BCMSN

Building Cisco Multilayer Switched Networks

Student Guide

Version 2.1

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Table of Contents

Volume 1

Table of Contents

Volume 2

Table of Contents

Volume 3

BCMSN

Course Introduction

Overview

In *Building Cisco Multilayer Switched Networks* (BCMSN) v2.1, network administrators learn how to build networks using multilayer switching technologies over high-speed Ethernet. This course includes both routing and switching, covering Layer 2 and Layer 3 technologies. BCMSN is part of the recommended training path for those students seeking the Cisco CCNP®, CCDP®, CCIP™, and CCIE® certifications.

Outline

The Course Introduction includes these topics:

- Course Objectives
- Cisco Certifications
- Learner Skills and Knowledge
- Learner Responsibilities
- General Administration
- Course Flow Diagram
- Icons and Symbols
- Learner Introductions

Course Objectives

This topic lists the course objectives.

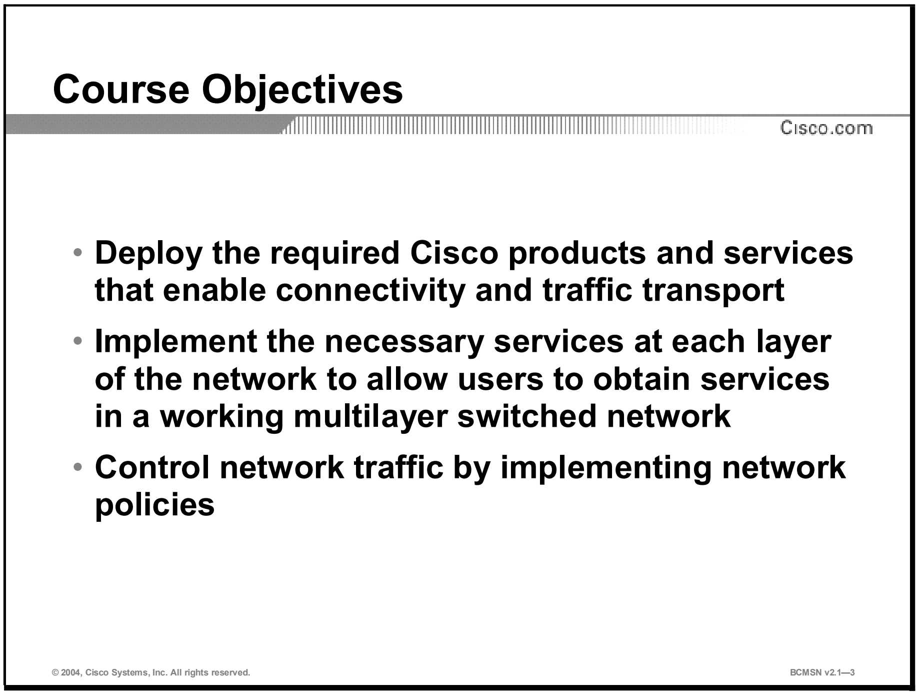

Course Objectives

- Deploy the required Cisco products and services that enable connectivity and traffic transport
- Implement the necessary services at each layer of the network to allow users to obtain services in a working multilayer switched network
- Control network traffic by implementing network policies

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—3

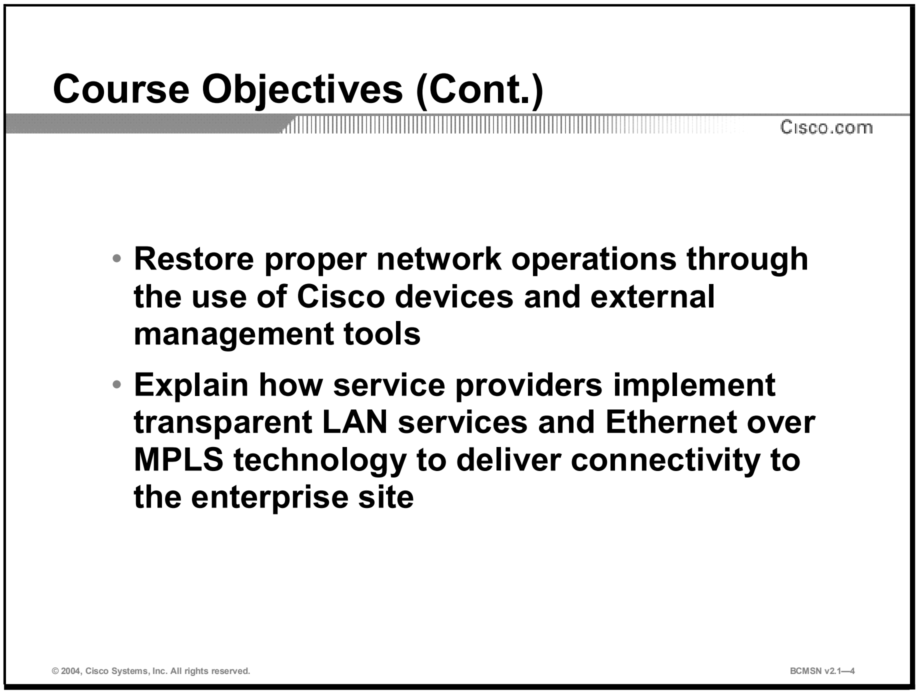

Course Objectives (Cont.)

- Restore proper network operations through the use of Cisco devices and external management tools
- Explain how service providers implement transparent LAN services and Ethernet over MPLS technology to deliver connectivity to the enterprise site

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—4

Upon completing this course, you will be able to:

- Deploy the required Cisco products and services that enable connectivity and traffic transport, given a network design that includes multilayer switching over various Ethernet technologies
- Implement the necessary services at each layer of the network to allow users to obtain services in a working multilayer switched network
- Control network traffic by implementing network policies, given a working multilayer switched network
- Restore proper network operations through the use of Cisco devices and external management tools, when presented with an incorrectly working multilayer switched network
- Explain how service providers implement transparent LAN services and Ethernet over MPLS technology to deliver connectivity to the enterprise site

Cisco Certifications

This topic lists the certification requirements of this course.

Cisco Certifications

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—5

Cisco provides three levels of general career certifications for IT professionals with several different tracks to meet individual needs. Cisco also provides focused Cisco Qualified Specialist (CQS) certifications for designated areas such as cable communications, voice, and security.

There are many paths to Cisco certification, but only one requirement—passing one or more exams demonstrating knowledge and skill. For details, go to http://www.cisco.com/go/certifications.

Learner Skills and Knowledge

This topic lists the recommended course prerequisites.

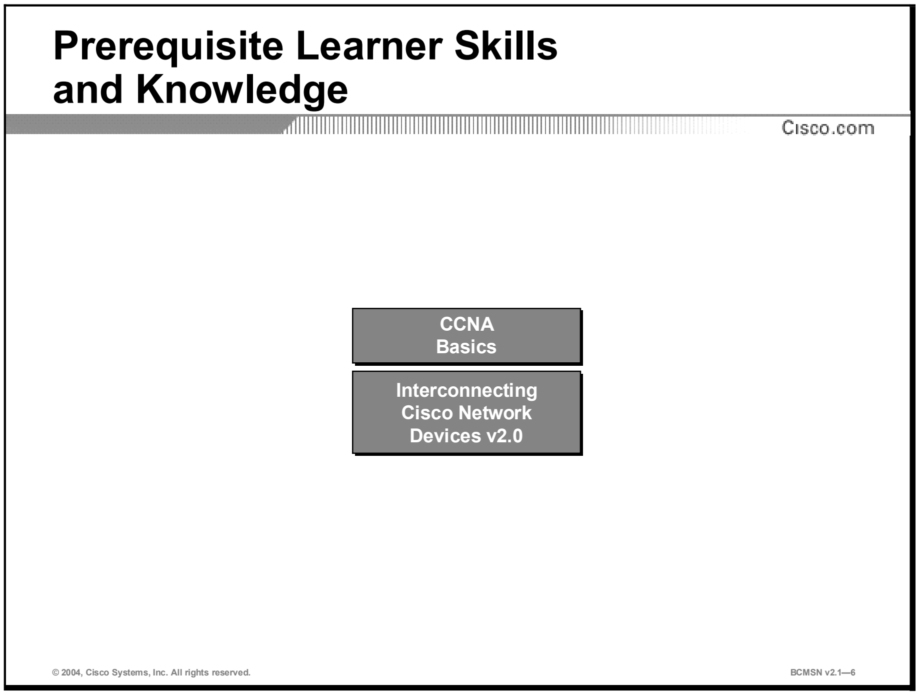

Prerequisite Learner Skills and Knowledge

Cisco.com

CCNA Basics

Interconnecting Cisco Network Devices v2.0

BCMSN v2.1—6

Before taking the BCMSN course, learners should be familiar with internetworking technologies, Cisco products, and Cisco IOS features. Specifically, learners should be familiar with the following before attending this course:

- Basic router configuration
- Basic switch configuration
- Basic VLAN configuration
- Basic Spanning Tree Protocol configuration
- Basic trunking configuration
- Standard access list configuration

Learner Responsibilities

This topic discusses the responsibilities of the learners.

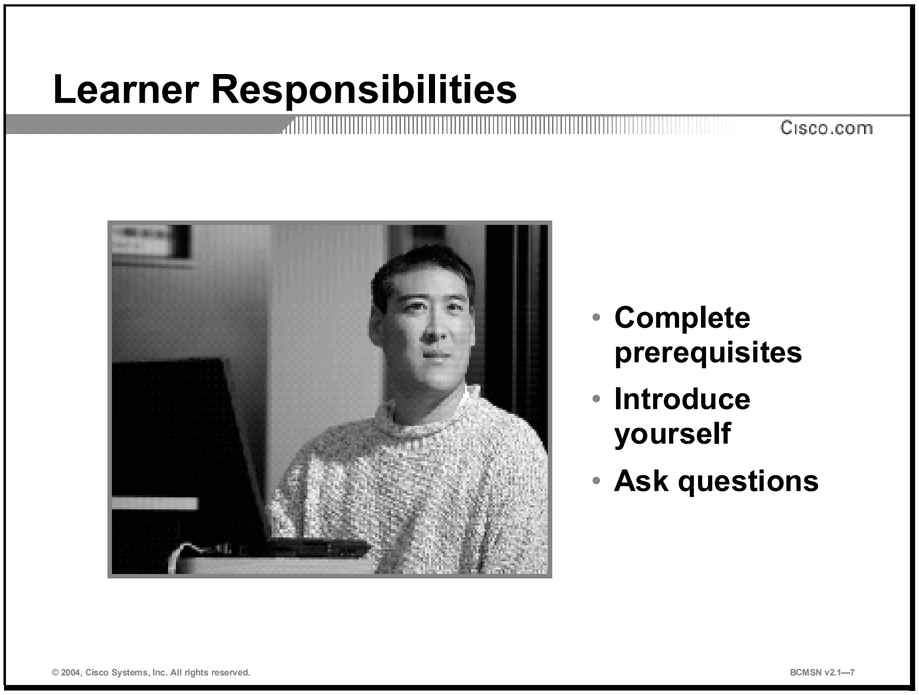

Learner Responsibilities

Cisco.com

- Complete prerequisites
- Introduce yourself
- Ask questions

BCMSN v2.1—7

To take full advantage of the information presented in this course, you must have completed the prerequisite requirements.

In class, you are expected to participate in all lesson exercises and assessments.

In addition, you are encouraged to ask any questions relevant to the course materials.

If you have pertinent information or questions concerning future Cisco product releases and product features, please discuss these topics during breaks or after class. The instructor will answer your questions or direct you to an appropriate information source.

General Administration

This topic lists the administrative issues for the course.

