic VTP advertisements and synchronize their databases with other switches in the domain.
- **Transparent**—It can create, delete, and rename VLANs, but its VLANs are only local. It does not originate advertisements or synchronize its database with any other switches. It forwards VTP advertisements out its trunk links, however.

VTP Pruning

By default, switches flood broadcasts, multicasts, and unknown unicasts across trunk links. Suppose a host in VLAN 10 on Switch B sends a broadcast. Hosts in VLAN 10 on Switch C need to see that broadcast, but Switch A has no ports in VLAN 10, so it doesn't need to receive the broadcast traffic.

Enabling VTP pruning causes the switch to keep track of VLAN port assignments in its downstream switches. The switch then sends flooded traffic only on trunks toward switches that have ports assigned to the

VLAN IMPLEMENTATION

VLAN originating the traffic. It prunes flooded traffic from all other trunks. VTP pruning increases the available bandwidth by preventing unnecessary traffic on trunk links.

There are two versions of VTP: Version 1 and Version 2. To use Version 2, all switches in the domain must be capable of using it. Configure one server for Version 2, and the information is propagated through VTP. Version 2 has the following added features:

- It supports Token Ring VLANs.
- Transparent switches pass along messages from both versions of VTP.
- Consistency checks are performed only when changes are configured through the CLI or SNMP.

Configuring VTP

VTP configuration is done at the global config mode. To configure the switch's VTP mode:

```
(config)#vtp {server | client |transparent}
```

To configure the VTP domain name:

```
(config)#vtp domain name
```

To configure a VTP password (all switches in the domain must use the same password):

```
(config)#vtp password password
```

To configure the switch to use VTP Version 2:

```
(config)#vtp version 2
```

To enable pruning:

```
vtp pruning
```

To specify which VLANs are to be pruned:

```
(config-if)#switchport trunk pruning vlan {add | except | none  
| remove} vlan-list [,vlan[,vlan[, , ,]]
```

Verifying and Monitoring VTP

To get basic information about the VTP configuration, use **show vtp status**. The example shows the default settings:

```
# show vtp status
VTP Version      : 1
Configuration Revision      : 0
Maximum VLANs supported locally      : 1005
Number of existing VLANs      : 5
VTP Operating Mode      : Server
VTP Domain Name      :
(config)#
VTP Pruning Mode      : Disabled
VTP V2 Mode      : Disabled
VTP Traps Generation      : Disabled
MD5 digest      :
```

CHAPTER 2

VLAN IMPLEMENTATION

Troubleshooting VTP

The following are some common things to check when troubleshooting problems with VTP:

- Make sure you are trunking between the switches. VTP is sent only over trunk links.
- Make sure the domain name matches on both switches (name is case sensitive).
- If the switch is not updating its database, make sure it is not in transparent mode.
- If using passwords, make sure they all match. To remove a password, use **no vtp password**.

Adding a New Switch to a VTP Domain

Adding a new switch in client mode does not prevent it from propagating its incorrect VLAN information. A server synchronizes to a client if the client has the higher configuration revision number. You must reset the revision number back to 0 on the new switch. The easiest way to do this is to change the domain name. Then change it back to the correct one, and attach the switch to the network.

Spanning Tree

Ethernet network design balances two separate imperatives. First, Ethernet has no capacity for detecting circular paths. If such paths exist, traffic loops around and accumulates until new traffic is shut out (this is called a broadcast storm). Second, having secondary paths is good preparation for inevitable link failure.

Spanning Tree is a protocol that prevents loop formation by detecting redundant links and disabling them until needed. Designers can therefore build redundant links and the protocol will allow one to pass traffic and keep the other in reserve. When the active link fails, the secondary link is enabled quickly.

Understanding the Spanning Tree Protocol

Switches either forward or filter Layer 2 frames. The way they make the forwarding/filtering decision can lead to loops in a network with redundant links. Spanning Tree is a protocol that detects potential loops and breaks them.


A Layer 2 switch is functionally the same thing as a transparent bridge. Transparent bridges:

- Learn MAC (Media Access Control) addresses by looking at the source address of incoming frames. They build a table mapping MAC address to port number.

- Forward broadcasts and multicasts out all ports except the one which they came. (This is called flooding.)
- Forward unknown unicasts out all ports except the one on which they came. An unknown unicast is a message bound for a unicast MAC address that is not in the switch's table of addresses and ports.
- Do not make any changes to the frames as they forward them.

Spanning Tree Protocol (STP) works by selecting a root bridge, then selecting one loop-free path from the root bridge to every other switch. (STP uses the term bridge because it was written before there were switches.) Consider the following switched network (see Figure 3-1).

FIGURE 3-1 EXAMPLE SWITCHED TOPOLOGY



SPANNING TREE

Spanning Tree must select:

- One root bridge
- One root port per nonroot bridge
- One designated port per network segment

Spanning Tree Election Criteria

Spanning Tree builds paths out from a central point along the fastest available links. It selects path according to the following criteria:

1. Lowest root bridge ID (BID)
2. Lowest path cost to the root
3. Lowest sender bridge ID
4. Lowest sender port ID (PID)

When reading the path selection criteria, remember the following:

- Bridge ID—Bridge priority: Bridge MAC address.
- Bridge priority—2-byte value, 0–65,535 (0–0xFFFF).
- Default priority is 32,768 (0x8000).
- Port ID—Port priority: port number.
- Port priority—A 6-bit value, 0–63, default is 32.

- Path cost—This is the cumulative value of the cost of each link between the bridge and the root. Cost values were updated in 2000 and you should see only new cost values, but both are given in the following table (see Table 3-1). Old and new switches work together.

TABLE 3-1: Spanning Tree Costs

Link Speed	Old Cost	New Cost
10 Mbps	100	100
100 Mbps	10	19
1 Gbps	1	4
10 Gbps	1	2

The STP Election

Spanning Tree builds paths out from a starting point, the “root” of the tree. The first step in selecting paths is to identify this root device. Then, each device selects its best path back to the root, according to the criteria laid out in the previous sections (lowest root BID, lowest cost, lowest advertising BID, lowest port).

Root Bridge Election

Looking at Figure 3-1, first select the root bridge. Assume each switch uses the default priority.

- Switch A BID = 80-00-00-0c-11-11-00-11
- Switch B BID = 80-00-00-0c-26-78-10-10
- Switch C BID = 80-00-00-0c-32-1a-bc-de
- Switch D BID = 80-00-00-0c-81-81-11-22
- Switch E BID = 80-00-00-0c-26-79-22-22

Switch A has the lowest BID, so it is the root. Each nonroot switch must now select a root port.

Root Port Election

The root port is the port that leads back to the root. Continuing with Figure 3-1, once A is acknowledged as the root, the remaining bridges sort out their lowest cost path back to the A.

- **Switch B**—Uses the link to A with a cost of 19 (link speed of 100 Mbps).
- **Switch C**—The connected link has a cost of 100 (Ethernet), the link through B has a path cost of 38 (two 100 Mbps links), and so B is chosen.
- **Switch D**—The link through B has a path cost of 119, the path cost through C to A is 119, the path through C then B is 57, so C is chosen.

- **Switch E**—The lowest path cost is the same for both ports (76 through D to C to B to A). Next check sender BID—sender for both ports is D, so that it does not break the tie. Next check sender Port ID. Assuming default port priority, the PID for 0/1 is lower than the PID for 0/2, so the port on the left is the root port.


Designated Port Election

Designated ports are ports that lead away from the root. Obviously, all ports on the root bridge are designated ports (A-B and A-C in Figure 3-1).

- **Segment B-D**—B has the lowest path cost to root (19 vs 119), so it is designated for this segment.
- **Segment C-D**—C has the lowest path cost to the root (100 vs 119), so it is designated for this segment.
- **Segment B-C**—B has the lowest path cost to the root (19 vs 100), so it is designated for this segment.
- **Both segments D-E**—D has the lowest cost to the root (57 vs 76), so it is designated for both segments.

Now the looped topology has been turned into a tree with A at the root. Notice that there are no more redundant links.

FIGURE 3-2 THE ACTIVE TOPOLOGY AFTER SPANNING TREE IS COMPLETE



Bridge Protocol Data Units (BPDUs)

Switches exchange BPDUs. There are two types of BPDUs: Configuration and Topology Change (TCN).

Configuration BPDUs are sent every two seconds from the root toward the downstream switches. They:

- Are used during an election.
- Maintain connectivity between switches.
- Send timer information from the root.

TCN BPDUs are sent toward the root when:

- There is a link failure.
- A port starts forwarding, and there is already a designated port.
- The switch receives a TCN from a neighbor.

When a switch receives a TCN BPDU, it acknowledges that with a configuration BPDU that has the TCN Acknowledgment bit set.

When the root bridge receives a TCN, it starts sending configuration BPDUs with the TCN bit set for a period of time equal to max age plus forward delay. Switches that receive this change their MAC table aging time to the Forward Delay time, causing MAC addresses to age faster. The topology change also causes an election of the root bridge, root ports, and designated ports.

BPDU Fields

Some of the fields in the BPDU include:

- Root bridge ID—The BID of the current root.
- Sender's root path cost—The cost to the root.
- Sender's bridge ID—Sender's priority concatenated to MAC.
- Sender's port ID—The port number, transmitted as final tie-breaker.
- Hello time—Two seconds by default.
- Forward Delay—15 seconds by default.
- Max Age—20 seconds by default.

Spanning Tree Port States

When a port is first activated, it transitions through the following stages shown in Table 3-2.

TABLE 3-2: Spanning Tree Port States

Port State	Timer	Actions
Blocking	Max Age (20 sec)	Discards frames, does not learn MAC addresses, receives BPDUs.
Listening	Forward Delay (15 sec)	Discards frames, does not learn MAC addresses, receives BPDUs to determine its role in the network.

TABLE 3-2: Spanning Tree Port States

Port State	Timer	Actions
Learning	Forward Delay (15 sec)	Discards frames, does learn MAC addresses, receives and transmits BPDUs.
Forwarding		Accepts frames, learns MAC addresses, receives and transmits BPDUs.

Designing for Spanning Tree

To optimize data flow in the network, design and configure switches for the following STP roles:

- Primary and secondary root bridges (set priority values)
- Designated and root ports (set port priorities/path cost)
- Enable STP enhancements, such as Root Guard

Spanning Tree and PVST

With PVST (Per Vlan STP), there is a different instance of STP for each VLAN. To derive the VLAN BID, the switch picks a different MAC address from its base pool for each VLAN. Each VLAN has its own root bridge, root port, and so on. You can configure these so that data flow is optimized, and traffic load is balanced among the switches.

Spanning Tree is enabled by default on every VLAN.