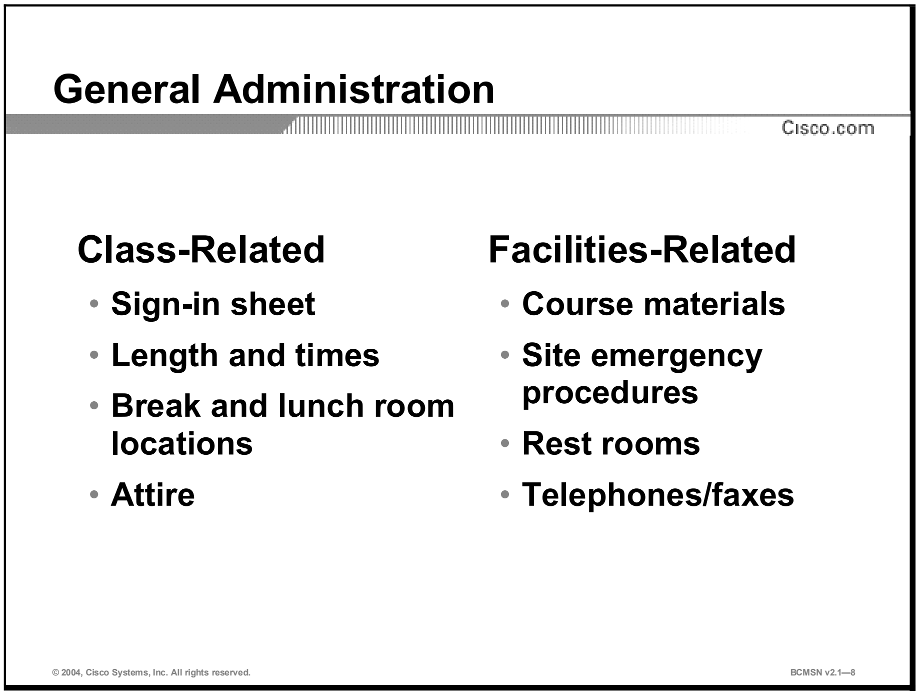

General Administration

Cisco.com

Class-Related

- **Sign-in sheet**
- **Length and times**
- **Break and lunch room locations**
- **Attire**

Facilities-Related

- **Course materials**
- **Site emergency procedures**
- **Rest rooms**
- **Telephones/faxes**

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—8

The instructor will discuss these administrative issues so that you know exactly what to expect from the class:

- Sign-in process
- Starting and anticipated ending times of each class day
- Class breaks and lunch facilities
- Appropriate attire during class
- Materials that you can expect to receive during class
- What to do in the event of an emergency
- Location of the rest rooms
- How to send and receive telephone and fax messages

Course Flow Diagram

This topic covers the suggested flow of the course materials.

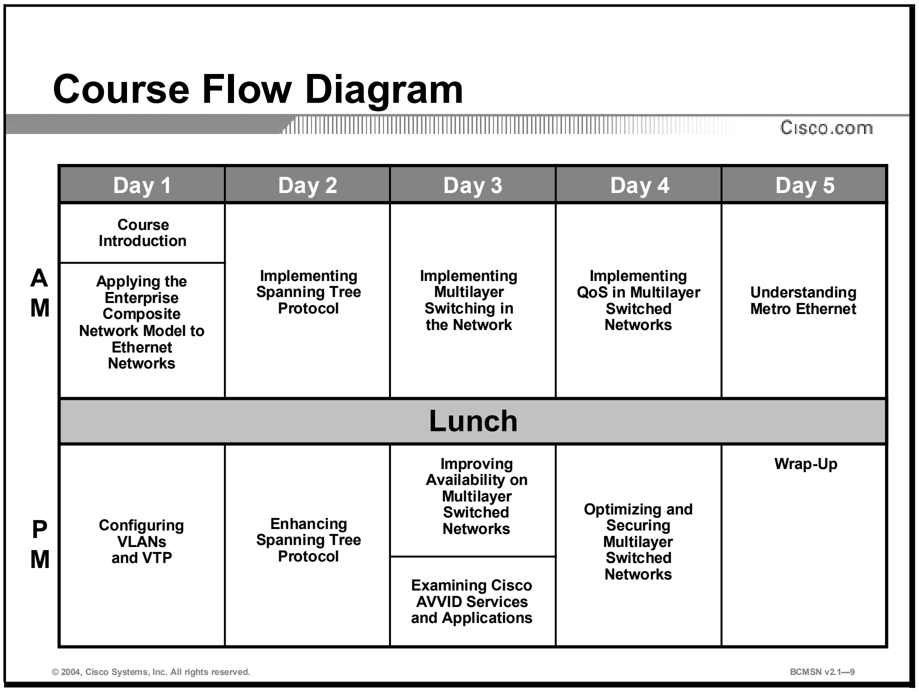

Course Flow Diagram

Cisco.com

	Day 1	Day 2	Day 3	Day 4	Day 5
AM	Course Introduction Applying the Enterprise Composite Network Model to Ethernet Networks	Implementing Spanning Tree Protocol	Implementing Multilayer Switching in the Network	Implementing QoS in Multilayer Switched Networks	Understanding Metro Ethernet
	Lunch				
PM	Configuring VLANs and VTP	Enhancing Spanning Tree Protocol	Improving Availability on Multilayer Switched Networks Examining Cisco AVVID Services and Applications	Optimizing and Securing Multilayer Switched Networks	Wrap-Up

BCMSN v2.1—9

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the lesson assessments and exercises. The exact timing of the subject materials and labs depends on the pace of your specific class.

Icons and Symbols

This topic shows the Cisco icons and symbols that are used in this course.

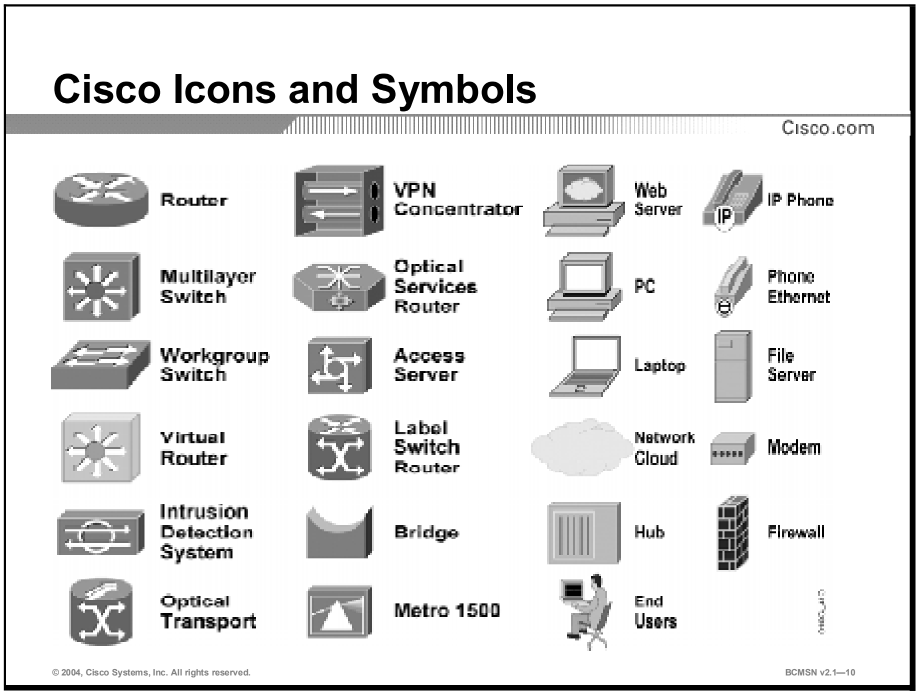

Cisco Icons and Symbols

Cisco.com

Router

Multilayer Switch

Workgroup Switch

Virtual Router

Intrusion Detection System

Optical Transport

VPN Concentrator

Optical Services Router

Access Server

Label Switch Router

Bridge

Metro 1500

Web Server

PC

Laptop

Network Cloud

Hub

End Users

IP Phone

Phone Ethernet

File Server

Modem

Firewall

BCMSN v2.1—10

Learner Introductions

This is the point in the course where you introduce yourself.

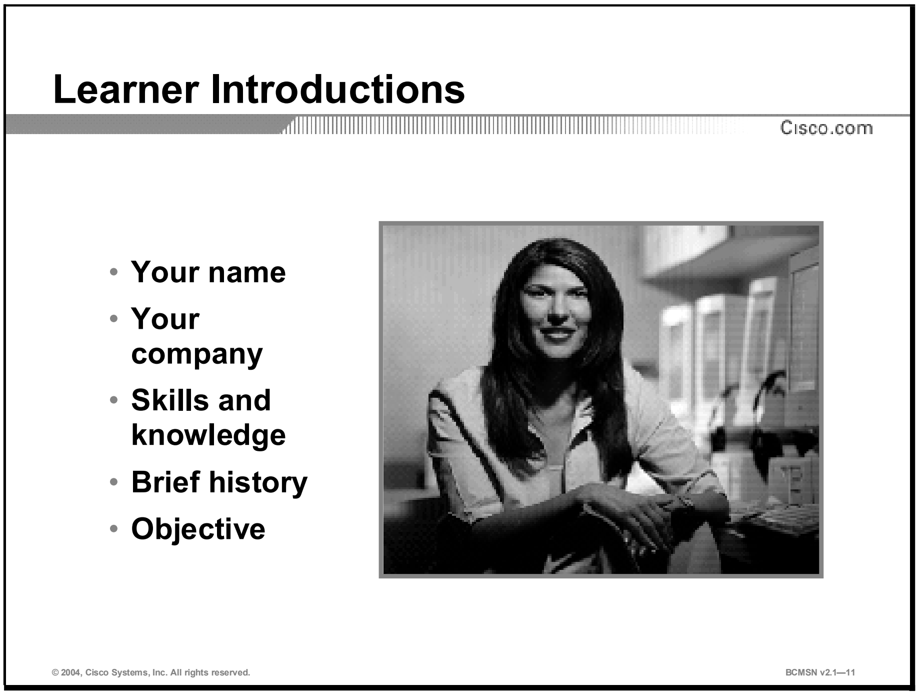

Learner Introductions

Cisco.com

- Your name
- Your company
- Skills and knowledge
- Brief history
- Objective

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—11

Prepare to share this information:

- Your name
- Your company
- The prerequisite skills that you have
- A profile of your experience
- What you would like to learn from this course

Module 1

Applying the Enterprise Composite Network Model to Ethernet Networks

Overview

Multilayer switches combine traditional Layer 2 switching with Layer 3 routing in a single product through a fast hardware implementation. Advances in hardware have enabled the recent rise of the multilayer switch. New higher-density application-specific integrated circuits (ASICs) allow real-time switching and forwarding at wirespeed. This increased throughput is accomplished at a lower cost than with traditional software-based routers built around general-purpose processors.

Upon completing this module, you will be able to:

- Select the most appropriate factor that would solve a specific problem in an Ethernet network
- Label the parts of an Enterprise Composite Network model based on the functionality of each part
- Select whether a switch or a router best solves a specific problem in a Campus Infrastructure module
- Select the correct data link layer technology with the most appropriate location on a diagram of the Enterprise Composite Network model

Outline

The module contains these components:

- Using Switching and Routing in Ethernet Networks
- Addressing Common Network Problems
- Using Multilayer Switches in the Campus Infrastructure Module
- Using Data-Link Layer Technologies in an Enterprise Composite Network Model
- Lesson Assessments

Using Switching and Routing in Ethernet Networks

Overview

Networks that do not exhibit a layered or hierarchical design are subject to several issues that can be eliminated with the implementation of a layered approach to the design. Layer 2 switches deliver the ability to increase bandwidth. Because the routing functionality operates at the Layer 3 level, routing provides connectivity between Layer 2 physical and virtual LANs (VLANs) and will ultimately be required at some point in a network design.

Relevance

To build a cost-effective and efficient multilayer switched network, it is critical to understand how and why switches are incorporated into the various components of the network.

Objectives

Upon completing this lesson, you will be able to:

- Match communication features to the correct communication type
- Select the issues that can occur in an unlayered network
- Select the correct features of a Layer 2 switch
- Select the issues that can occur in a Layer 2 network
- Select the correct features of routers
- Select the correct issues that can occur in a routed network
- Select the correct features of multiplayer switches
- Identify the correct issue that can occur in a poorly designed network

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the *Interconnecting Cisco Network Devices* (ICND) course or passing the Cisco CCNA® certification exam

Outline

This lesson includes these topics:

- Overview
- Identifying the Communication Types in a Large Network
- Discovering the Issues That Can Occur in an Unlayered Network
- Identifying the Features of Layer 2 Switches
- Discovering the Issues That Can Occur in a Layer 2 Network
- Identifying the Features of Routers
- Discovering the Issues That Can Occur in a Routed Network
- Identifying the Features of Multilayer Switches
- Discovering the Issues That Can Occur in a Poorly Designed Network
- Summary
- Quiz

Identifying the Communication Types in a Large Network

This topic identifies the features of each communication type in a large network.

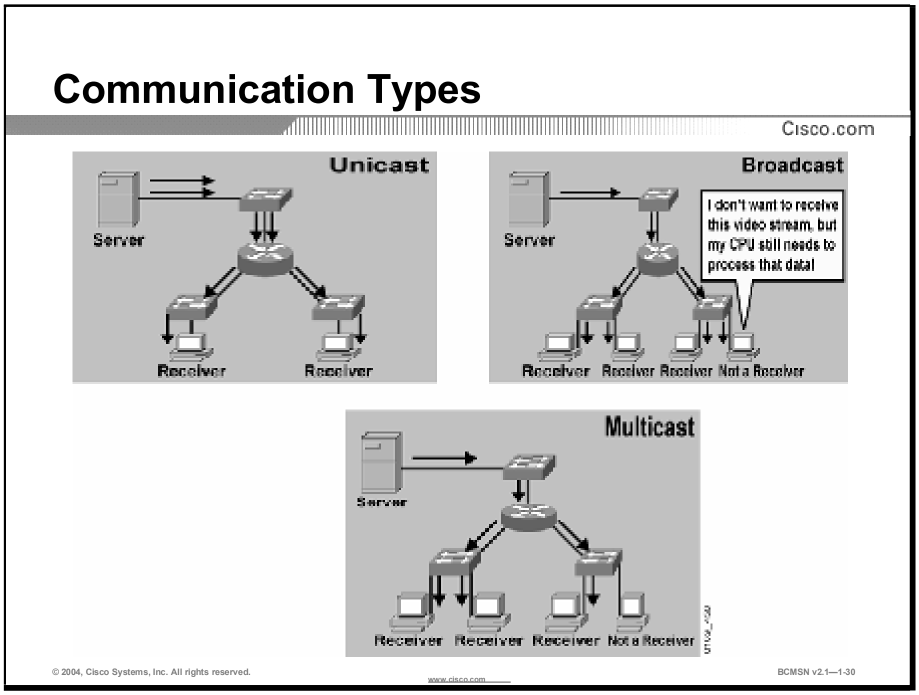

Communication Types

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

www.cisco.com

BCMSN v2.1—1-30

There are three types of communication in a large network:

- Unicast
- Multicast
- Broadcast

Unknown unicasts and multicasts are types of Layer 2 broadcasts.

Unicasts are communications from a source to one specific destination.

Examples of broadcasts are IP Address Resolution Protocol (ARP) requests, NetBIOS name requests, or Internetwork Packet Exchange (IPX) Get Nearest Server (GNS) requests. These types of broadcasts typically flood the entire subnet, with the intention of having only the target destination device respond directly to the broadcast.

Multicast traffic can also consume bandwidth. Multicast traffic is treated the same as broadcasts except that the intended destination is a specific group or subset of destinations. Depending on the number of users in a group or the type of application data contained in the multicast packet, this type of broadcast can consume most, if not all, of the network resources. Examples of multicast implementations are the Cisco IP/TV application using multicast packets to distribute multimedia data, and Novell 5 on IP using multicast packets to locate services.

Discovering the Issues That Can Occur in an Unlayered Network

This topic identifies the issues that can occur in an unlayered network.

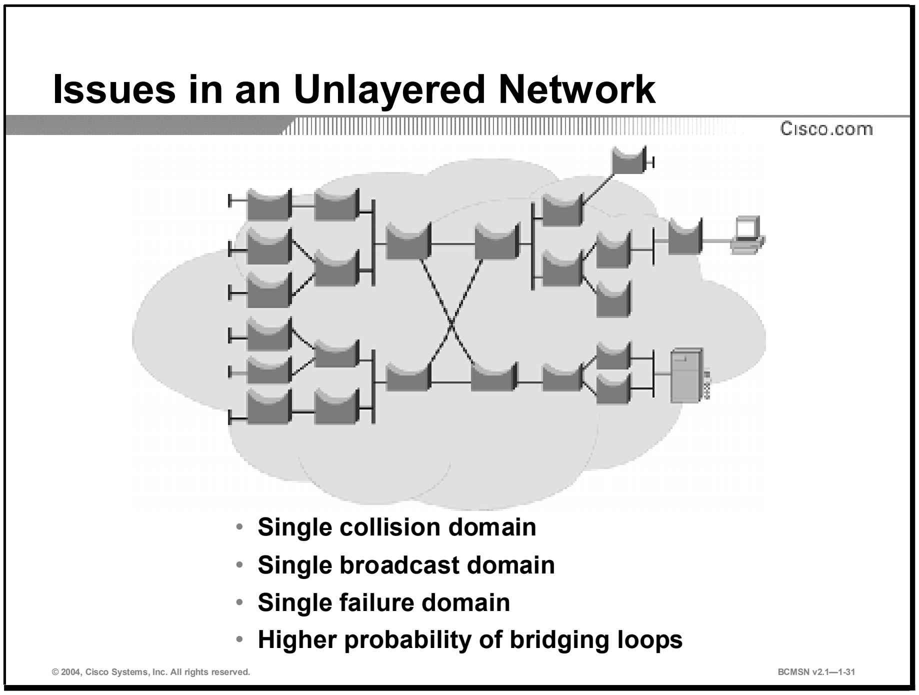

Issues in an Unlayered Network

Cisco.com

- Single collision domain
- Single broadcast domain
- Single failure domain
- Higher probability of bridging loops

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-31

Networks that do not exhibit a layered or hierarchical design are subject to several issues that can be effectively eliminated only with the implementation of a layered approach to their design.

Unlayered flat networks consist of large network segments that extend large broadcast and failure domains. Because of the carrier sense multiple access collision detect (CSMA/CD) nature of Ethernet and the inefficient use of bandwidth a flat network introduces, performance, network resources, and reliability decrease. As network availability decreases, the occurrence of broadcast domains, failure domains, and network congestion increases—sometimes to the point of network failure.

Because an unlayered, flat network is larger in size, its boundaries and member devices are not always clearly defined. As a result, the potential for miswiring increases. Wiring errors can lead to Layer 2 bridging loops that create erroneous forwarding paths and broadcast storms.

When problems or issues do occur, unlayered, flat networks do not easily lend themselves for problem isolation and determination.

Identifying the Features of Layer 2 Switches

This topic identifies the features of a Layer 2 switch.

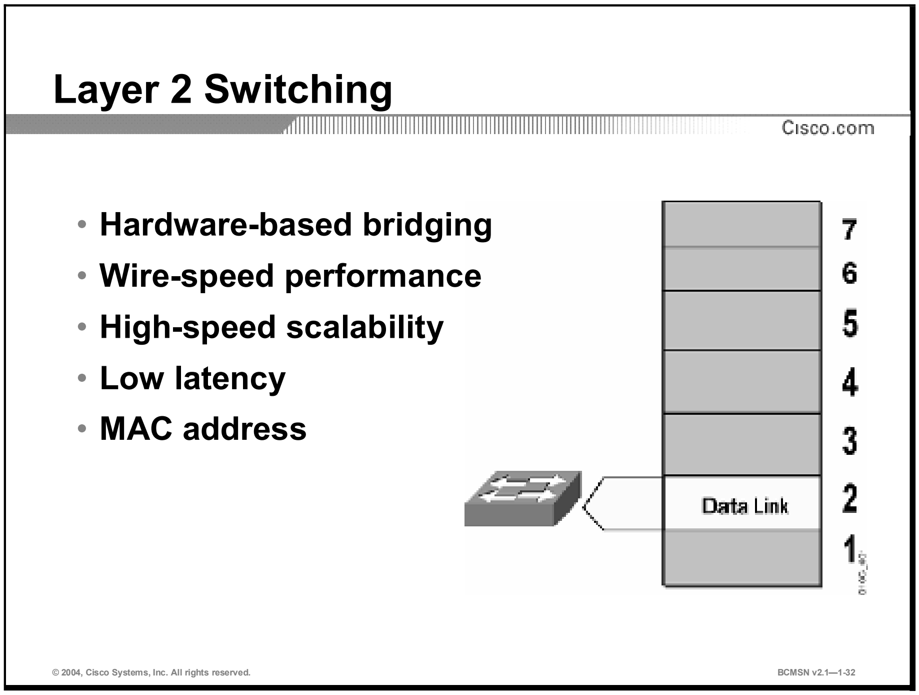

Layer 2 Switching

Cisco.com

- Hardware-based bridging
- Wire-speed performance
- High-speed scalability
- Low latency
- MAC address

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—1-32

Layer 2 switches provide high-speed scalability to the wiring closet and the Campus Backbone submodule.

Layer 2 switching is hardware-based bridging. In a switch, specialized hardware chips called application-specific integrated circuits (ASICs) handle frame forwarding.

Layer 2 switches deliver the ability to increase bandwidth to the wiring closet attached end-user nodes, without adding unnecessary complexity to the network. Layer 2 data frames consist of both infrastructure content, such as MAC addresses, and end-user content. At Layer 2, no modification is required to the frame or its content when going between Layer 1 interfaces, such as Fast Ethernet to 10 Gigabit Ethernet. Frame switching takes advantage of the relatively simple process of examining MAC addresses within the frame.

Discovering the Issues That Can Occur in a Layer 2 Network

This topic identifies the issues that can occur in a Layer 2 network.

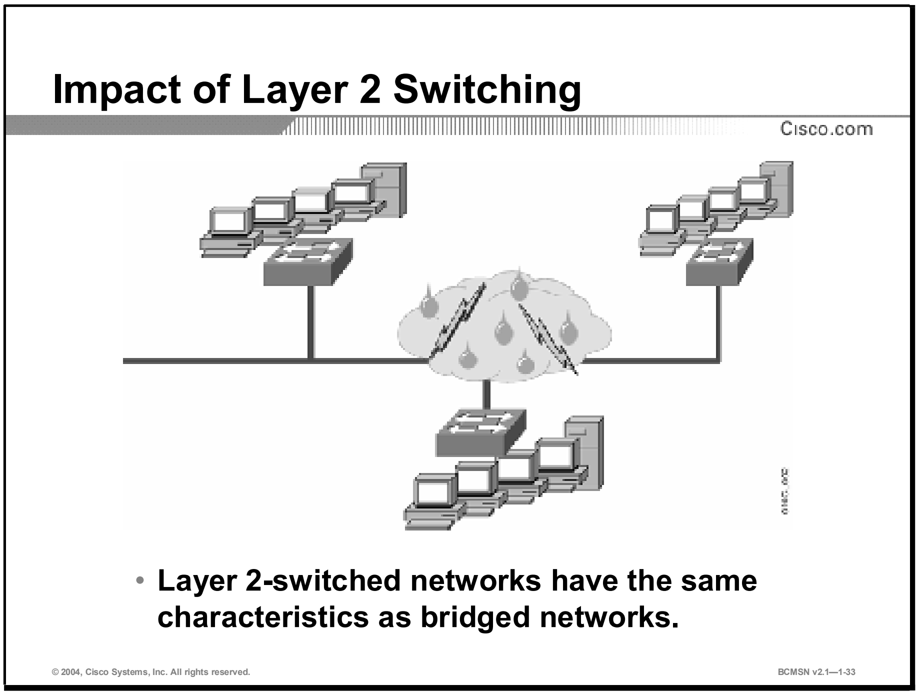

Impact of Layer 2 Switching

Cisco.com

- **Layer 2-switched networks have the same characteristics as bridged networks.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-33

The features and functionality of a Layer 2 switch effectively eliminate collision domains. In addition, these features and functionality are used to produce network designs that decrease the number of hosts per network segment to reduce the size of any one broadcast domain. Decreasing the hosts per segment leads to a design with more segments in the network. This technique is called “segmentation.”

However, for all its advantages, Layer 2 switching has all the same characteristics and limitations as bridging, and it cannot provide any kind of connectivity between the segments.

Identifying the Features of Routers

This topic identifies the features of routers.

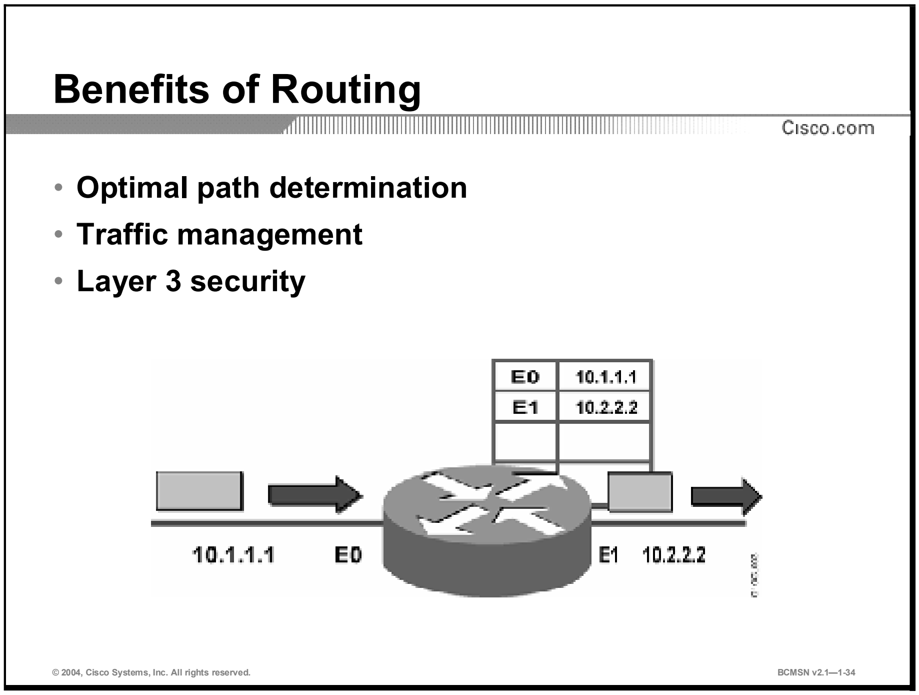

Benefits of Routing

Cisco.com

- Optimal path determination
- Traffic management
- Layer 3 security

E0	10.1.1.1
E1	10.2.2.2

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-34

Because the routing functionality operates at Layer 3, routing provides connectivity between Layer 2 physical and logical network segments (virtual LANs or VLANs) and will be required at some point in a network design.