Configuring Spanning Tree

To change the STP priority value, use the following:

```
Switch (config)#spanning-tree vlan vlan_no. priority value
```

To configure a switch as root without manually changing priority values, use the following:

```
Switch (config)# spanning-tree vlan vlan_no. root {primary | secondary}
```

To change the STP port cost for an access port, use the following:

```
Switch(config-if)# spanning-tree cost value
```

To change the STP port cost for a VLAN on a trunk port, use the following:

```
Switch(config-if)# spanning-tree vlan vlan_no. cost value
```

To display STP information for a VLAN, use the following:

```
Switch# show spanning-tree vlan vlan_no.
```

To display the STP information for an interface, use the following:

```
Switch # show spanning-tree interface interface_no. [detail]
```

To verify STP timers, use the following:

```
Switch #show spanning-tree bridge brief
```

Spanning Tree Enhancements

Cisco has some proprietary enhancements to Spanning Tree that help speed up network convergence. They include:

- PortFast
- UplinkFast
- BackboneFast

Portfast

Portfast is for access (user) ports only. It causes the port to bypass the STP listening and learning states and transition directly to forwarding. Connecting a switch to a Portfast port can cause loops to develop.

```
(config-if)#spanning-tree portfast
```

UplinkFast

UplinkFast is for speeding convergence when a direct link to an upstream switch fails. The switch identifies backup ports for the root port (these are called an uplink group). If the root port fails, then one of the ports in the uplink group is unblocked and transitions immediately to forwarding—it bypasses the listening and learning stages. It should be used in wiring closet switches with at least one blocked port.

The command to enable uplinkfast is shown below. Please note that uplinkfast is enabled globally, so the command affects all ports and all VLANs.

```
(config)# spanning-tree uplinkfast
```

BackboneFast

BackboneFast is used for speeding convergence when a link fails that is not directly connected to the switch. It helps the switch detect indirect failures. If a switch running BackboneFast receives an inferior BPDU from its designated bridge, it knows a link on the path to the root has failed. (An inferior BPDU is one that lists the same switch for root bridge and designated bridge.)

The switch then tries to find an alternate path to the root by sending a Root Link Query (RLQ) frame out all alternate ports. The root then responds with an RLQ response, and the port receiving this response can transition to forwarding. Alternate ports are determined in this way:

- If the inferior BPDU was received on a blocked port, then the root port and any other blocked ports are considered alternates.
- If the inferior BPDU was received on the root port, then all blocked ports are considered alternates.
- If the inferior BPDU was received on the root port and there are no blocked ports, the switch assumes it has lost connectivity with the root and advertises itself as root.

Configure this command on all switches in the network:

```
(config)#spanning-tree backbonefast
```

Rapid Spanning Tree (RSTP)

Rapid Spanning Tree (RSTP) 802.1w is a standards-based, non-proprietary way of speeding STP convergence. Switch ports exchange an explicit handshake when they transition to forwarding. RSTP describes different port states than regular STP, as shown in the Table 3-3.

TABLE 3-3: Comparing 802.1d and 802.1w Port States

STP Port State	Equivalent RSTP Port State
Disabled	Discarding
Blocking	Discarding
Listening	Discarding
Learning	Learning
Forwarding	Forwarding

RSTP Port Roles

RSTP also defines different Spanning Tree roles for ports:

- Root port—The best path to the root (same as STP).
- Designated port—Same role as with STP.
- Alternate port—A backup to the root port.
- Backup port—A backup to the designated port.
- Disabled port—One not used in the Spanning Tree.
- Edge port—One connected only to an end user.

BPDUs Differences in RSTP

In regular STP, BPDUs are originated by the root and relayed by each switch. In RSTP, each switch originates BPDUs, whether or not it receives a BPDU on its root port. All eight bits of the BPDU type field are used by RSTP. The TC and TC Ack bits are still used. The other six bits specify the port's role and its RSTP state, and are used in the port handshake. The RSTP BPDU is set to Type 2, Version 2. PVST is done by Rapid PVST+ on Catalyst switches.

RSTP Fast Convergence

The Rapid Spanning tree process understands and incorporates topology changes much quicker than the previous version.

- RSTP uses a mechanism similar to BackboneFast—When an inferior BPDU is received, the switch accepts it. If the switch has another path to the root, it uses that and informs its downstream switch of the alternate path.
- Edge ports work the same as Portfast ports—They automatically transition directly to forwarding.
- Link type—If you connect two switches through a point-to-point link and the local port becomes a designated port, it exchanges a handshake with the other port to quickly transition to forwarding. Full-duplex links are assumed to be point-to-point, half-duplex links are assumed to be shared.

- Backup and alternate ports—Ports that can transition to forwarding when no BPDUs are received from a neighbor switch (similar to UplinkFast).

If an RSTP switch detects a topology change, it sets a TC timer to twice the hello time and sets the TC bit on all BPDUs sent out to its designated and root ports until the timer expires. It also clears the MAC addresses learned on these ports.

If an RSTP switch receives a TC BPDU, it clears the MAC addresses on that port and sets the TC bit on all BPDUs sent out its designated and root ports until the TC timer expires.

Multiple Spanning Tree (MST)

With Multiple Spanning Tree (MST), you can group VLANs and run one instance of Spanning Tree for a group of VLANs. This cuts down on the number of root bridges, root ports, designated ports, and BPDUs in your network. Switches in the same MST Region share the same configuration and VLAN mappings. Configure MST with these commands:

```
(config)# spanning-tree mode mst
(config)# spanning-tree mst configuration
(config-mst)# name region_name
(config-mst)# revision number
(config-mst)# instance number vlan vlan_range
(config-mst)# end
```

To be compatible with 802.1Q trunking, which has one common Spanning Tree (CST) for all VLANs, MST runs one instance of an

CHAPTER 3

SPANNING TREE

Internal Spanning Tree (IST). The IST appears as one bridge to a CST area and is MST instance number 0. The original MST Spanning Trees (called M-Trees) are active only in the region—they combine at the edge of the CST area to form one.

EtherChannels

EtherChannel is a way of combining several physical links between switches into one logical connection. Normally, Spanning Tree blocks redundant links; EtherChannel gets around that and allows load balancing across those links. Load is balancing on the basis of such things as source or destination MAC address or IP address. The Etherchannel load-balancing method is configured at global configuration mode.

```
(config)#port-channel load-balance type
```

A logical interface—the Port Channel interface—is created. Configuration can be applied to both the logical and physical interfaces.

Some guidelines for EtherChannels are as follows:

- Interfaces in the channel do not have to be physically next to each other or on the same module.
- All ports must be the same speed and duplex.
- All ports in the bundle should be enabled.
- None of the bundle ports can be a SPAN port.
- Assign an IP address to the logical Port Channel interface, not the physical ones.

- Put all bundle ports in the same VLAN, or make them all trunks. If they are trunks, they must all carry the same VLANs and use the same trunking mode.
- Configuration you apply to the Port Channel interface affects the entire EtherChannel. Configuration you apply to a physical interface only affects that interface.

Configuring an EtherChannel

Basically, for a Layer 3 EtherChannel, you should configure the logical interface and then put the physical interfaces into the channel group:

```
(config)#interface port-channel number
(config-if)#no switchport
(config-if)#ip address address mask
```

Then, at each port that is part of the EtherChannel, use the following:

```
(config)#interface { number | range interface - interface }
(config-if)#channel-group number mode { auto | desirable | on }
```

Putting the IP address on the Port Channel interface creates a Layer 3 EtherChannel. Simply putting interfaces into a channel group creates a Layer 2 EtherChannel, and the logical interface is automatically created.

The Cisco proprietary Port Aggregation Protocol (PAgP) dynamically negotiates the formation of a channel. There are three PAgP modes:

- On—The port channels without using PAgP negotiation. The port on the other side must also be set to On.

CHAPTER 3

SPANNING TREE

- Auto—Responds to PAgP messages but does not initiate them. Port channels if the port on the other end is set to Desirable. This is the default mode.
- Desirable—Port actively negotiates channeling status with the interface on the other end of the link. Port channels if the other side is Auto or Desirable.

There is also a non-proprietary protocol called Link Aggregation Control Protocol (LACP), IEEE 802.3ad, which does the same thing. LACP has two modes:

- Active—Port actively negotiates channeling with the port on the other end of the link. A channel forms if the other side is Passive or Active.
- Passive—Responds to LACP messages but does not initiate them. A channel forms if the other end is set to Active.

If you want to use LACP, specify it under the interface and put the interface in either active or passive mode:

```
(config-if)#channel-protocol lacp
```

Verifying an EtherChannel

Some typical commands for verifying include:

- `#show running-config interface number`
- `#show interfaces number etherchannel`

- `#show etherchannel number port-channel`

- `#show etherchannel summary`

Additional Spanning Tree Features

Some additional features available to help you tune Spanning Tree include:

- BPDU Guard
- BPDU Filtering
- Root Guard
- UDLD
- Loop Guard

BPDU Guard

BPDU Guard is used to prevent loops if another switch is attached to a Portfast port. When BPDU Guard is enabled on an interface, it is put into an error-disabled state (basically, shut down) if a BPDU is received on the interface. It can be enabled at either global config mode—in which case it affects all Portfast interfaces, or at interface mode. Portfast does not have to be enabled for it to be configured at a specific

CHAPTER 3

SPANNING TREE

interface. The following configuration example shows BPDU guard being enabled.

```
(config)#spanning-tree portfast bpduguard default
(config-if)#spanning-tree bpduguard enable
```

BPDU Filtering

BPDU filtering is another way of preventing loops in the network. It also can be enabled either globally or at the interface, and functions differently at each. In global config, if a Portfast interface receives any BPDUs, it is taken out of Portfast status. At interface config mode, it prevents the port from sending or receiving BPDUs. The commands are:

- (config)# **spanning-tree portfast bpdudfilter default**
- (config-if)# **spanning-tree bpdudfilter enable**

Root Guard

Root Guard is meant to prevent the wrong switch from becoming the Spanning Tree root. It is enabled on ports other than the root port and on switches other than the root. If a Root Guard port receives a BPDU that might cause it to become a root port, then the port is put into “root-inconsistent” state and does not pass traffic through it. If the port stops receiving these BPDUs, it automatically re-enables itself.

```
(config-if)# spanning-tree guard root
```

Unidirectional Link Detection (UDLD)

A switch notices when a physical connection is broken by the absence of Layer 1 electrical keepalives (Ethernet calls this a link beat). However, sometimes a cable is intact enough to maintain keepalives, but not to pass data in both directions. This is a Unidirectional Link. Unidirectional Link Detection (UDLD) detects a unidirectional link by sending periodic hellos out to the interface. It also uses probes, which must be acknowledged by the device on the other end of the link. UDLD operates at Layer 2. The port is shut down if a unidirectional link is found.

To enable UDLD on all fiber-optic interfaces, use the following command:

```
(config)# udld enable
```

Although this command is given at global config mode, it applies only to fiber ports.

To enable UDLD on non-fiber ports, give the same command at interface config mode.