Routing provides for an optimal path determination process or “routing process” and examines each incoming packet to determine what route the packet should take through the network and across network segments. A router that has an interface in each network segment performs routing decisions. Routers also provide for the connectivity between the segments. Routing between segments involves a determination of the next network point to which a Layer 3 packet should be forwarded toward its destination.

Routing is a well-defined set of packet manipulations that do not forward broadcasts by default. This characteristic completes the ability to constrain broadcast and failure domains to a single limited segment and or VLAN while still enabling network communications.

A network layer address identifies an entity. This entity can be a source, a destination, or an intermediate Layer 3 device’s interface, which is called a “logical address.” For the TCP/IP protocol stack, the logical address is called an “IP address.” Routers and other internetworking devices require one IP address per physical or logical network segment to provide connection for each network layer protocol supported.

Discovering the Issues That Can Occur in a Routed Network

This topic identifies the problems that can occur in a routed network.

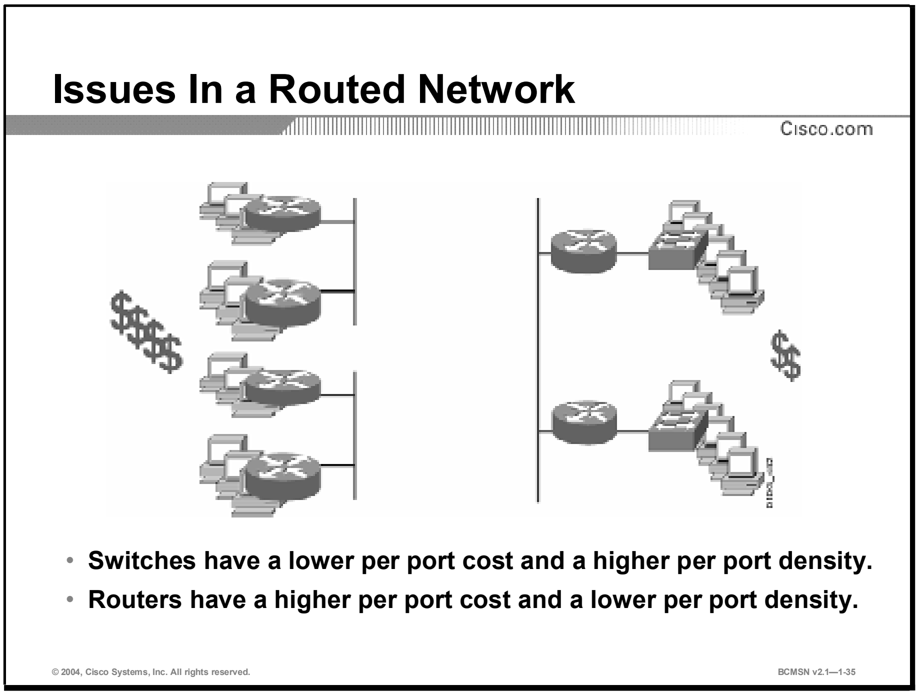

Issues In a Routed Network

Cisco.com

- Switches have a lower per port cost and a higher per port density.
- Routers have a higher per port cost and a lower per port density.

BCMSN v2.1—1-35

At one time, routed points within a network were seen as bottlenecks to be avoided. At that time, routers were slow because they were software-based. Those routers were also expensive. Modern routers employ the same or similar forwarding architectures and technologies that switches do. Routers now do most of their processing and path determination via specialized hardware circuitry.

Routers are still more expensive per port than switches. Switches are designed for port density and operate at the lower OSI layer, making their forwarding decisions less process intensive. While some routers are inexpensive, those routers are generally limited in performance and in the number of interfaces and features. Because of these factors, routers are still limited in the specific role they can fulfill.

Identifying the Features of Multilayer Switches

This topic identifies the features of multilayer switches.

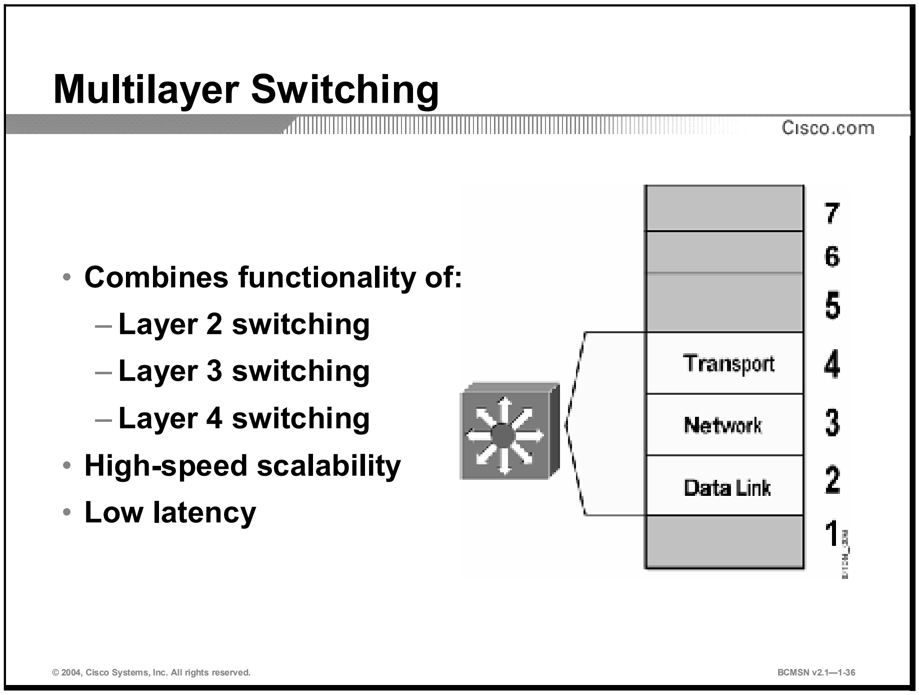

Multilayer Switching

Cisco.com

- Combines functionality of:
  - Layer 2 switching
  - Layer 3 switching
  - Layer 4 switching
- High-speed scalability
- Low latency

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-36

Multilayer switching is hardware-based switching and routing integrated into a single platform. In some cases, the frame and packet forwarding operation is handled by the same specialized hardware ASICs and other specialized circuitry. A multilayer switch does everything to a frame and packet that a traditional switch and router does, such as the following:

- Micro-segments collision domains
- Provides multiple simultaneous switching paths
- Segments broadcast and failure domains
- Provides destination specific frame forwarding based on Layer 2 information
- Determines the forwarding path based on Layer 3 information
- Validates the integrity of the Layer 2 frame and Layer 3 packet via checksums and other methods
- Verifies packet expiration and updates accordingly
- Processes and responds to any option information
- Updates forwarding statistics in the Management Information Base (MIB)
- Applies security controls, if required
- Provides optimal path determination

The primary differences between the packet-forwarding operation of a router and multilayer switching operation of a multilayer switch are the physical implementation and the OSI layer of operations. The technologies of Layer 2 switching and Layer 3 forwarding are applied in a single platform.

Because it is designed to handle high-performance LAN traffic, a multilayer switch can be placed anywhere within the network, cost-effectively replacing the traditional switches and routers. Generally, however, a multilayer switch is more than is required in the Building Access submodule.

Discovering the Issues That Can Occur in a Poorly Designed Network

This topic identifies the issues that occur in a poorly designed network.

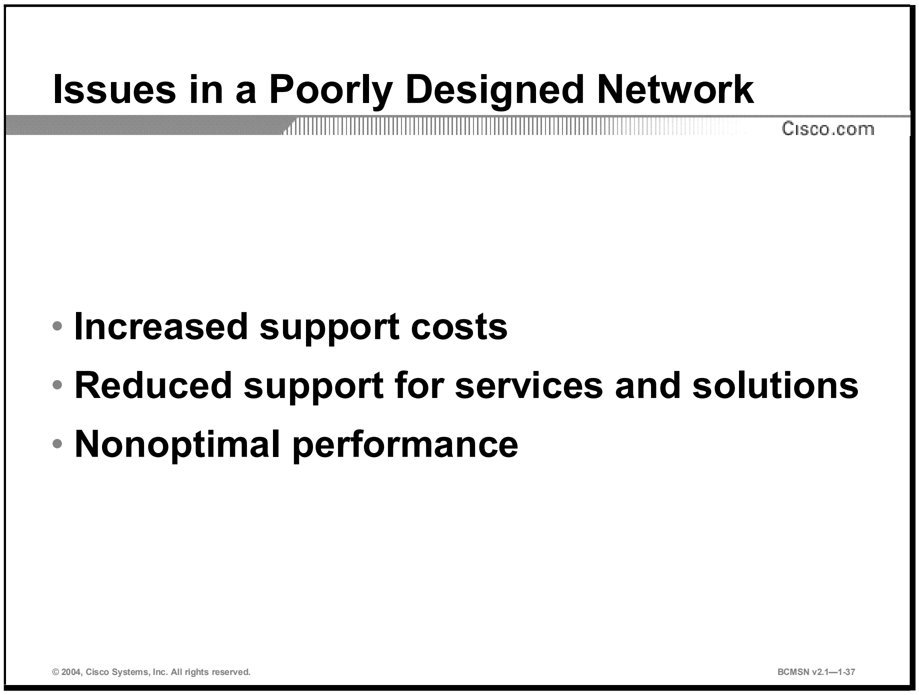

Issues in a Poorly Designed Network

Cisco.com

- Increased support costs
- Reduced support for services and solutions
- Nonoptimal performance

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-37

A poorly designed network has increased support costs, reduced services and solutions that can be supported, and nonoptimal performance issues that most likely will affect end users directly. Here are some of the issues that stem from a poorly designed network:

- **Failure domains:** One of the most important reasons to implement an effective design is to reduce the impact of network problems resulting from s errors in protocol-stack implementations, network configuration, faulty NICs transmitting bad frames or packets, and broadcast-intensive applications. These issues could have an impact on the entire network in a flat or poorly designed switched environment, potentially to the point of a total failure. Failure domains coexist with broadcast domains and can be bound to a single network segment.
- **Broadcast domains:** Broadcast domains exist in every network. Many applications and network operations require broadcasts to function properly. Therefore, it is not possible to completely eliminate broadcast domains. However, just as with failure domains, it is crucial to minimize any impact on a network. Broadcasts consume end-user resources and affect bandwidth availability. End-user resources are consumed because every broadcast must be processed to some extent by the end nodes receiving the broadcast, regardless of whether the broadcast was intended for that node. A poorly designed network will propagate broadcasts over a large number of end-user nodes affecting a large portion of a network population.

- **Isolation of unknown MAC unicast floods:** Internally, Catalyst switches implement VLANs by maintaining a MAC table on a per VLAN basis. This limits frame forwarding to ports only in the same VLAN. In the case of unicast traffic, this is limited to only that specific switch port on which the destination device is located. However, if a frame arrives for a destination MAC address that has not been recorded in the MAC table, the switch will perform what is called an "unknown MAC unicast flood." The frame is flooded to all switch ports within the same VLAN in an effort to deliver the frame. A unicast flood also has the same effect as a broadcast in that every end node must receive the frame to determine whether the date is destined for that node. Again, a proper network design limits the impact of such floods with a small, well-defined subset of the network.
- **Isolation of multicast traffic:** IP multicast is a technique that allows IP traffic to be propagated from one source to a number of destinations, or from many sources to many destinations. Rather than sending one packet to each destination, one packet is sent to the multicast group identified by a single IP and MAC destination group address pair. Similar to unicast flooding and broadcasting, multicasts can congest and have a negative impact on a poorly designed network. A proper design allows for multicast functionality without the negative effects.
- **Ease of management and support:** Because a poorly designed network is unorganized and lacks the deterministic features found within the Enterprise Composite Network model, support, maintenance, and problem resolution become extremely time consuming and arduous tasks.

A poorly designed network always has a negative impact and becomes a burden for any organization in terms of support and related costs.

Summary

This topic summarizes the key points discussed in this lesson.

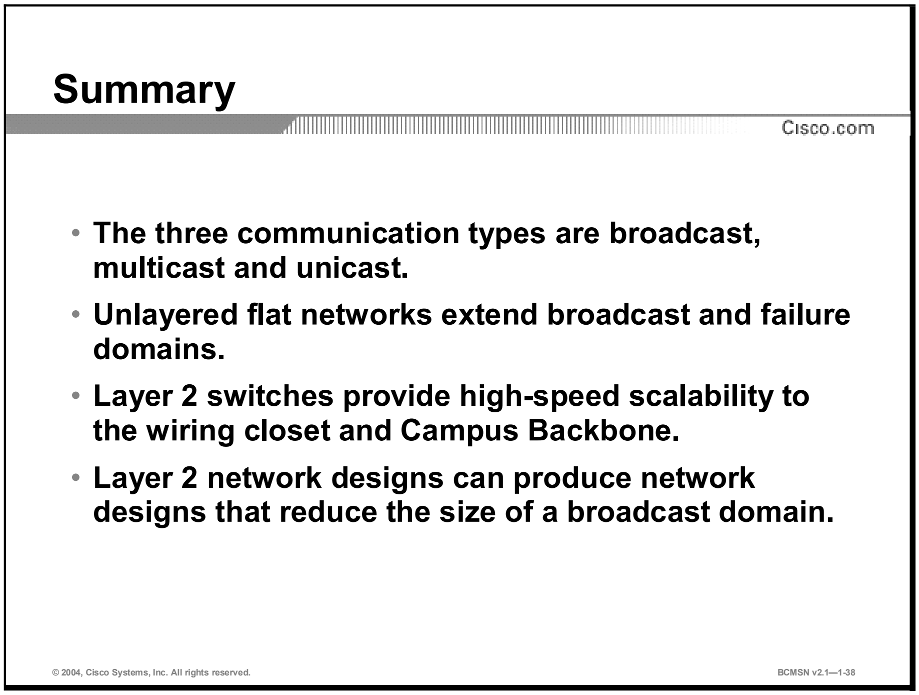

Summary

Cisco.com

- The three communication types are broadcast, multicast and unicast.
- Unlayered flat networks extend broadcast and failure domains.
- Layer 2 switches provide high-speed scalability to the wiring closet and Campus Backbone.
- Layer 2 network designs can produce network designs that reduce the size of a broadcast domain.

BCMSN v2.1—1-38

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Match the type of communication with the appropriate feature.

_____ 1. unicast

_____ 2. broadcast

_____ 3. multicast

A) broadcasts to one user

B) broadcasts to a specific group

C) can propagate throughout a network

Q2) Select the issues that can occur in an unlayered network. (Choose two.)

A) There is an inefficient use of switches.

B) Large broadcast and failure domains are extended.

C) Wiring errors can lead to blocked bridges.

D) Problem isolation and determination are difficult.

Q3) Select the correct features of a Layer 2 switch. (Choose two.)

A) It provides high-speed scalability to the Access and Distribution modules.

B) It supports hardware-based bridging.

C) A slight modification is required to the packet infrastructure content.

D) It increases bandwidth to the wiring closet.

Q4) Select the issues that can occur in a Layer 2 network. (Choose two.)

A) cannot provide any kind of connectivity between segments

B) has same limitations as bridging

C) slow performance

D) more expensive

Q5) Select the correct features of routers. (Choose two.)

A) Routers increase bandwidth to the wiring closet.

B) Routers determine the next network point to which a packet should be forwarded toward its destination.

C) Routers provide connectivity between Layer 2 virtual LANs.

D) Frame forwarding is handled by specialized hardware called ASICs.

Q6) Select the issues that can occur in a routed network. (Choose two.)

A) expensive

B) boundaries not clearly defined

C) minimal connectivity

D) uses specialized hardware

Q7) Select the features of multiplayer switches. (Choose two.)

A) combines Layer 2 switching and Layer 3 routing functionality

B) provides high-speed scalability

C) determines forwarding path based on Layer 2 information

D) verifies packet information and updates hourly

Q8) Select the issues that can occur in a poorly designed network. (Choose two.)

A) Network errors can affect all hosts on a network segment.

B) Broadcasts are propagated over a large number of end-user nodes.

C) Network is divided into defined subsets.

D) Network errors affect only key segments.

Quiz Answer Key

Q1) 1=A, 2=C, 3=B

Relates to: Identifying the Communication Types in a Large Network

Q2) B, C

Relates to: Discovering the Issues That Can Occur in an Unlayered Network

Q3) B, C

Relates to: Identifying the Features of Layer 2 Switches

Q4) A, B

Relates to: Discovering the Issues That Can Occur in a Layer 2 Network

Q5) B, C

Relates to: Identifying the Features of Routers

Q6) A, C

Relates to: Discovering the Issues That Can Occur in a Routed Network

Q7) A, B

Relates to: Identifying the Features of Multilayer Switches

Q8) A, B

Relates to: Discovering the Issues That Can Occur in a Poorly Designed Network

Addressing Common Network Problems

Overview

The Enterprise Composite Network model provides a framework for designing the components of an enterprise network. The model relies on the hierarchical approach, allowing for flexibility in network design, and facilitates implementation and troubleshooting.

Relevance

Developing a common vocabulary and architecture is critical to implementing modular enterprise networks that provide performance, scalability, and availability.

Objectives

Upon completing this lesson, you will be able to:

- Match a layer of a hierarchical network model with its correct function
- Select the correct features for the Enterprise Composite Network model
- Match the correct feature with the appropriate functional area
- Match the correct functional area with the appropriate feature
- Match the correct submodule with the appropriate function
- Select the correct Cisco components that support voice and data in the Enterprise Composite Network model

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the *Interconnecting Cisco Network Devices* (ICND) course or passing the Cisco CCNA® certification exam

Outline

This lesson includes these topics:

- Overview
- Identifying the Layers of a Hierarchical Network Model
- Identifying the Functional Areas of the Enterprise Composite Network Model
- Identifying the Features of the Enterprise Composite Network Model
- Distinguishing the Modules of the Enterprise Campus Functional Areas
- Distinguishing the Submodules in the Campus Infrastructure Module
- Integrating Voice and Data in the Enterprise Composite Network Model
- Summary
- Quiz

Identifying the Layers of a Hierarchical Network Model

This topic identifies the function of each layer in the hierarchical network model.

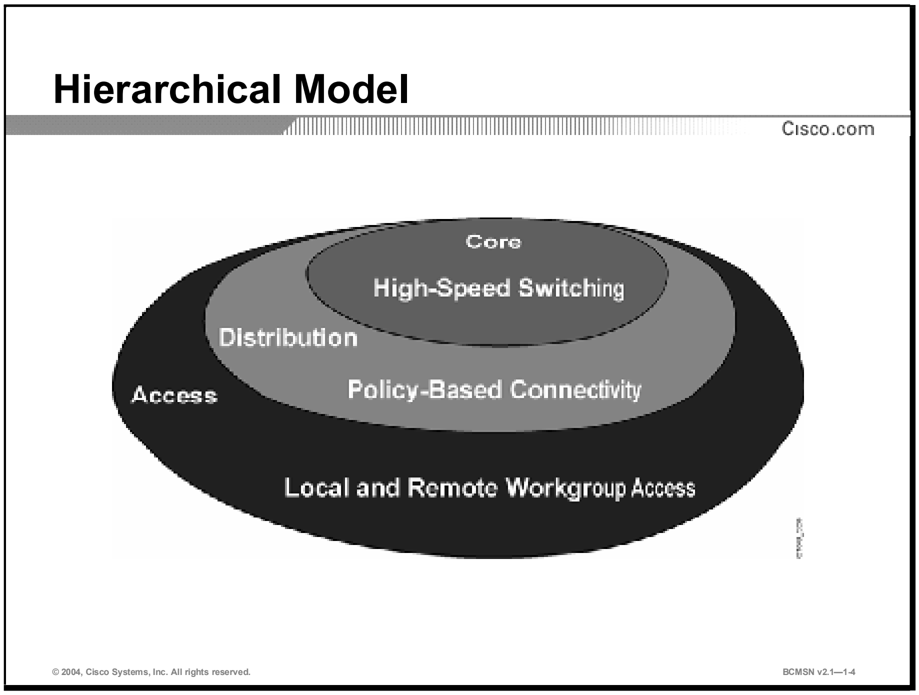

Hierarchical Model

Cisco.com

BCMSN v2.1—1-4

The Enterprise Composite Network model provides a modular framework for designing networks. The modularity within the model allows flexibility in network design and facilitates implementation and troubleshooting. The hierarchical model divides networks into the Building Access, Building Distribution, and Building Core layers, as follows:

- **Building Access layer:** The Building Access layer is used to grant user access to network devices. At a network campus, the Building Access layer generally incorporates switched LAN devices with ports that provide connectivity to workstations and servers. In the WAN environment, the Building Access layer can provide sites with access to the corporate network using a WAN technology.
- **Building Distribution layer:** The Building Distribution layer aggregates the wiring closets and uses switches to segment workgroups and isolate network problems. Routing and packet manipulation occur in the Building Distribution layer.
- **Building Core layer:** The Building Core layer is a high-speed backbone and is designed to switch packets as fast as possible. Because the core is critical for connectivity, it must provide a high level of availability and must adapt to changes very quickly.

A simple hierarchical model is useful but has weaknesses when implementing large, complex enterprise networks.

Identifying the Functional Areas of the Enterprise Composite Network Model

This topic identifies the functional areas of the Enterprise Composite Network model.

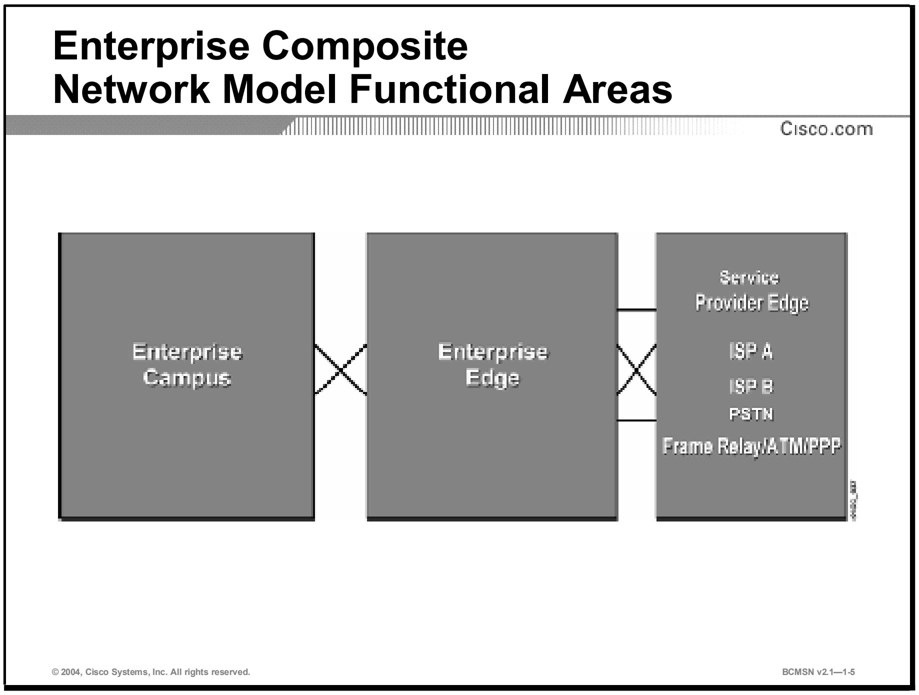

Enterprise Composite Network Model Functional Areas

Cisco.com

Enterprise Campus

Enterprise Edge

Service Provider Edge

ISP A

ISP B

PSTN

Frame Relay/ATM/PPP

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-5

The Enterprise Composite Network model introduces additional modularity into the network structure. The entire network is divided into functional areas that contain the hierarchical model access, distribution, and core layers in each functional area.

The Enterprise Composite Network model contains these three major functional areas:

- **Enterprise Campus:** The Enterprise Campus functional area contains the modules required to build a hierarchical, highly robust campus network that offers performance, scalability, and availability. This functional area contains the network elements required for independent operation within a single campus. This functional area does not offer remote connections or Internet access.

  A campus is defined as one or more buildings, with multiple virtual and physical networks, connected across a high-performance, multilayer switched backbone.

- **Enterprise Edge:** The Enterprise Edge functional area aggregates connectivity from the various elements at the Edge of the enterprise network. This functional area filters traffic from the Edge modules and routes it into the Enterprise Campus functional area. The Enterprise Edge functional area contains all of the network elements for efficient and secure communication between the Enterprise Campus and remote locations, remote users, and the Internet.

- **Service Provider Edge:** The Service Provider Edge functional area provides functionality implemented by service providers. This functional area enables communication with other networks using different WAN technologies and Internet service providers (ISPs).