To disable UDLD on a specific fiber port, use the following command:

```
(config-if)# udld disable
```

To disable UDLD on a specific non-fiber port, use the following command:

```
(config-if)#no udld enable
```


SPANNING TREE

To re-enable all interfaces shut by UDLD, use the following:

```
#udld reset
```

To verify UDLD status, use the following:

```
#show udld interface
```

Loop Guard

Loop Guard prevents loops that might develop if a port that should be blocking inadvertently transitions to the forwarding state. This can happen if the port stops receiving BPDUs (perhaps because of a unidirectional link or a software/configuration problem in its neighbor switch). When one of the ports in a physically redundant topology stops receiving BPDUs, the STP conceives the topology as loop-free. Eventually, the blocking port becomes designated and moves to forwarding state, thus creating a loop. With Loop Guard enabled, an additional check is made.

If no BPDUs are received on a blocked port for a specific length of time, Loop Guard puts that port into “loop inconsistent” blocking state, rather than transitioning to forwarding state. Loop Guard should be enabled on all switch ports that have a chance of becoming root or designated ports. It is most effective when enabled in the entire switched network in conjunction with UDLD.

To enable Loop Guard for all point-to-point links on the switch, use the following command:

```
(config)# spanning-tree loopguard default
```

To enable Loop Guard on a specific interface, use the following:

```
(config-if)# spanning-tree guard loop
```

Loop Guard automatically re-enables the port if it starts receiving BPDUs again.

Troubleshooting STP

Some common things to look for when troubleshooting Spanning Tree Protocol include:

- Duplex mismatch—When one side of a link is half-duplex and the other is full-duplex. This causes late collisions and FCS errors.
- Unidirectional link failure—The link is up but data flows only in one direction. It can cause loops.
- Frame corruption—Physical errors on the line cause BPDUs to be lost, and the port incorrectly begins forwarding. This is caused by duplex mismatch, bad cable, or cable too long.
- Resource errors—STP is implemented in software, so a switch with an overloaded CPU or memory might neglect some STP duties.
- Port Fast configuration errors—Connecting a switch to two ports that have Port Fast enabled. This can cause a loop.
- STP tuning errors—Max age or forward delay set too short can cause a loop. A network diameter that is set too low causes BPDUs to be discarded and affects STP convergence.

Identifying a Bridging Loop

Suspect a loop if you see the following:

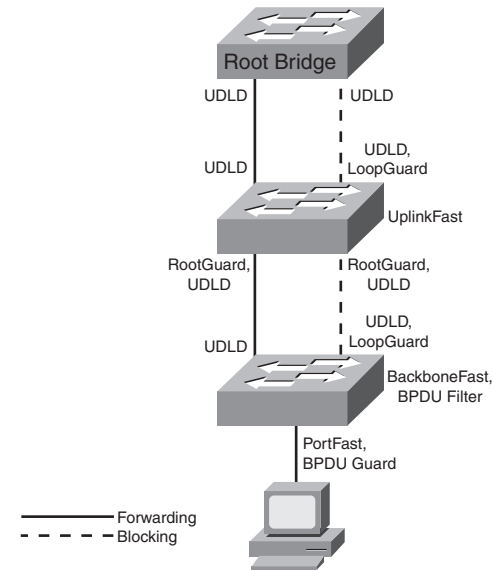
- You capture traffic on a link, and see the same frames multiple times.
- All users in a bridging domain have connectivity problems at the same time.
- There is abnormally high port utilization.

To remedy a loop quickly, shut redundant ports and then enable them one at a time. Some switches allow debugging of STP (not 3550/2950) to help in diagnosing problems.

What to Use Where

Confused by all the acronyms and STP features? Figure 3-3 shows the STP features you might use in your network and where you might use them.

FIGURE 3-3 EXAMPLE SWITCHED TOPOLOGY



InterVLAN Routing

VLANs divide the network into smaller broadcast domains, but also prohibit communication between domains. To enable communication between those groups—without also passing broadcasts—routing is used.

InterVLAN Routing Using Multilayer Switches

Port roles

- Virtual LAN (VLAN) Port—Acts as layer 2 switching port with a VLAN.
- Static VLAN—Use the **switchport** command to identify VLAN.
- Dynamic VLAN—Use VLAN Membership Policy Server (VMPS).
- Trunk Port—Passes multiple VLANs and differentiates by tagging.

Use the **switchport** command to set parameters:

- ISL(Interswitch Link) or 802.1Q
- Switched Virtual Interface (SVI)—Virtual routed port in a VLAN
 - Use to route or fallback bridge between VLANs
 - Default SVI for VLAN 1 automatically created
 - Associate with VLAN using **interface vlan#**

- Routed port—Acts as layer 3 routed port
 - Place in layer 3 mode with no **switchport**
 - Not associated with VLAN
 - Turn on routing using **ip routing**
 - Assign address and enable routing protocols as needed

InterVLAN Routing

Multilayer switches do the following:

- Enable IP routing using **ip routing**
- Create SVI using **interface vlan#**
- Assign an IP address to each interface

A router on a stick attaches the router to the switch using a trunk line (ISL or 802.1Q). Following are features of these:

- Easy to implement
- Use existing equipment
- Much more latency than Multi-layer switching (MLS) solution
- Configure by creating subinterface with **interface fastether-net 1/0.7**
- Associate the VLAN to the interface with command **encapsulation isl 7** or **encapsulation dot1q 7**

CHAPTER 4

INTERVLAN ROUTING

- ISL—No address on main interface
- 802.1Q—Address on main interface for native (untagged) VLAN

Multilayer Switching

This next section walks through the switching process and focuses on order of operations. The order things happen is extremely important for two reasons. First, order of events is good test material. Second, understanding the processing order allows you to evaluate how the various filtering and forwarding mechanisms interact (examples include error checking, access-lists, VLAN access-lists, routing, and QoS).

Understanding the Switching Process

Steps involved in layer 2 forwarding are as follows:

- Input
 1. Receive frame.
 2. Verify frame integrity.
 3. Apply inbound VLAN ACL (Virtual Local Area Network Access List).
 4. Look up destination MAC (Media Address Code).
- Output
 1. Apply outbound VLAN ACL.
 2. Apply outbound QoS ACL.

3. Select output port.
4. Queue on port.
5. Rewrite.
6. Forward.

Steps involved in layer 3 forwarding are as follows:

- Input
 1. Receive frame.
 2. Verify frame integrity.
 3. Apply inbound VLAN ACL.
 4. Look up destination MAC.
- Routing
 1. Input ACL.
 2. Switch if entry cached.
 3. Identify exit interface and next-hop address using routing table.
 4. Output ACL.
- Output
 1. Apply outbound VLAN ACL.
 2. Apply outbound QoS ACL.
 3. Select output port.
 4. Queue on port.

CHAPTER 4

INTERVLAN ROUTING

5. Rewrite source and destination MAC, IP checksum and frame check sequence, and decrement TTL (Time to Live field in the IP header).
6. Forward.

Understanding the Switching Table

Content Addressable Memory (CAM) is used for MAC tables for layer two switching.

- Used for Catalyst 4500 layer 2 forwarding tables
- Used for Catalyst 6500 layer 2 and Netflow forwarding tables
- Contains binary values (0 or 1)
- Match must be exact

In comparison, MLS uses Ternary Content Addressable Memory (TCAM).

- Used for Catalyst 3500/3700, 4500, and 6500 layer 3 switching
- Ternary (3) values (0, 1, or wildcard)
- Entries are in VMR form
 - Value—Pattern to be matched.
 - Mask—Masking bits associated with pattern.
 - Result—Consequences of a match (permit/deny or more complex information).

Understanding Switch Forwarding Architectures

In a Centralized Forwarding model, the CPU controls forwarding decisions:

- Decision made by single table
- Used by 4500 and 6500

With Distributed Forwarding, the forwarding decisions are spread throughout the interface ASICs:

- Decision made at port or module
- Used by 3500/3700 and 6500 with distributed forwarding card
- NetFlow switching
- Decision made cooperatively by Route Processor and MLS
- First packet switched in software, result cached
- Subsequent packets switched in hardware

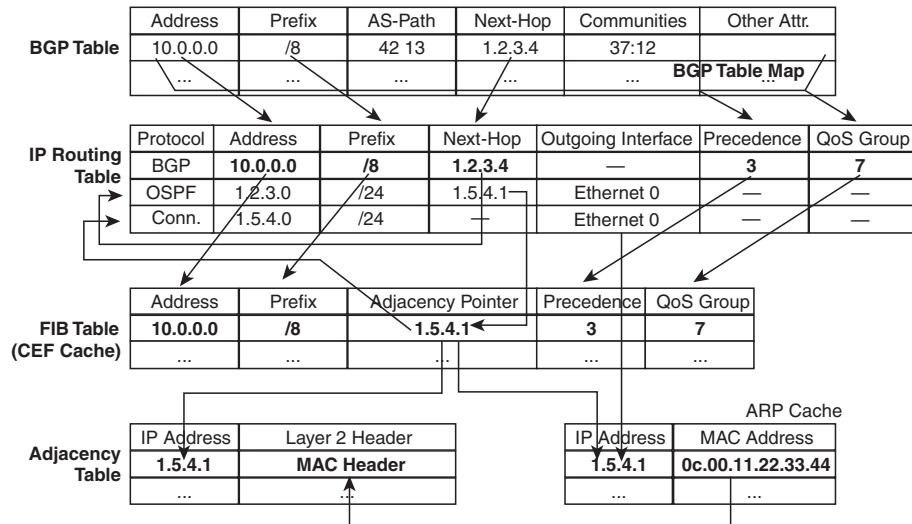
Cisco Express Forwarding (CEF) uses a different kind of memory to facilitate forwarding:

- Uses TCAM
- Topology-based switching (via Forwarding Information Base [FIB])
- Can be centralized or distributed

Multilayer Switching

Multilayer Switching (MLS) is a switch feature that allows the switch to route traffic between VLANs and routed interfaces in a highly optimized and efficient manner. Cisco Express Forwarding (CEF) is an example technology used to facilitate MLS (see Figure 4-1). Cisco Express Forwarding (CEF) does the following:

FIGURE 4-1 CISCO EXPRESS FORWARDING



- Separates control plane hardware from data plane hardware.
- Controls plane runs in software and builds FIB and adjacency table.
- The data plane uses hardware to forward most IP unicast traffic.

- Handles traffic that must be forwarded in software (much slower) and includes:
 - Packets originating from device.
 - Packets with IP header options.
 - Tunneled traffic.
 - 802.3 (IPX) frames.
 - Load sharing traffic.
- FIB is an optimized routing table, stored in TCAM.
- Builds adjacencies from ARP data.
- Eliminates recursive loops.