Identifying the Features of the Enterprise Composite Network Model

This topic identifies the features of the Enterprise Composite Network model.

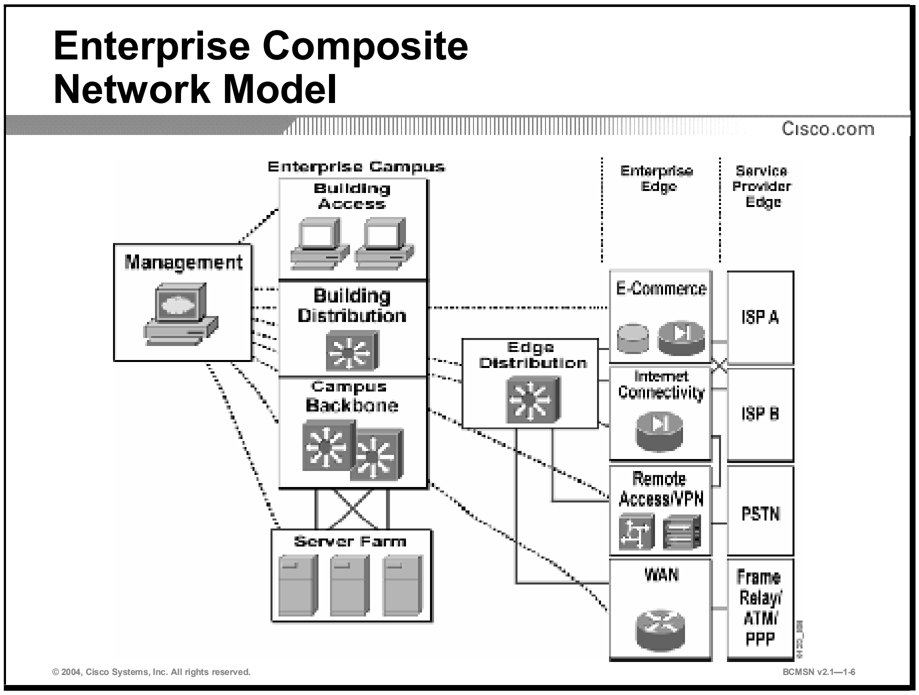

Enterprise Composite Network Model

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-6

To scale the hierarchical model, Cisco introduced the Enterprise Composite Network model, which further divides the enterprise network into physical, logical, and functional boundaries. The Enterprise Composite Network model contains functional areas, each of which having its own Building Access, Building Distribution, and Building Core layers.

The Enterprise Composite Network model meets these criteria:

- It defines a deterministic network with clearly defined boundaries between modules. The model also has clear demarcation points, so that the designer knows exactly where traffic is located.
- It increases network scalability and eases the design task by making each module discrete.
- It provides scalability by allowing enterprises to add modules easily. As network complexity grows, designers can add new functional modules.
- It offers more network integrity in network design, allowing the designer to add services and solutions without changing the underlying network design.

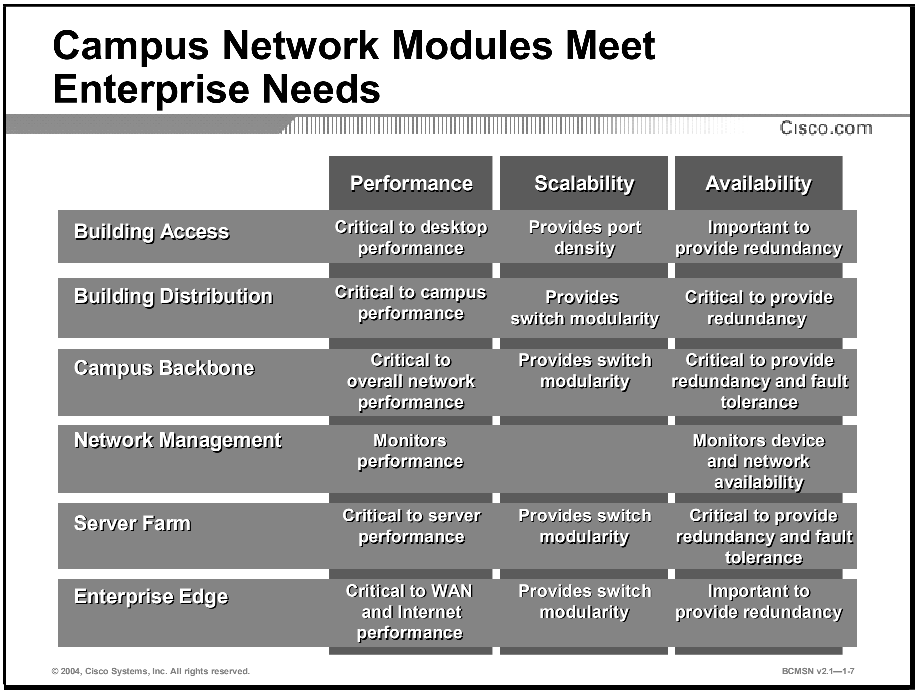

Campus Network Modules Meet Enterprise Needs

Cisco.com

	Performance	Scalability	Availability
Building Access	Critical to desktop performance	Provides port density	Important to provide redundancy
Building Distribution	Critical to campus performance	Provides switch modularity	Critical to provide redundancy
Campus Backbone	Critical to overall network performance	Provides switch modularity	Critical to provide redundancy and fault tolerance
Network Management	Monitors performance		Monitors device and network availability
Server Farm	Critical to server performance	Provides switch modularity	Critical to provide redundancy and fault tolerance
Enterprise Edge	Critical to WAN and Internet performance	Provides switch modularity	Important to provide redundancy

BCMSN v2.1—1-7

The campus network meets enterprise network needs for performance, scalability, and availability.

Distinguishing the Modules of the Enterprise Campus Functional Areas

This topic identifies the features of each functional area in the Enterprise Campus.

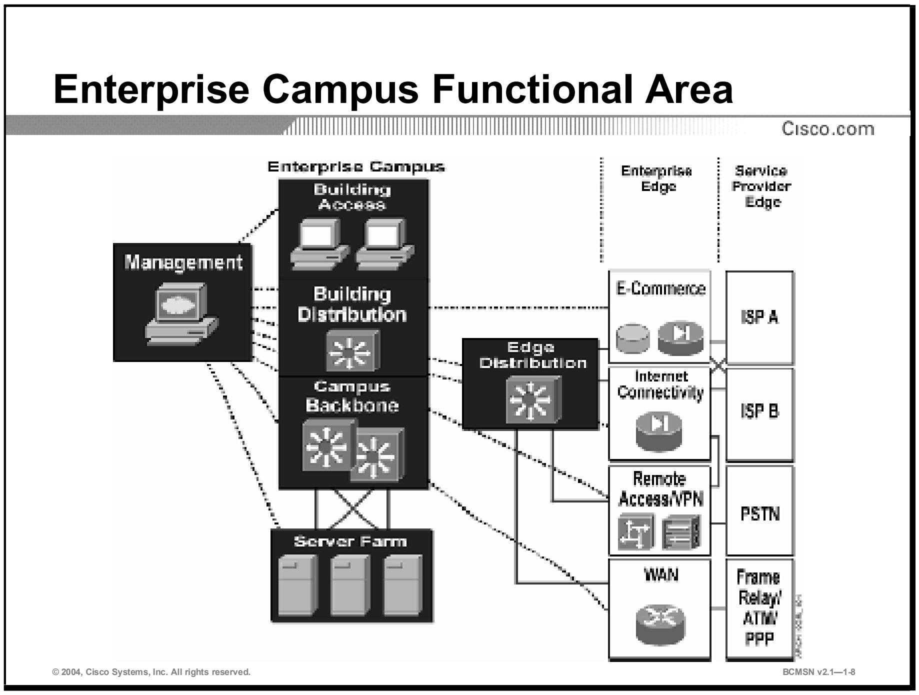

Enterprise Campus Functional Area

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-8

The Enterprise Campus functional area includes the Campus Infrastructure, Network Management, Server Farm, and Edge Distribution modules. Each module has a specific function within the campus network.

The Campus Infrastructure module connects users within a campus with the Server Farm and Edge Distribution modules. This module is composed of one or more floors or buildings connected to the Campus Backbone submodule. Each building contains a Building Access and Building Distribution submodule.

The Network Management module performs system logging and authentication as well as network monitoring and general configuration management functions.

The Server Farm module contains e-mail and corporate servers providing application, file, print, e-mail, and Domain Name System (DNS) services to internal users.

The Edge Distribution module aggregates the connectivity from the various elements at the Enterprise Edge functional area and routes the traffic into the Campus Backbone submodule.

Distinguishing the Submodules in the Campus Infrastructure Module

This topic identifies the features of each of the submodules in the Campus Infrastructure module.

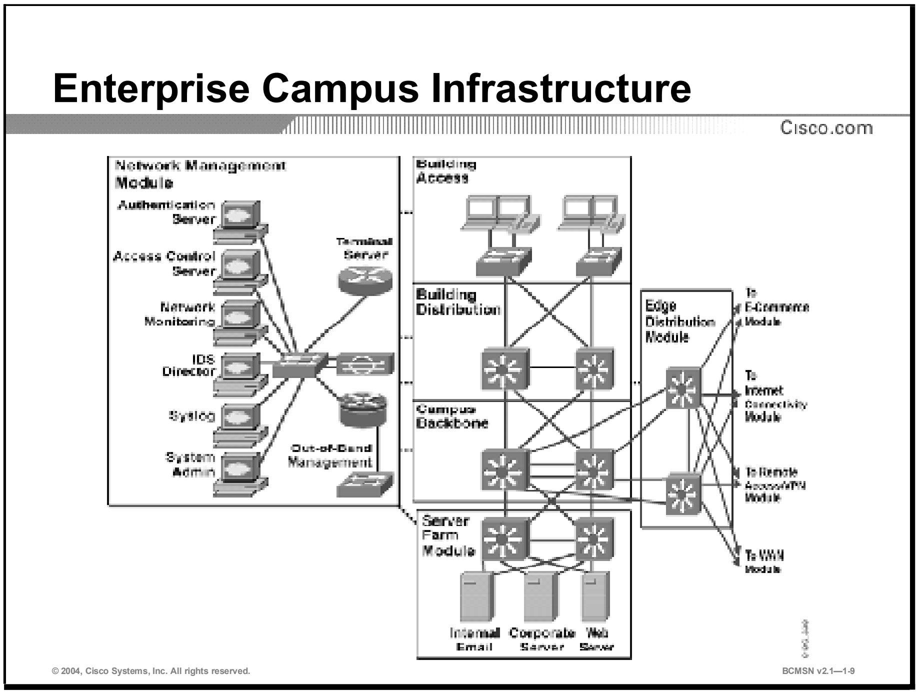

Enterprise Campus Infrastructure

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-9

The Campus Infrastructure module connects users within a campus with the Server Farm and Edge Distribution modules. This module is composed of one or more floors or buildings connected to the Campus Backbone submodule. Each building contains a Building Access and Building Distribution submodule.

The Campus Infrastructure module includes these submodules:

- **Building Access submodule (also known as Building Access layer):** Contains end-user workstations, IP Phones, and Layer 2 access switches that connect devices to the Building Distribution submodule. The Building Access submodule performs important services, such as Layer 2 and Layer 3 broadcast suppression, protocol filtering, network access, and quality of service (QoS).
- **Building Distribution submodule (also known as Building Distribution layer:** Provides aggregation of Building Access devices, often using Layer 3 switching. The Building Distribution submodule performs routing, QoS, and access control. Requests for data flow into the Building Distribution switches and to the campus core. This submodule provides fast failure recovery, because each Building Distribution switch maintains two equal-cost paths in the routing table to every destination network. When one connection to the campus core fails, all routes immediately switch over to the remaining path after the link failure is detected.

- **Campus Backbone submodule (also known as Building Core layer):** Provides redundant and fast-converging connectivity between buildings, as well as with the Server Farm and Edge Distribution modules. The Campus Backbone submodule routes and switches traffic as fast as possible from one module to another. This module uses Layer 2 or Layer 3 switches for high throughput functions with added routing, QoS, and security features.

Integrating Voice and Data in the Enterprise Composite Network Model

This topic identifies the Cisco components that support voice and data in the Enterprise Composite Network model.

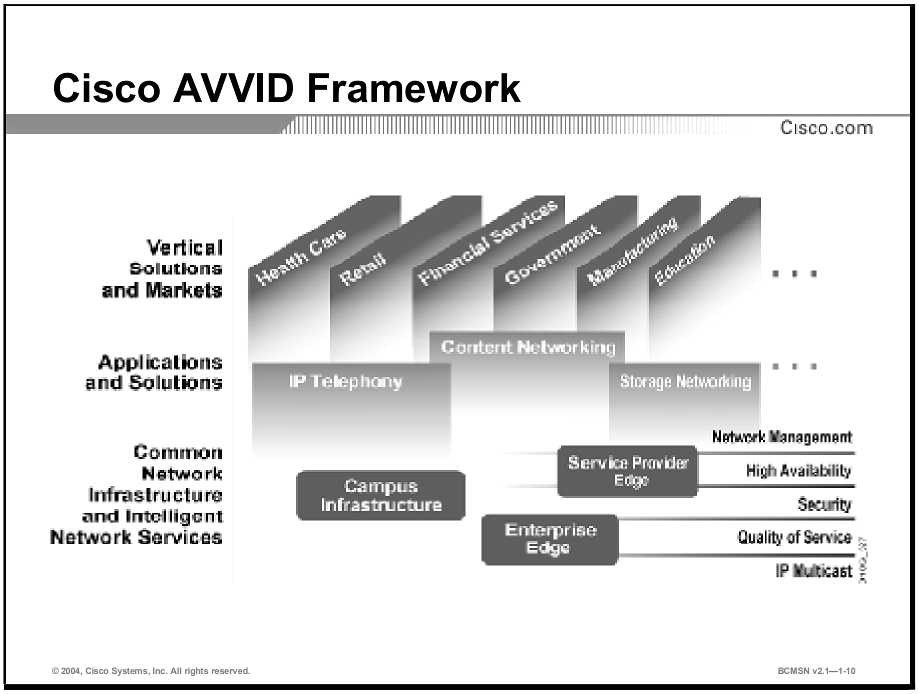

Cisco AVVID Framework

Cisco.com

Vertical Solutions and Markets

Health Care

Retail

Financial Services

Government

Manufacturing

Education

. . .

Applications and Solutions

Content Networking

IP Telephony

Storage Networking

. . .

Common Network Infrastructure and Intelligent Network Services

Campus Infrastructure

Service Provider Edge

Enterprise Edge

Network Management

High Availability

Security

Quality of Service

IP Multicast

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-10

Cisco Architecture for Voice, Video and Integrated Data (AVVID) is an enterprise-wide, standards-based network architecture that provides a roadmap for combining business and technology strategies into a cohesive model. The primary components of Cisco AVVID architecture include network infrastructure, intelligent network services, and network solutions.

A network architecture is a roadmap and guide for ongoing network planning, design, and implementation. As such, a network architecture provides a coherent framework that unifies disparate solutions onto a single foundation.

The Cisco AVVID framework supports these key components:

- **Network infrastructure:** Network infrastructure includes the hardware and software used to send, receive, and manage datagrams that are transmitted between end-user devices throughout the enterprise. Examples of these devices are routers, LAN switches, WAN switches, and PBXs.
- **Intelligent network services:** Intelligent network services add intelligence to the network infrastructure beyond just moving a datagram between two points. Intelligent network services allow for application awareness and include security, network management, QoS, IP multicast, and high availability.
- **Network solutions:** Network solutions allow enterprises to make business decisions about the business itself as well as about networks and the technologies and applications that run on them. Network-based applications enable an enterprise organization to interact more effectively with customers, suppliers, partners, and employees.

Summary

This topic summarizes the key points discussed in this lesson.

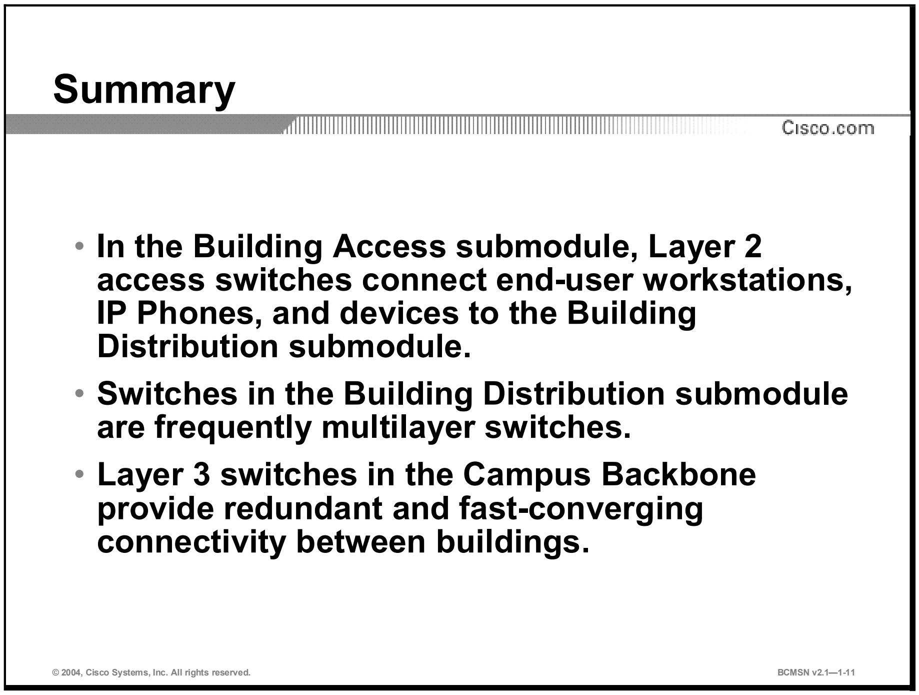

Summary

Cisco.com

- In the Building Access submodule, Layer 2 access switches connect end-user workstations, IP Phones, and devices to the Building Distribution submodule.
- Switches in the Building Distribution submodule are frequently multilayer switches.
- Layer 3 switches in the Campus Backbone provide redundant and fast-converging connectivity between buildings.

BCMSN v2.1—1-11

References

For additional information, refer to this resource:

- “Cisco AVVID: Enabling E-Business” at http://www.cisco.com/en/US/netsol/netwarch/ns19/net_solution_home.html

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Match each hierarchical network model layer with the appropriate function.

_____ 1. Building Access layer

_____ 2. Building Distribution layer

_____ 3. Building Core layer

A) designed to switch packets as fast as possible

B) grants user access to network devices

C) uses Layer 2 and Layer 3 switching to segment workgroups

Q2) Match each Enterprise Composite Network model functional area with its correct feature.

_____ 1. Enterprise Edge

_____ 2. Enterprise Campus

_____ 3. Service Provider Edge

A) aggregates connectivity from the various elements at the Edge of the enterprise network

B) enables communication with other networks using different WAN technologies and Internet service providers

C) contains the modules required to build a hierarchical, highly robust campus network that offers performance, scalability, and availability

Q3) Select the correct features of the Enterprise Composite Network model. (Choose two.)

A) This model defines a network with loosely defined boundaries between modules.

B) Each functional area has its own hierarchical layers.

C) This model allows enterprises to add modules easily.

D) This model requires only slight changes to the underlying network design when adding services and solutions.

Q4) Match each module with its appropriate feature.

_____ 1. Campus Infrastructure module

_____ 2. Edge Distribution module

_____ 3. Server Farm module

_____ 4. Network Management module

A) provides Domain Name System services to internal users

B) provides network monitoring and general configuration management functions

C) routes traffic into the Campus Backbone submodule

D) connects users within a campus with the Server Farm and Edge Distribution modules

Q5) Match each Campus Infrastructure submodule with the appropriate function.

_____ 1. Campus Backbone

_____ 2. Building Distribution

_____ 3. Building Access

A) It performs important services such as Layer 2 and Layer 3 broadcast suppression, protocol filtering, network access, and QoS.

B) It provides redundant and fast-converging connectivity between buildings.

C) Each switch maintains two equal-cost paths in the routing table to every destination network, thus providing fast failure recovery.

Q6) Which Cisco AVVID component includes security, network management, and quality of service?

A) vertical solutions

B) network solutions

C) network infrastructure

D) intelligent network services

Quiz Answer Key

Q1) 1=B, 2=C, 3=A

Relates to: Identifying the Layers of a Hierarchical Network Model

Q2) 1=A, 2=C, 3=B

Relates to: Identifying the Functional Areas of the Enterprise Composite Network Model

Q3) B, C

Relates to: Identifying the Features of the Enterprise Composite Network Model

Q4) 1=D, 2=C, 3=A, 4=B

Relates to: Distinguishing the Modules of the Enterprise Campus Functional Areas

Q5) 1=B, 2=C, 3=A

Relates to: Distinguishing the Submodules in the Campus Infrastructure Module

Q6) D

Relates to: Integrating Voice and Data in the Enterprise Composite Network Model

Using Multilayer Switches in the Campus Infrastructure Module

Overview

Traditionally, switches provided strictly Layer 2 functionality. In effect, a Layer 2 switch is hardware-based bridging. Multilayer switches expand this capability to provide functionality beyond Layer 2 throughout the network. In many situations, a multilayer switch can take the place of a traditional router.

Relevance

To build a cost-effective and efficient multilayer switched network, it is critical to understand how and why switches are incorporated into the various components of the network.

Objectives

Upon completing this lesson, you will be able to:

- Select the correct roles of multilayer switches in the Building Access submodule
- Select the correct roles of multilayer switches in the Building Distribution submodule
- Select the correct roles of multilayer switches in the Building Core submodule
- Select the correct role of multilayer switches in the Server Farm module
- Select the correct role of multilayer switches in the Edge Distribution module

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Identifying the Role of Multilayer Switches in the Building Access Submodule
- Identifying the Role of Multilayer Switches in the Building Distribution Submodule
- Identifying the Role of Multilayer Switches in the Building Core Submodule
- Identifying the Role of Multilayer Switches in the Server Farm Module
- Identifying the Role of Multilayer Switches in the Edge Distribution Module
- Summary
- Quiz

Identifying the Role of Multilayer Switches in the Building Access Submodule

This topic identifies the role of multilayer switches in the Building Access submodule.

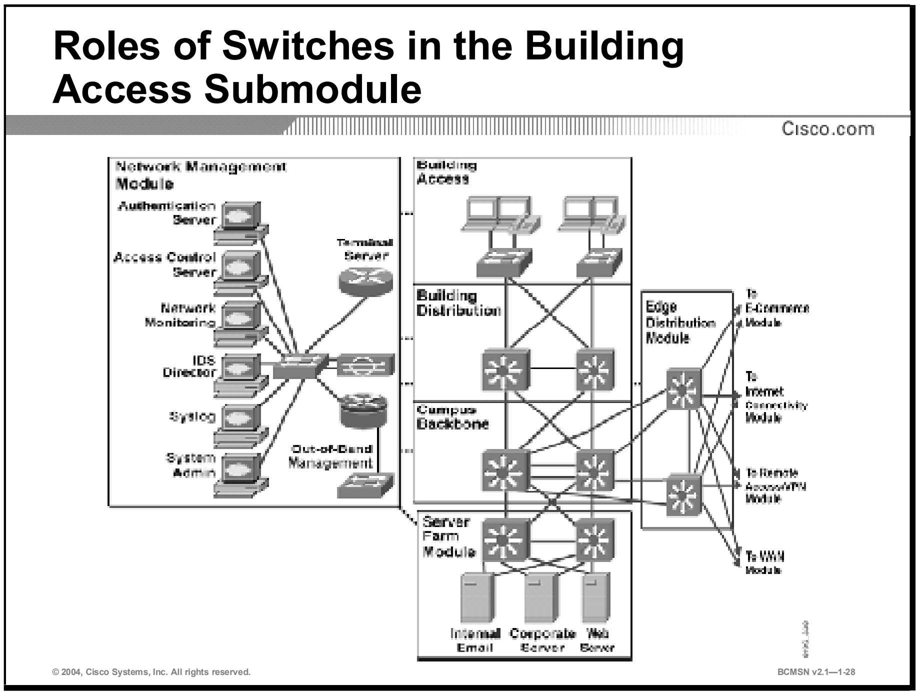

Roles of Switches in the Building Access Submodule

Cisco.com

BCMSN v2.1—1-28

In the Building Access submodule**,** Layer 2 access switches connect end-user workstations, IP Phones, and devices to the Building Distribution submodule.