ARP Throttling

ARP throttling is a tool to limit ARPs into a VLAN. ARPs, you may recall, are sent as broadcast. Once an ARP is sent for a given IP, the switch prevents repetitive ARPs for a short period of time:

- First packet to destination forwarded to Route Processor.
- Subsequent traffic dropped until MAC is resolved.
- It prevents overwhelming the Route Processor (RP) with redundant ARP requests.
- It helps during Denial of Service attacks.
- It is removed when MAC is resolved or in two seconds.

CHAPTER 4

INTERVLAN ROUTING

Configuring and Troubleshooting CEF

By default, CEF is on and supports per destination load sharing.

To disable CEF:

- 4500—Use `(config)#no ip cef`.
- 3500/3700—On each interface, use `(config)#no ip route-cache cef`.
- 6550 with policy feature card, distributed FC, and multilayer switch FC—cannot be disabled.

View CEF information with the following:

```
#show interface fastethernet 2/2 | begin L3
```

View switching statistics with the following:

```
#show interface fastethernet 2/2 | include switched
```

View FIB with the following:

```
#show ip cef
```

View detailed CEF FIB entry with the following:

```
#show ip cef fastethernet 2/2 10.0.0.1 detail
```

Troubleshoot CEF drops with the following:

```
#debug ip cef drops
```

Troubleshoot packets not forwarded by CEF with the following:

```
#debug ip cef receive
```

Troubleshoot CEF events with the following:

```
#debug ip cef events
```

CHAPTER 5

Layer 3 Redundancy

Specifying a default gateway leads to a single point of failure. Proxy Address Resolution Protocol (ARP) is one method for hosts to dynamically discover gateways, but it has issues in a highly-available environment. With Proxy ARP:

- Hosts ARP for all destinations, even remote.
- Router responds with its MAC.
- Problem: Slow failover because ARP entries take minutes to timeout.

Instead of making the host responsible for choosing a new gateway, Layer 3 redundancy protocols allow two or more routers to support a shared MAC address. If the primary router is lost, the backup router assumes control of traffic forwarded to that MAC. This section refers to routers, but includes those Layer 3 switches that can also implement Layer 3 redundancy.

Hot Standby Router Protocol (HSRP)

HSRP is a Cisco proprietary protocol.

With HSRP, two or more devices support a virtual router with a fictitious MAC address and unique IP address. Hosts use this IP address as their default gateway, and the MAC address for the Layer 2 header. The virtual router's MAC address is 0000.0c07.ACxx, where xx is the HSRP group. Multiple groups (virtual routers) are allowed.

The *Active* router forwards traffic. The *Standby* is backup. The standby monitors periodic hellos (multicast to 224.0.0.2, UDP port 1985) to detect a failure of the active router. On failure, the standby device starts answering messages sent to the IP and MAC addresses of the virtual router.

The active router is chosen because it has the highest HSRP priority (default priority is 100). In case of a tie, the router with the highest configured IP address wins the election. A new router with a higher priority does not cause an election unless it is configured to *preempt*—that is, take over from a lower priority router. Configuring a router to preempt also insures that the highest priority router regains its active status if it goes down but then comes back online again.

Interface tracking reduces the active router's priority if a specified circuit is down. This allows the standby router to take over even though the active router is still up.

HSRP States

HSRP devices move between these states:

- Initial—HSRP is not running.
- Learn—The router does not know the virtual IP address and is waiting to hear from the active router.
- Listen—The router knows the IP and MAC of the virtual router, but it is not the active or standby router.
- Speak—Router sends periodic HSRP hellos and participates in the election of the active router.

LAYER 3 REDUNDANCY

- Standby—Router monitors hellos from active router and assumes responsibility if active router fails.
- Active—Router forwards packets on behalf of the virtual router.

Configuring HSRP

To begin configuring HSRP, use the **standby group-number ip virtual-IP-address** command in interface configuration mode. Routers in the same HSRP group must belong to the same subnet/virtual LAN (VLAN.) Give this command under the interface connecting to that subnet or VLAN. For instance, use the following to configure the router as a member of HSRP group 39 with virtual router IP address 10.0.0.1:

```
Router(config-if)#standby 39 ip 10.0.0.1
```

Tune HSRP with four options: Priority, Preempt, Timers, and Interface Tracking.

Manually select the active router by configuring its priority higher than the default of 100:

```
Router(config-if)#standby 39 priority 150
```

Along with configuring priority, configure **preempt** to allow a router to take over if the active router has lower priority, as shown in the following commands. This helps lead to a predictable data path through the network. The second command shown delays preemption until the router or switch has fully booted, and the routing protocol has converged. Time how long it takes to boot and add 50 percent to get the delay value in seconds:

```
Router(config-if)#standby 39 preempt
Router(config-if)#standby 39 preempt delay minimum 90
```

Speed convergence by changing the hello and hold timers. The following sets the hello interval to 2 seconds and the hold time to 7 seconds. They can be set between 1–255 seconds (the default hello is 3 seconds and hold time is 10 seconds):

```
Router(config-if)#standby 39 timers 2 7
```

Tracking an interface can trigger an election if the active router is still up, but a critical interface (such as the one to the Internet) is down. In the following, if serial 1/0/0 is down, the router's HSRP priority is decremented by 100:

```
Router(config-if)#standby 39 track s1/0/0 100
```

Note

The standby router must be configured with the preempt command for it to take control.

Multiple HSRP standby groups can be configured, and the same router can be active for some groups and standby for others by adjusting priorities. You can have a maximum of 255 groups. When using Layer 3 switches, configure the same switch as the primary HSRP router and the Spanning Tree root.

To view the HSRP status, use the **show standby interface interface** command, or **show standby brief**. To monitor HSRP activity, use the **debug standby** command.

Virtual Router Redundancy Protocol (VRRP)

Virtual Router Redundancy Protocol (VRRP) is similar to HSRP, but it is an open standard (RFC 2338). Two or more devices act as a virtual router. With VRRP, however, the IP address used can be either a virtual one or the actual IP address of the primary router.

The VRRP *Master* router forwards traffic. The master is chosen because 1) it owns the real address, or 2) it has the highest priority (default is 100). If a real address is being supported, the owner of real address *must* be master. A *Backup* router takes over if the master fails, and there can be multiple backup routers. They monitor periodic hellos multicast by the master to 224.0.0.18, using UDP port 112, to detect a failure of the master router.

Multiple VRRP groups are allowed, just as with HSRP.

Routers in the same VRRP group must belong to the same subnet/VLAN. To enable VRRP, give this command **vrrp group-number ip virtual-IP-address** under the interface connecting to that subnet or VLAN:

```
Router(config-if)#vrrp 39 ip 10.0.0.1
```

Control the master and backup elections by configuring priority values from 1–255. If a master VRRP router is shutdown, it advertises a priority of 0. This triggers the backup routers to hold an election without waiting for the master's hellos to time out.

```
Router(config-if)#vrrp 39 priority 175
```

VRRP uses the following timers:

- Advertisement, or hello, interval in seconds. Default is 1 second.
- Master down interval. Equals (3 x advertisement interval) plus skew time. Similar to a hold or dead timer.
- Skew time. $(256 - \text{priority}) / 256$. This is meant to ensure that the highest priority backup router becomes master, since higher priority routers have shorter master down intervals.

To change the timers on the master, use the following command because it is the router that advertises the hellos:

```
Router(config-if)#vrrp 39 timers advertise 5
```

To change the timers on the backup routers, use the following command because they hear the hellos from the master:

```
Router(config-if)#vrrp 39 timers learn
```

GLBP

One issue with both HSRP and VRRP is that only the primary router is in use, the others must wait for the primary to fail before they are used. These two protocols use groups to get around that limitation. However, Gateway Load Balancing Protocol (GLBP) allows the simultaneous use of up to four gateways, thus maximizing bandwidth. With GLBP, there is still one virtual IP address. However, each participating router has a virtual MAC address, and different routers' virtual MAC addresses are sent in answer to ARPs sent to the virtual IP address. GLBP can also use groups up to a maximum of 1024 per physical interface.

CHAPTER 5

LAYER 3 REDUNDANCY

The load sharing is done in one of three ways:

- Weighted load balancing—Traffic is balanced proportional to a configured weight.
- Host-dependent load balancing—A given host always uses the same router.
- Round-robin load balancing—Each router MAC is used to respond to ARP requests in turn.

GLBP routers elect an Active Virtual Gateway (AVG). It is the only router to respond to ARPs. It uses this capacity to balance the load among the GLBP routers. The highest priority router is the AVG; the highest configured IP address is used in case of a tie.

The actual router used by a host is its Active Virtual Forwarder (AVF). GLBP group members multicast hellos every 3 seconds to IP address 224.0.0.102, UDP port 3222. If one router goes down, another router answers for its MAC address.

Configure GLBP with the interface command **glbp group-number ip virtual-IP-address**, as shown:

```
Router(config-if)#glbp 39 ip 10.0.0.1
```

To ensure deterministic elections, each router can be configured with a priority. The default priority is 100:

```
Router(config-if)#glbp 39 priority 150
```

Hello and hold (or dead) timers can be configured for each interface with the command **glbp group-number timers [msec] hello-time [msec] hold-time**. Values are in seconds unless the **msec** keyword is used.

GLBP can also track interfaces; if an interface goes down, another router answers for the first router's MAC address.

CHAPTER 6

Using Wireless LANs

Wireless LAN Overview

Devices on a wireless LAN (WLAN) transmit and receive data using radio or infrared signals, sent through an access point (AP). WLANs function similarly to Ethernet LANs with the access point providing connectivity to the rest of the network as would a hub or switch.

WLANs use an Institute of Electrical and Electronics Engineers (IEEE) standard that defines the physical and data link specifications, including the use of Media Access Control (MAC) addresses. The same protocols (such as IP) and applications (such as IPSec) can run over both wired and wireless LANs.

WLANs are local to a building or a campus, use customer-owned equipment, and are not usually required to have radio frequency (RF) licenses.

Service Set Identifiers (SSID) correspond to a VLAN and can be used to segment users. SSIDs can be broadcast by the access point, or statically configured on the client, but the client must have the same SSID as the AP to register with it. SSIDs are case sensitive. Clients associate with access points as follows:

- Step 1.** The client sends a probe request.
- Step 2.** The AP sends a probe response.
- Step 3.** The client initiates an association to an AP. Authentication and any other security information is sent to the AP.

Step 4. The AP accepts the association.

Step 5. The AP adds the client's MAC address to its association table.

Characteristics of Wireless LANs

The following lists some characteristics of wireless LANs, and the data transmitted over wireless networks.

- WLANs use Carrier Sense Multi-Access/Collision Avoidance (CSMA/CA). Wireless data is half-duplex. CSMA/CA uses Request to Send (RTS) and Clear to Send (CTS) messages to avoid collisions.
- WLANs use a different frame type than Ethernet.
- Radio waves have unique potential issues. They are susceptible to interference, multipath distortion, and noise. Their coverage area can be blocked by building features, such as elevators. The signal might reach outside the building and lead to privacy issues.
- WLAN hosts have no physical network connection. They are often mobile and often battery-powered. The wireless network design must accommodate this.
- WLANs must adhere to each country's RF standards.

Clients can roam between APs that are configured with the same SSIDs/VLANs. Layer 2 roaming is done between APs on the same subnet; Layer 3 roaming is done between APs on different subnets.

WLAN Topologies

Use of the Cisco Aironet line of wireless products falls into three categories:

- Client access, which allows mobile users to access the wired LAN resources
- Wireless connections between buildings
- Wireless mesh

Wireless connections can be made in *ad-hoc* mode or *infrastructure* mode. *Ad-hoc* mode (or Independent Basic Service Set [IBSS]) is simply a group of computers talking wirelessly to each other with no access point (AP). It is limited in range and functionality. *Infrastructure* mode's BSS uses one AP to connect clients. The range of the AP's signal, called its microcell, must encompass all clients. The Extended Service Set (ESS) uses multiple APs with overlapping microcells to cover all clients. Microcells should overlap by 10–15 percent for data, and 15–20 percent for voice traffic. Each AP should use a different channel.

Wireless repeaters extend an AP's range. They use the same channel as their AP, they must be configured with the AP's SSID, and they should have 50 percent signal overlap.

Workgroup bridges connect to devices without a wireless network interface card (NIC) to allow them access to the wireless network.

Wireless mesh networks can span large distances because only the edge APs connect to the wired network. The intermediate APs connect wire-

lessly to multiple other APs and act as repeaters for them. Each AP has multiple paths through the wireless network. The Adaptive Wireless Path (AWP) protocol runs between APs to determine the best path to the wired network. APs choose backup paths if the best path fails.

WLAN Standards

WLANs use three unlicensed frequency bands: 900 MHz, 2.4 GHz, and 5 GHz. These bands are all in the Industrial, Scientific, and Medical (ISM) frequency range. Higher frequency bands allow greater bandwidth, but have smaller transmission ranges. Within all bands, the data rate decreases as the client moves away from the AP.

802.11b Standard

802.11b is a widely adopted standard that operates in the 2.4 GHz range and uses Direct Sequence Spread Spectrum (DSSS). It has four data rates: 1, 2, 5.5, and 11 Mbps. 802.11b provides from 11–14 channels, depending on country standards, but only three channels have nonoverlapping frequencies: 1, 6, and 11. Cisco recommends a maximum of 25 users per cell; expect an actual peak throughput of about 6.8 Mbps.

Note

Japan provides a 14 channel, which does not overlap with channel 11 and gives a fourth available nonoverlapping channel.

802.11a Standard

802.11a operates in the 5 GHz range and uses Orthogonal Frequency-Division Multiplexing (OFDM). It has eight data rates: 6, 9, 12, 18, 24, 36, 48, and 54 Mbps. 802.11a provides from 12–23 nonoverlapping channels, depending on country regulations. Portions of the 5 GHz range are allocated to radar, so 802.11a uses Dynamic Frequency Selection (DFS) to check for radar signals and choose a different channel if it detects them. It also uses Transmit Power Control (TMC) to adjust client power, so that they use only enough to stay in contact with the AP. DFS and TMC are part of the 802.11h specification. Cisco recommends a maximum of 15 users per cell; expect an actual peak throughput of about 32 Mbps.

802.11g Standard

802.11g operates in the same 2.4 GHz range as 802.11b and uses the same three nonoverlapping channels: 1, 6, and 11. It can provide higher data rates; however. 802.11g uses DSSS to provide 1, 2, 5.5, and 11 Mbps throughput, which makes it backward compatible with 802.11b. It uses OFDM to provide 6, 9, 12, 18, 24, 36, 48, and 54 Mbps throughput, as does 802.11a.

802.11b/g access points can register both 802.11b and 802.11g clients. Because 802.11b clients do not understand OFDM messages, when 802.11b clients register, the AP implements an RTS/CTS protection mechanism against collisions. When a client wants to talk, it sends an RTS message. The AP must answer with a CTS message before the client is allowed to transmit. This creates overhead for the AP and

causes a drop in overall throughput for all clients. Cisco recommends a maximum of 20 users per cell; expect an actual peak throughput of about 32 Mbps.

Wireless Security

Wireless security methods, listed from weakest to strongest, include:

- **Wired Equivalent Privacy (WEP)**—It uses static keys, weak authentication, and is not scalable.
- **802.1x Extensible Authentication Protocol (EAP)**—Uses RADIUS for authentication, dynamic keys, and stronger encryption. Cisco supports it via Lightweight EAP (LEAP) and Protected EAP (PEAP).
- **Wi-Fi Protected Access (WPA)**—This is a Wi-Fi Alliance standard. Uses Temporal Key Integrity Protocol (TKIP) for encryption, dynamic keys, and 802.1x user authentication. Cisco supports it via Lightweight EAP (LEAP), Protected EAP (PEAP), and Extensible Authentication Protocol-Flexible Authentication via Secure Tunneling (EAP-FAST).
- **WPA2**—The Wi-Fi Alliance’s implementation of the 802.11i standard, which specifies the use of Advanced Encryption Standard (AES) for data encryption and uses 802.1x authentication methods. Can also use TKIP encryption.

WPA/WPA2 Authentication

When a host wanting WLAN access needs to be authenticated in a network using WPA or WPA2, the following steps occur:

- Step 1.** An 802.1x/EAP supplicant on the host contacts the AP (or WLAN controller, if it is a lightweight AP) using 802.1x.
- Step 2.** The AP or WLAN controller uses RADIUS to contact the AAA server, and attempts to authentication the user.
- Step 3.** If the authentication succeeds, all traffic from the client to the AP is encrypted.

Cisco Wireless Network Components

This section is mainly concerned with Cisco products and is quite marketing oriented. Cisco supported two types of wireless solutions: one using autonomous access points, and one using lightweight (or “dumb”) access points in combination with WLAN controllers. The wired network infrastructure is the same for both types: switches and routers.

Cisco Unified Wireless Network

The Cisco Unified Wireless Network concept has five components that work together to create a complete network, from client devices to network infrastructure, to network applications. Cisco has equipment appropriate to each component. Table 6-1 lists components and equipment.

TABLE 6-1 Cisco Unified Wireless Network Components

Component	Description and Devices
Client Devices	Cisco Aironet client, and Cisco compatible third-party vendor clients.
Mobility Platform	Aironet APs and bridges, using LWAPP.
Network Unification	Leverages existing wired network. 2000- and 4400-series WLAN controllers and switch and router modules.
World-Class Network Management	Visualize and secure the WLAN. WCS for location tracking, RF management, wireless IPS, and WLC management.
Unified Advanced Services	Applications such as wireless IP phones, location appliances, and RF firewalls.

You should review the following link for more information on Cisco wireless controllers and access points before you take the exam: http://www.cisco.com/en/US/products/hw/wireless/products_category_buyers_guide.html Wireless Clients.

Cisco has a wireless NIC that can be installed on Windows 2000 and Windows XP systems. It comes with some utilities: Aironet Desktop Utility (ADU), Aironet Client Monitor (ACM), and Aironet Client Administration Utility (ACAU). Cisco recommends using the ADU and ACM utilities to control your wireless card, rather than the built-in Windows controls to get the increased functionality Cisco provides. The Cisco ACAU allows loading and configuration of the Cisco client software over the network, using encrypted files. There is also an Aironet Site Survey Utility to scan for APs and get information about them.

Cisco wireless IP phones have the same features as Cisco wired IP phones and can use LEAP for authentication.

The Cisco Compatible Extensions Program tests other vendors' devices for compatibility with Cisco wireless products. Using products certified by this program ensures full functionality of Cisco enhancements and proprietary extensions. A list of these products can be found at www.cisco.com/go/ciscocompatible/wireless.

Autonomous APs

Autonomous APs run Cisco IOS, are programmed individually, and act independently. They can be centrally managed with the CiscoWorks Wireless LAN Solution Engine (WLSE) and can use Cisco Secure Access Control Server (ACS) for RADIUS and TACAS+ authentication. Redundancy consists of multiple APs.

Lightweight Access Points

Lightweight APs divide the 802.11 processing between the AP and a Cisco Wireless LAN Controller (WLC). This is sometimes called “split MAC,” because they split the functions of the MAC layer—Layer 2. Their management components also include the Wireless Control System (WCS) and a location-tracking appliance. Redundancy consists of multiple WLCs. The AP handles real-time processes, and the WLC handles processes such as:

- Authentication
- Client association/mobility management

- Security management
- QoS policies
- VLAN tagging
- Forwarding of user traffic

The Lightweight Access Point Protocol (LWAPP) supports the split MAC function in traffic between a lightweight AP and its controller. LWAPP uses AES-encrypted control messages and encapsulates, but does not encrypt, data traffic. LWAPP operates at Layer 2, and also at Layer 3 over UDP. (However, Layer 2 operation has been deprecated by Cisco.) The controller can be either in the same broadcast domain and IP subnet or in a different broadcast domain and IP subnets for Layer 3 operation. The AP follows this process to discover its controller:

- Step 1.** The AP requests a DHCP address. The DHCP response includes the management IP address of one or more WLCs.
- Step 2.** The AP sends an LWAPP Discovery Request message to each WLC.
- Step 3.** The WLCs respond with an LWAPP Discovery Response that includes the number of APs currently associated to it.
- Step 4.** The AP sends a Join Request to the WLC with the fewest APs associated to it.
- Step 5.** The WLC responds with a Join Response message, the AP and the controller mutually authenticate each other and

derive encryption keys to be used with future control messages. The WLC then configures the AP with settings, such as SSIDs, channels, security settings, and 802.11 parameters.

The Cisco Aironet 2000 series WLC can handle up to six APs; thus, it is sized for small- to medium-sized operations.

The Cisco Aironet 4400 series WLC supports medium to large facilities with the 4402 handling up to 50 APs, and the 4404 handling up to 100 APs.

Wireless LAN Antennas

Several concepts are important in understanding wireless antennas:

- **Gain**—The energy an antenna adds to the RF signal.
- **Directionality**—How the radio coverage is distributed.
- **Polarization**—The physical orientation the RF element. Cisco Aironet antennas use vertical polarization.
- **Multipath Distortion**—Receiving both direct and reflected signals arriving from different directions.
- **Effective Isotropic Radiated Power (EIRP)**—The AP radio's effective transmission power. Includes gain from the antenna and loss from the antenna's cable.

Gain

Cisco measures gain in dBi, which stands for *decibel isotropic* and is a measure of decibels relative to an isotropic source in free space. A *decibel* is the ratio between two signal levels. An *isotropic* antenna is a theoretical one in which the signal spreads out evenly in all directions from one point. Thus, dBi is the ratio of an antenna's signal to that of an isotropic antenna.

Directionality

Omnidirectional antennas have signals that theoretically extend in all directions, both vertically and horizontally. When gain is increased, the signal expands horizontally, but decreases vertically. One omnidirectional example is the dipole “Rubber Duck” antenna.

Directional antennas aim their signal in a specific direction. Signals can spread fairly wide in one direction or can be narrowly focused. Some examples include the Diversity Patch Wall Mount Antenna, Yagi, and dish antennas.

Multipath Distortion

Because radio waves are transmitted in many directions, not all go in a straight line to every client's antenna. Some bounce off walls or other objects and arrive at the client in varying intervals. Thus, the client receives several copies of the same RF signal, which can cause degraded data quality. This is *multipath distortion*, or *multipath interference*. Diversity systems try to minimize this by using two antennas; you might try moving antennas or changing the frequency if this is a

USING WIRELESS LANS

problem in your facility. OFDM uses multiple frequencies operating together to increase performance in multipath situations.

EIRP

EIRP is the actual power of the signal that comes from the antenna, measured in Decibel Milliwatts (dBm). (0 dBm equals 1 milliwatt of power.) EIRP is calculated by taking the transmitter power, subtracting the amount of signal lost traversing the cable between the transmitter and antenna, and adding the antenna's gain. This can be expressed:

$$\text{EIRP} = (\text{power} - \text{cable loss}) + \text{antenna gain.}$$

Different countries have different rules about the amount of EIRP allowed. For instance, the maximum in the United States is 36 dBm. To minimize signal loss, use the shortest low-loss cable possible. Wider cables conserve more signal but are also more expensive.

Power over Ethernet (PoE) Switches

Access points can receive their power over Ethernet cables from Power over Ethernet (PoE) switches, routers with PoE switch modules, or midspan power injectors, thus alleviating the need for electrical outlets near them. APs require up to 15W of power, so plan your power budget accordingly. Two power standards are the Cisco Prestandard PoE and the IEEE's 802.3af standard. Both have a method for sensing that a powered device is connected to the port. 802.3af specifies a method for determining the amount of power needed by the device. Cisco devices,

when connected to Cisco switches, can additionally use CDP to send that information. Power can be supplied over the data pairs—1, 2, 3, and 6—or over the unused pairs of 4, 5, 7, and 8.

Cisco PoE switches are configured by default to automatically detect and provide power. To disable this function, or to re-enable it, use the interface command **power inline {never | auto}**. To view interfaces and the power allotted to each, use **show power inline [interface]**.

Configuring Wireless LAN Devices

Autonomous APs must be configured individually, while the WLC provides configuration to lightweight APs. WLAN clients must also be configured; this process varies depending on the client software used.

Configuring Autonomous Access Points

Autonomous APs can be configured in one of three ways:

- IOS Command Line—Either via Telnet or the console port.
- Web browser—This is the Cisco preferred way.
- CiscoWorks WLSE—For centralized configuration control.

The AP must already have an IP address to use any of these except the console port. It attempts to obtain one via DHCP by default. This link

USING WIRELESS LANS

has directions and screen shots for both the command line and web browser configuration:

http://www.cisco.com/en/US/products/ps6087/products_installation_and_configuration_guides_list.html.

Aironet 1100, 1200, and 1300 series APs perform various functions:

- Wireless AP
- Root bridge
- Nonroot bridge
- Repeater
- Scanner
- Workgroup bridge

Configuring a WLAN Controller

Cisco lightweight APs receive their configuration from the Wireless LAN Controller, which must be configured first. Initial configuration of the lightweight WLC can be done via command line using the console port or via web browser using the service port. Subsequent configuration can be done via:

- IOS Command Line—Either by Telnet, SSH, or the console port.
- Web browser—Using the WLC's IP address and Internet Explorer.
- Cisco Wireless Control System—For centralized configuration control.

You need to configure the WLC with information such as VLANs, SSIDs, and security policies. It downloads a configuration to its associated APs, and you can also configure, monitor, or reset individual APs through the web browser of the WLC. Review the material at this link for screen shots and WLC configuration information:

http://www.cisco.com/en/US/products/ps6366/products_configuration_guide_book09186a00806b0077.html.

WLCs use several different types of physical and logical interfaces that are described in Table 6-2.

TABLE 6-2 Wireless LAN Controller Interfaces

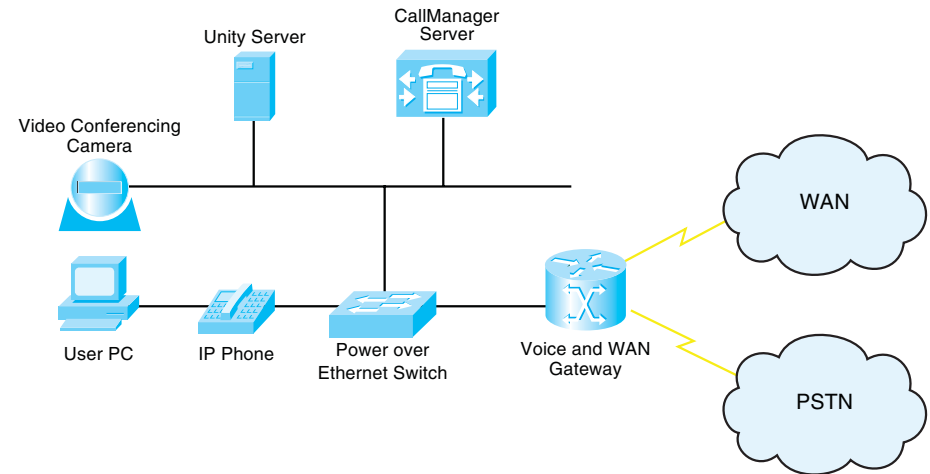
Interface Type	Description
Service Port	Used for out of band management and initial setup. Must be a unique subnet. Not present on the 2006 WLC.
Management Interface	Used by the APs to find their WLC and associate with it. One per WLC.
AP-Manager Interface	Used for LWAPP traffic between controller and APs. Can have multiple AP-Manager interfaces.
Virtual Interface	IP address used for mobility group when implementing layer 3 roaming.
User Interface	Used to carry data traffic from users. One per VLAN.

VoIP in a Campus Network

Many companies are integrating Voice over IP (VoIP) into their networks. Figure 7-1 shows some components of a VoIP system, which can include the following:

- IP phones—Provide voice and applications to the user.
- Voice gateways—Translates between PSTN and IP calls and provides backup to the Cisco CallManager (IP PBX, or Call Agent).
- Gatekeepers—An optional component that can do call admission control, allocate bandwidth for calls, and resolve phone numbers into IP addresses.
- Cisco CallManager—Serves as an IP PBX. Registers phones, controls calls.
- Video conferencing unit—Allows voice and video in the same phone call.
- Multipoint control unit—Allows multiple participants to join an audio and/or video conference call.
- Application server—Provides services such as Unity voice mail.

FIGURE 7-1 SOME COMPONENTS OF A VOIP SYSTEM



Voice and data have different network requirements. Although TCP data adjusts to dropped packets, packet loss is one of the biggest enemies of voice transmissions and is often caused by jitter and congestion. Jitter (variable delay) causes buffer over- and under-runs. Congestion at the interface can be caused by traffic from a fast port being switched to exit out a slower port, which causes the transmit buffer to be overrun.

VoIP traffic consists of two types: voice bearer and call control signaling. Voice bearer traffic is carried over the UDP-based Real Time Protocol (RTP). Call control uses one of several different protocols to communicate between the phone and CallManager and between the CallManager and the voice gateways.

CHAPTER 7

VOIP IN A CAMPUS NETWORK

Preparing the Network for VoIP

When adding voice or video to an existing network, you should examine several things in advance to provide the high level of availability users expect in their phone system:

- What features are needed?—Power for IP phones, voice VLANs on the switches, network redundancy for high availability, security for voice calls, and Quality of Service (QoS) settings.
- The physical plant—Cabling at least CAT-5.
- Electrical power for the IP phones—Use either inline power from Catalyst switch or power patch panel. Need uninterruptible power supply (UPS) with auto-restart, monitoring, and 4-hour response contract. May need generator backup. Maintain correct operating temperatures.
- Bandwidth—Commit no more than 75 percent of bandwidth. Consider all types of traffic—voice, video, and data. Have more than enough bandwidth if possible. Include both voice and call-control traffic in your planning.
- Network management—Need to monitor and proactively manage the network so that it does not go down.

Network and Bandwidth Considerations

The network requirements for VoIP include:

- Maximum delay of 150–200 ms (one-way)
- No more than 1 percent packet loss
- Maximum average jitter of 30 ms
- Bandwidth of 21–106 kbps per call, plus about 150 bps per phone for control traffic

A formula to use when calculating bandwidth needed for voice calls is as follows:

(Packet payload + all headers) * Packet rate per second

Auxiliary (or Voice) VLANs

Cisco switches can be configured to dynamically place IP telephones into a VLAN separate from the data VLANs. They can do this even when the phone and PC are physically connected to the same switch port. This is called an auxiliary VLAN or a voice VLAN. Voice VLANs allow phones to be dynamically placed in a separate IP subnet from hosts, to have QoS (using 802.1Q/p headers) and security policies applied, and makes troubleshooting easier.

QoS for VoIP

QoS gives special treatment to certain traffic at the expense of others. Using QoS in the network has several advantages:

- Prioritizes access to resources, so that critical traffic can be served.
- Allows good management of network resources.
- Allows service to be tailored to network needs.
- Allows mission-critical applications to share the network with other data.

People sometimes think that there is no need for QoS strategies in a LAN. However, switch ports can experience congestion because of port speed mismatches, many people trying to access the switch backbone, and many people trying to send traffic to the same switch port (such as a server port).

QoS Actions

Three QoS strategies are commonly implemented on interfaces where traffic enters the switch:

- Classification—Distinguishing one type of traffic from another. After traffic is classified, other actions can be performed on it. Some classification methods include access lists, ingress interface, and NBAR.
- Marking—At layer 2, placing 802.1p class of service (CoS) value within the 802.1Q tag. At layer 3, setting IP Precedence or Differentiated Services Code Point (DSCP) values on the classified traffic.

- Policing—Determining whether or not a specific type of traffic is within preset bandwidth levels. If so, it is usually allowed and might be marked. If not, the traffic is typically marked or dropped. CAR and class-based policing are examples of policing techniques.

Other QoS techniques are typically used on outbound interfaces:

- Traffic shaping and conditioning—Attempts to send traffic out in a steady stream at a specified rate. Buffers traffic that goes above that rate and sends it when there is less traffic on the line.
- Queuing—After traffic is classified and marked, one way it can be given special treatment is to be put into different queues on the interface to be sent out at different rates and times. Some examples include priority queuing, weighted fair queuing, and custom queuing. The default queuing method for a switch port is FIFO.
- Dropping—Normally interface queues accept packets until they are full and then drop everything after that. You can implement prioritized dropping, so that less important packets are dropped before more important ones—such as with Weighted Random Early Detection (WRED).

DSCP Values

Differentiated services provide levels of service based on the value of certain bits in the IP or ISL header or the 802.1Q tag. Each hop along the way must be configured to treat the marked traffic the way you want—this is called per-hop behavior (PHB).

CHAPTER 7

VOIP IN A CAMPUS NETWORK

In the Layer 3 IP header, you use the 8-bit ToS field. You can set either IP Precedence using the top 3 bits or Differentiated Services Code Points (DSCP) using the top 6 bits of the field. The bottom 2 bits are set aside for congestion notification. The default DSCP value is zero, which corresponds to best-effort delivery.

The six DSCP bits can be broken down into two sections: The first 3 bits define the DiffServ Assured Forwarding (AF) class, and the next 2 bits define the drop probability within that class. The sixth bit is 0 and unused. AF classes 1–4 are defined, and within each class, 1 is low drop probability, 2 is medium, and 3 is high (meaning that traffic is more likely to get dropped if there is congestion). These are shown in Table 7-1. Each hop still needs to be configured for how to treat each AF class.

TABLE 7-1 DSCP Assured Forwarding Values

	Low Drop	Medium Drop	High Drop
Class 1	AF11	AF12	AF13
Class 2	AF21	AF22	AF23
Class 3	AF31	AF32	AF33
Class 4	AF41	AF42	AF43

Voice bearer traffic uses an Expedited Forwarding value of DSCP 46 to give it higher priority within the network.

Trust Boundaries

When IP traffic comes in already marked, the switch has some options about how to handle it. It can:

- Trust the DSCP value in the incoming packet, if present.
- Trust the IP Precedence value in the incoming packet, if present.
- Trust the CoS value in the incoming frame, if present.
- Classify the traffic based on an IP access control list or a MAC address access control list.

Mark traffic for QoS as close to the source as possible. If the source is an IP telephone, it can mark its own traffic. If not, the building access module switch can do the marking. If those are not under your control, you might need to mark at the distribution layer. Classifying and marking slows traffic flow, so do not do it at the core. All devices along the path should then be configured to trust the marking and provide a level of service based on it. The place where trusted marking is done is called the *trust boundary*.

Configuring VoIP Support on a Switch

Manual Configuration

To associate a voice VLAN with a switch port, use the following:

```
Switch(config-if)#switchport voice vlan vlan-ID
```

To configure an IOS switch to trust the markings on traffic entering an interface, use the following:

```
Switch(config-if)#mls qos trust {dscp | cos}
```

To configure the switch to trust the traffic markings only if a Cisco phone is connected, use the following:

```
Switch(config-if)#mls qos trust device cisco-phone
```

To set a COS value for frames coming from a PC attached to the phone, use the following:

```
Switch(config-if)#switchport priority extend cos cos-value
```

To verify the interface parameters, use the following:

```
Switch(config-if)#show interfaces interface switchport
```

To verify the QoS parameters on an interface, use the following:

```
Switch(config-if)#show mls qos interface interface
```

Using AutoQoS

When AutoQoS is enabled, the switch configures its interfaces based on a best-practices template. AutoQoS has the following benefits:

- Automatic discovery and classification of network applications.
- Creates QoS policies for those applications.
- Configures the switch to support Cisco IP phones as well as network applications. Manual configuration can be done afterward, also.
- Sets up SNMP traps for network reporting.
- Configures consistently across your network when used on all routers and switches.

CDP must be enabled for AutoQoS to function properly with Cisco IP phones.

AutoQoS commands for switches running the Catalyst OS are listed in Table 7-2.

TABLE 7-2 AutoQoS Commands for Catalyst OS

Command	Description
set qos autoqos	Globally enables AutoQoS on the switch.
set port qos <i>mod/port</i> autoqos trust [cos dscp]	Configures the port to trust either the COS or DSCP markings of all traffic coming in the port.
set port qos <i>mod/port</i> autoqos voip [ciscosoftphone ciscoipphone] [trust]	Configures the port to trust traffic markings only if a Cisco phone or a computer with a Cisco softphone is connected to the port. Requires that CDP be enabled.

CHAPTER 7

VOIP IN A CAMPUS NETWORK

AutoQoS commands for switches running Native IOS are shown in Table 7-3.

TABLE 7-3 AutoQoS Commands for IOS

Command	Description
(config-if)# auto qos voip trust	Configures the port to trust the COS on all traffic entering the port.
(config-if)# auto qos voip cisco-phone	Configures the port to trust traffic markings only if a Cisco phone is connected to the port. Requires that CDP be enabled.
# show auto qos [interface interface]	Shows the AutoQoS configuration. Does not show any manual QoS configuration—use show run to see that.

CHAPTER 8

Campus Network Security

Attention has traditionally been paid to network perimeter security, such as firewall, and to mitigating Layer 3 attacks. However, networks must be protected against Layer 2 attacks, also. These are launched from devices inside the network by either a rogue device or a legitimate device that has been compromised. Rogue devices might be placed maliciously or might just be connected to an access switch by an employee wanting more switch port or wireless access. They include:

- Wireless routers or hubs
- Access switches
- Hubs

A switch might become the Spanning Tree root bridge, and disrupt user traffic. Use **root guard** and **bpdu guard** commands to prevent this. (Spanning tree security is discussed later in this chapter.)

There are four typical types of attacks against a switched network:

- MAC-based attacks, such as MAC address flooding
- VLAN-based attacks, such as VLAN hopping and attacks against devices on the same VLAN
- Spoofing attacks, such as DHCP spoofing, MAC spoofing, Address Resolution Protocol (ARP) spoofing, and Spanning Tree attacks
- Attacks against the switch, such as Cisco Discovery Protocol (CDP) manipulation, Telnet attacks, and Secure Shell (SSH) attacks

MAC Address Flooding

In a MAC address flooding attack, the attacker fills the switch's Content Addressable Memory (CAM) table with invalid MAC addresses. After the table is full, all traffic with an address not in the table is flooded out all interfaces. This has two bad effects—more traffic on the LAN and more work for the switch. Additionally, the intruder's traffic is also flooded, so they have access to more ports than they would normally have. After the attack stops, CAM entries age out and life returns to normal. However, meanwhile the attacker might have captured a significant amount of data.

Port security and port-based authentication can help mitigate MAC address attacks.

Port Security

Port security limits the number of MAC addresses allowed per port and can also limit which MAC addresses are allowed. Allowed MAC addresses can be manually configured or the switch can sticky learn them. Table 8-1 lists port security commands; these are given at the interface.

TABLE 8-1 Port Security Commands

Command	Description
switchport port-security	Enables port security on that interface.
switchport port-security maximum value	Specifies the max MAC addresses allowed on this port. Default is 1.

CAMPUS NETWORK SECURITY

TABLE 8-1 Port Security Commands

Command	Description
switchport port-security violation {shutdown restrict protect}	Configures the action to be taken when the maximum number is reached and a MAC address not associated with the port attempts to use the port, or when a station whose MAC address is associated with a different port attempt to access this port. Default is shutdown .
switchport port-security mac-address mac-address	Statically associates a specific MAC address with a port.
switchport port-security mac-address sticky	Enables the switch port to dynamically learn secure MAC addresses. MAC addresses learned through that port, up to the maximum number, if a maximum is configured, are treated as secure MAC addresses.
show port security [interface interface address]	Verifies port security actions.

Port-Based Authentication

802.1x authentication requires a computer (called a client) to be authenticated before it is allowed access to the LAN. This can be combined with port security to allow only authenticated clients with specified MAC addresses to access a port. When a computer connects to a switch port configured for 802.1x authentication, the following steps occur:

- Step 1.** The port is in the *unauthorized* state, allowing only 802.1x EAP over LAN (EAPOL) traffic.
- Step 2.** The client connects to the port. The switch either requests authentication or the client sends an EAPOL frame to begin authentication.
- Step 3.** The switch relays authentication information between the client and a RADIUS server that acts in proxy for the client.
- Step 4.** If authentication succeeds, the port transitions to the *authorized* state, and normal LAN traffic is allowed through it.

Table 8-2 shows commands to configure 802.1x authentication on a switch.

TABLE 8-2 Configuring 802.1x Port Authentication

Command	Description
(config)# aaa new-model	Enables AAA on the switch.
(config)# aaa authentication dot1x default group radius	Creates a AAA method list that says to use 802.1x authentication by default, using a RADIUS server (configured separately).
(config)# dot1x system-auth-control	Globally enabled 802.1x authentication on the switch.
(config-if)# dot1x port-control auto	Enables 802.1x authentication on an interface of the switch.
show dot1x	Verifies 802.1x authentication.

VLAN-Based Attacks

VLAN-based attacks include VLAN hopping, in which a station is able to access a VLAN other than its own. This can be done with switch spoofing or with 802.1Q double-tagging.

Switch Spoofing

Switch spoofing involves a station configured to negotiate a trunk link between itself and the switch. By default, switches dynamically negotiate trunking status using Dynamic Trunking Protocol (DTP). If a computer is able to use DTP to establish a trunk link to the switch, it will receive all traffic bound for VLANs allowed on that trunk. By default, all VLANs are allowed on a trunk.


You can mitigate this by turning off DTP on all ports that should not become trunks, such as most access ports, using the interface command **switchport nonegotiate**. If the port should be an access port, configure it as such with the interface command **switchport mode access**. Additionally, shut down all unused ports and assign them to an unused VLAN. The commands to do this are:

```
Switch(config)#interface interface
Switch(config-if)#switchport mode access
Switch(config-if)#switchport access vlan vlan
Switch(config-if)#shutdown
```

802.1Q Double-Tagging

A double-tagging attack is possible with 802.1Q trunking because it does not tag frames from the native VLAN. In this attack, the attacking computer sets up a trunk port between itself and the switch, then generates frames with two 802.1Q tags. The first tag matches the native VLAN of the trunk port, and the second matches the VLAN of a host it wants to attack, as shown in Figure 8-1.

FIGURE 8-1 VLAN HOPPING BY 802.1Q DOUBLE-TAGGING



Switch A removes the first tag for VLAN 100, because it matches the native VLAN for that link. It forwards the frame out all links with the same native VLAN, including its link to Switch B. Switch B sees the frame come in with an 802.1Q tag for VLAN 200, so it forwards it out the VLAN 200 link to the victim computer.

To mitigate this type of attack, use the same strategies used for switch spoofing. You can also use VLAN access control lists, called *VACLs*, or implement Private VLANs.

VACLs

Cisco switches support of various kinds of ACLs:

- Traditional Router ACL (RACL)
- QoS ACL
- VACL

VLAN access control lists (VACLs) are similar to route-maps in that they are composed of statements that contain match and set conditions. In a VACL, the “set” conditions are called “actions.” Actions include **forward**, **drop**, and **redirect**. Like route-maps, VACL statements are numbered for ordering. After configuration, VACLs are applied to traffic to specified VLANs.

The following is a sample VACL that instructs the switch to drop traffic matching ACL 101 (not shown), and forward all other traffic:

```
Switch(config)#vlan access-map Drop101 5
Switch(config-access-map)#match ip address 101
Switch(config-access-map)#action drop
Switch(config-access-map)#vlan access-map Drop101 10
Switch(config-access-map)#action forward
!
Switch(config)#vlan filter Drop101 vlan_list 10
```

To view VACL settings, use the commands **show vlan access-map** *vacl_name* or **show vlan filter access-map** *vacl_name*.

Private VLANs

Private VLANs (PVLANS) allow service providers to isolate customers into separate multi-access domains. Using a VLAN for each customer is not scalable, because a switch’s maximum VLANs would limit the number of customers an ISP can have. Each VLAN requires a separate IP subnet, which could also be a limiting factor.

PVLANS divide a VLAN into secondary VLANs, letting you isolate a set of ports from other ports within the same VLAN. There are two types of secondary VLANs:

- Community VLANs—Ports can communicate with other ports in the same community VLAN.
- Isolated VLANs—Ports cannot communicate with each other.

Ports within a private VLAN can be one of three types:

- Community—Communicates with other community ports and with promiscuous ports.
- Isolated—Communicates only with promiscuous ports.
- Promiscuous—Communicates with all ports.

Table 8-3 shows the commands to configure a primary private VLAN, secondary PVLANS, and their associated ports.

CHAPTER 8

CAMPUS NETWORK SECURITY

TABLE 8-3 Configuring Private VLANs

Command	Description
vlan <i>vlan-id</i>	Enters VLAN configuration mode.
private-vlan { community isolated primary }	Configures the VLAN as a private VLAN and specifies the type. Repeat this command to configure all primary and secondary VLANs.
vlan <i>primary-vlan-id</i>	Enters configuration mode for the primary VLAN.
private-vlan association <i>secondary_vlan_list</i>	Associates secondary VLANs with the primary one. Separate the secondary VLAN numbers with a comma, no spaces.
switchport mode private-vlan { host promiscuous }	Configures a port as either a host port (for community or isolated) or a promiscuous port.
switchport private-vlan host-association <i>primary_vlan_ID secondary_vlan_ID</i>	Associates a host port with its primary and secondary PVLANS.
private-vlan mapping <i>primary_vlan_ID secondary_vlan_list</i>	Associates a promiscuous port with its primary and secondary PVLANS.
show interfaces <i>interface</i> switchport	Verifies the VLAN configuration.
show interfaces private-vlan mapping	Verify the private VLAN configuration.

Spoof Attacks

Spoof attacks include DHCP spoofing, MAC address spoofing, and ARP spoofing.

DHCP Spoofing

A DHCP spoofing attacker listens for DHCP requests and answers them, giving its IP address as the client default gateway. The attacker then becomes a “man-in-the-middle” as all off-net traffic flows through it.

DHCP snooping can prevent DHCP spoofing attacks. When DHCP snooping is enabled, only ports that uplink to an authorized DHCP server are trusted, and allowed to send all types of DHCP messages. All other ports on the switch are untrusted and can send only DHCP requests. If a DHCP response is seen on an untrusted port, the port is shut down. The switch can also be configured to send information, such as port ID, using DHCP option 82.

Note

DHCP snooping configuration is user impacting, because the switch drops all DHCP requests until the ports are configured. You should do this during off hours or during a maintenance window.

CAMPUS NETWORK SECURITY

Configure DHCP snooping with the following commands, either globally or for a particular VLAN. Configure only individual ports that uplink to DHCP servers as trusted ports.

```
Switch(config)#ip dhcp snooping
Switch(config)#ip dhcp snooping information option
Switch(config)#ip dhcp snooping vlan number number
Switch(config-if)#ip dhcp snooping trust
Switch#show ip dhcp snooping
```

To extend the protection further, IP Source Guard tracks the IP addresses of the host connected to each port and prevents traffic sourced from another IP address from entering that port. The tracking can be done based on just an IP address or on both IP and MAC addresses.

Enable IP Source Guard for both IP and MAC addresses on host access interfaces with the command **ip verify source vlan dhcpsnooping port-security**.

ARP Spoofing

In an ARP spoofing attack, the attacker sends out gratuitous (unsolicited) ARP messages giving the IP address of the local default gateway, with its own MAC address as the layer 2 address. Local devices overwrite their existing correct ARP information with the incorrect one, and, thus, they forward off-net traffic to the attacker (it becomes a “man-in-the-middle”). If the attacker then forwards it on to the legitimate router, this type of attack might go undetected by the users.

Dynamic ARP Inspection (DAI) can work with DHCP spoofing to stop ARP spoofing. DAI defines trusted and untrusted interfaces. It intercepts ARP messages on untrusted ports, and checks them against the IP address/MAC address bindings in the DHCP snooping database. They must match for the switch to forward the traffic. Access ports should be configured as untrusted, and ports that connect to other switches or to a router should be trusted.


Enable DAI on a VLAN, or multiple VLANs, and configure trusted interfaces. You can optionally configure a rate limit, or configure which addresses DAI matches against (the default is IP and MAC address). The basic commands are:

```
Switch(config)#ip arp inspection vlan vlan_id
Switch(config-if)#ip arp inspection trust
```

Securing Spanning Tree

Spanning Tree tuning can help prevent a rogue device from becoming root bridge or otherwise disrupting your user traffic. There are several tools at your disposal—Figure 8-2 shows where each could be used in a switched network.

FIGURE 8-2 SECURING SPANNING TREE



BPDUGuard

BPDUGuard prevents loops if another switch is attached to a Portfast port. When BPDUGuard is enabled on an interface, it is put into an error-disabled state (basically, it is shut down) if a BPDU is received on the interface. It can be enabled at either global configuration mode—in

which case it affects all Portfast interfaces—or at interface mode. Portfast does not have to be enabled for it to be configured at a specific interface.

```
Switch(config)#spanning-tree portfast bpduguard default
Switch(config-if)#spanning-tree bpduguard enable
```

BPDU Filtering

BPDU filtering is another way of preventing loops in the network. It also can be enabled either globally or at the interface, and it functions differently at each. In global configuration, if a Portfast interface receives any BPDUs, it is taken out of Portfast status. At interface configuration mode, it prevents the port from sending or receiving BPDUs. The commands are:

```
Switch(config)#spanning-tree portfast bpdupfilter default
Switch(config-if)#spanning-tree bpdupfilter enable
```

Root Guard

Root Guard is meant to prevent the wrong switch from becoming the spanning-tree root. It is enabled on ports other than the root port and on switches other than the root. If a Root Guard port receives a BPDU that causes it to become a root port, the port is put into a “root-inconsistent” state and does not pass traffic through it. If the port stops receiving these BPDUs, it automatically re-enables itself.

```
Switch(config-if)#spanning-tree guard root
Switch#show spanning-tree inconsistentports
```


Prevent Spanning Tree Loops

A switch notices when a physical connection is broken by the absence of Layer 1 electrical keepalives (Ethernet calls this a link beat).

However, sometimes a cable is intact enough to maintain keepalives, but not to pass data in both directions. This is a unidirectional link.

Unidirectional Link Detection (UDLD)

UDLD detects a unidirectional link by sending periodic hellos out the interface. It also uses probes, which must be acknowledged by the device on the other end of the link. UDLD operates at Layer 2. The port is shut down if a unidirectional link is found.

To enable UDLD on all fiber-optic interfaces, use this command:

```
Switch(config)#udld enable
```

Although this command is given at global configuration mode, it applies only to fiber ports. To enable UDLD on nonfiber ports, give the same command at interface config mode.

To disable UDLD on a specific fiber port, use this command:

```
Switch(config-if)#udld disable
```

To disable UDLD on a specific nonfiber port, use this command:

```
Switch(config-if)#no udld enable
```

To re-enable all interfaces shut by UDLD, use the following:

```
Switch#udld reset
```

To verify UDLD status, use the following:

```
Switch#show udld interface
```

Loop Guard

Loop Guard prevents loops that might develop if a port that should be blocking inadvertently transitions to the forwarding state. This can happen if the port stops receiving BPDUs (perhaps because of a unidirectional link or a software/configuration problem in its neighbor switch). When one of the ports in a physically redundant topology stops receiving BPDUs, the STP conceives the topology as loop-free. Eventually, the blocking port becomes designated and moves to a forwarding state, thus creating a loop. With Loop Guard enabled, an additional check is made.

If no BPDUs are received on a blocked port for a specific length of time, Loop Guard puts that port into loop inconsistent blocking state, rather than transitioning to a forwarding state. Loop Guard should be enabled on all switch ports that have a chance of becoming root or designated ports. It is most effective when enabled in the entire switched network, in conjunction with UDLD.

To enable Loop Guard for all point-to-point links on the switch, use the following command:

```
Switch(config)#spanning-tree loopguard default
```

To enable Loop Guard on a specific interface, use:

```
Switch(config-if)#spanning-tree guard loop
```

CHAPTER 8

CAMPUS NETWORK SECURITY

Loop Guard automatically re-enables the port if it starts receiving BPDUs once again.

Securing Your Switch

Here are some basic security suggestions for network devices:

- Use passwords that are not susceptible to a dictionary attack. Add numbers or substitute numbers and symbols for letters.
- Limit Telnet access using access lists.
- Use SSH instead of Telnet.
- Physically secure access to the device.
- Use banners that warn against unauthorized access.
- Remove unused services, such as finger, the TCP and UDP small servers, service config, and HTTP server.
- Set up and monitor Syslog.
- Disable automatic trunking on all nontrunk ports.
- Disable CDP on ports where it is not needed.

CCNP BCMSN Quick Reference Sheets

Brent Stewart
Denise Donohue

Copyright© 2007 Cisco Systems, Inc.

Published by: Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this digital short cut may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America 1 2 3 4 5 6 7 8 9 0

First Printing November 2006

ISBN: 1-58705-313-6

Warning and Disclaimer

This digital short cut is designed to provide information about networking. Every effort has been made to make this digital short cut as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc., shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the authors and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

All terms mentioned in this digital short cut that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this digital short cut should not be regarded as affecting the validity of any trademark or service mark.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers’ feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this digital short cut or otherwise alter it to better suit your needs, you can contact us through e-mail at feedback@ciscopress.com. Please make sure to include the digital short cut title and ISBN in your message.

We greatly appreciate your assistance.

Corporate and Government Sales

Cisco Press offers excellent discounts on this digital short cut when ordered in quantity for bulk purchases or special sales. For more information, please contact: U.S. Corporate and Government Sales 1-800-382-3419
corpsales@pearsontechgroup.com

For sales outside the U.S., please contact: International Sales international@pearsoned.com



Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

©2006 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0609R)