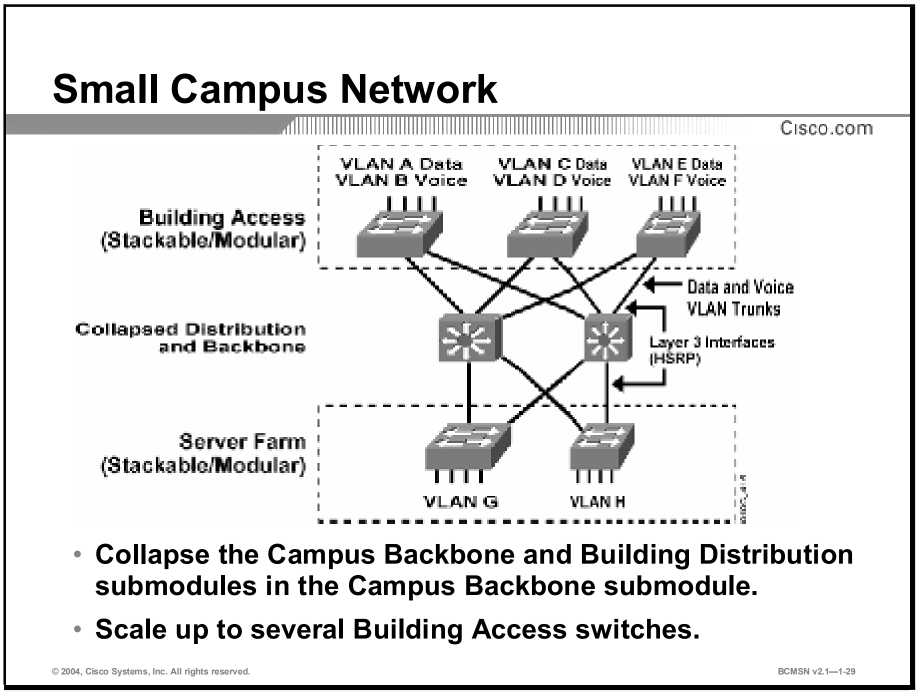

Small Campus Network

Cisco.com

- Collapse the Campus Backbone and Building Distribution submodules in the Campus Backbone submodule.
- Scale up to several Building Access switches.

BCMSN v2.1—1-29

Switches here, typically placed in a wiring closet, perform important services, such as broadcast suppression, protocol filtering, network access, and QoS marking.

Identifying the Role of Multilayer Switches in the Building Distribution Submodule

This topic identifies the role of multilayer switches in the Building Distribution submodule.

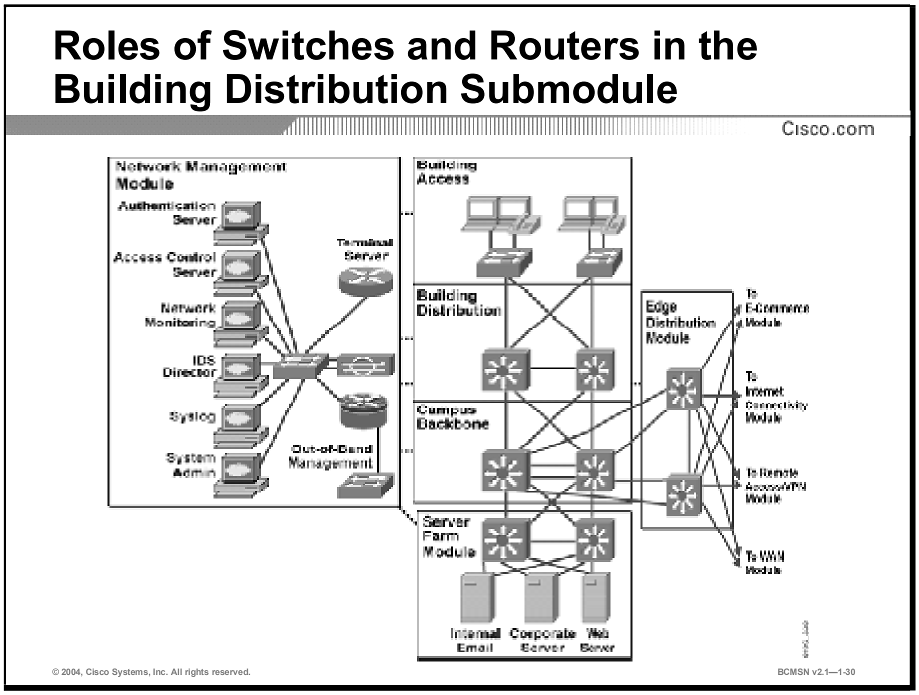

Roles of Switches and Routers in the Building Distribution Submodule

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-30

Switches in the Building Distribution submodule are multilayer switches that terminate broadcast and failure domains, keeping these domains to a manageable size. Multilayer switches provide aggregation of wiring closets, perform routing, QoS, and access control. This frees the Building Access and Campus Backbone submodules from dedicating resources to those tasks.

The Building Distribution submodule accepts requests for data flow to the Campus Backbone and other submodules and functional areas. In this way, the Building Access submodule switches can function at a lower cost switch and only provide a limited set of features. Because there are many more building access switches, to provide the port density required, this offloading of responsibility is extremely cost-effective. Also, the backbone is freed from the extensive frame and packet manipulations that the building distribution performs instead, such as routing and priority handling.

Identifying the Role of Multilayer Switches in the Building Core Submodule

This topic identifies the role of multilayer switches in the Building Core submodule.

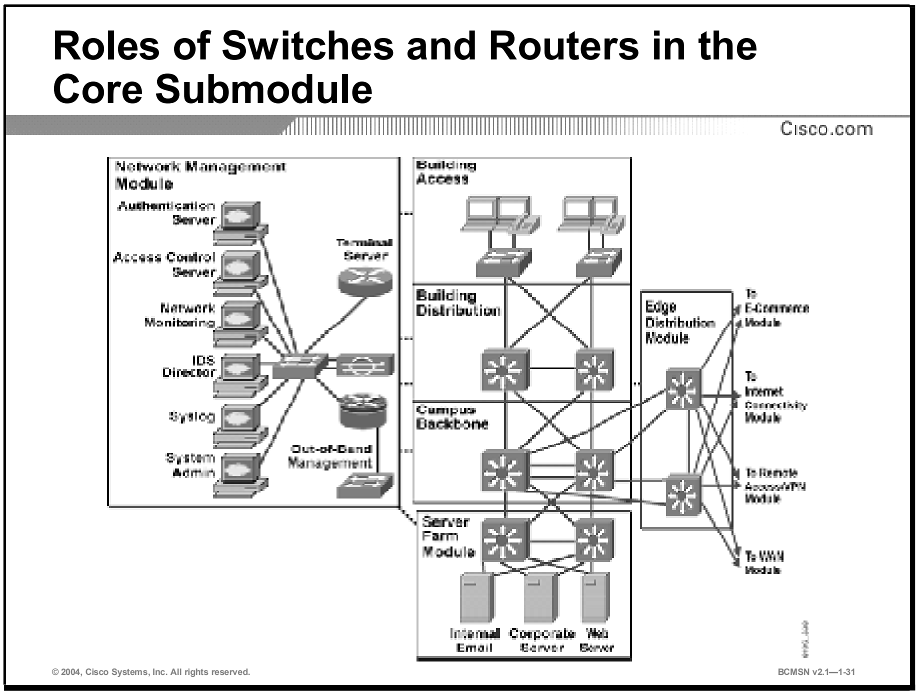

Roles of Switches and Routers in the Core Submodule

Cisco.com

BCMSN v2.1—1-31

Multilayer switches in the Campus Backbone submodule provide redundant and fast-converging connectivity between submodules and with the Server Farm and Edge Distribution modules. The Campus Backbone submodule uses Layer 2 or multilayer switches for high throughput functions according to the routing, QoS, and security features imposed at the Building Distribution layer.

Identifying the Role of Multilayer Switches in the Server Farm Module

This topic identifies the role of multilayer switches in the Server Farm module.

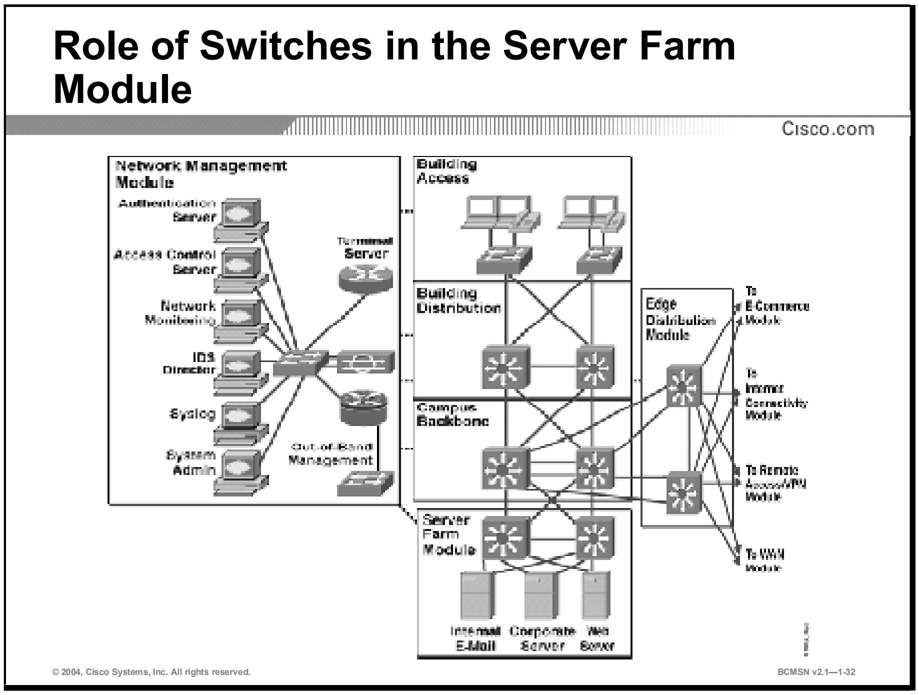

Role of Switches in the Server Farm Module

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-32

Within the Server Farm module, switches provide access between the Server Farm and the Core Backbone submodule. Switches may also provide access between servers and storage.

The Server Farm module contains internal e-mail and corporate servers that provide application, file, print, e-mail, and DNS services to internal users. Because access to these servers is vital, they are connected to two different switches, enabling full redundancy and load sharing. The Server Farm module switches are cross connected with Building Core layer switches, enabling high reliability and availability of all servers. Depending on the type of storage model deployed, additional switches may be used internally in the Server Farm module to provide access between servers and storage devices.

Identifying the Role of Multilayer Switches in the Edge Distribution Module

This topic identifies the role of multilayer switches in the Edge Distribution module.

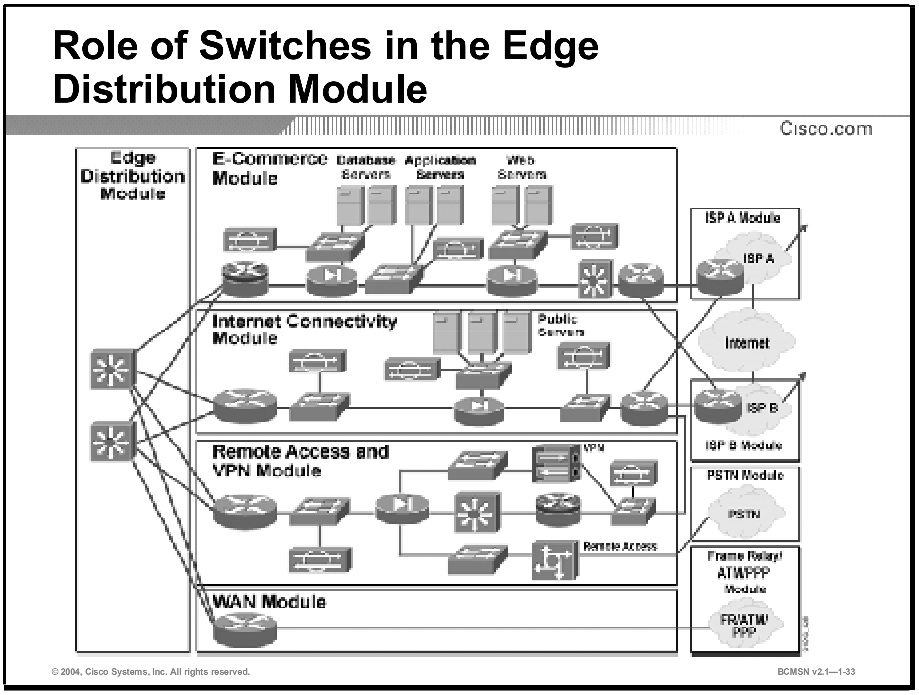

Role of Switches in the Edge Distribution Module

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-33

In the Edge Distribution module, switches perform the distribution function between the Enterprise Edge functional area and the Building Core layer. In addition, switches can play a role in the various Enterprise Edge modules.

- **E-Commerce:** All e-commerce transactions pass through a series of intelligent services to provide performance, scalability, and availability within the overall e-commerce network design. Switches can play a role in the server and storage aspects of an e-commerce solution and provide switching between the edge router and the rest of the module.
- **Remote Access and VPN:** Terminates VPN traffic, forwarded by the Internet Connectivity module, from remote users and remote sites. It also initiates VPN connections to remote sites through the Internet connectivity module. Layer 2 switches provide Layer 2 connectivity for devices.

Summary

This topic summarizes the key points discussed in this lesson.

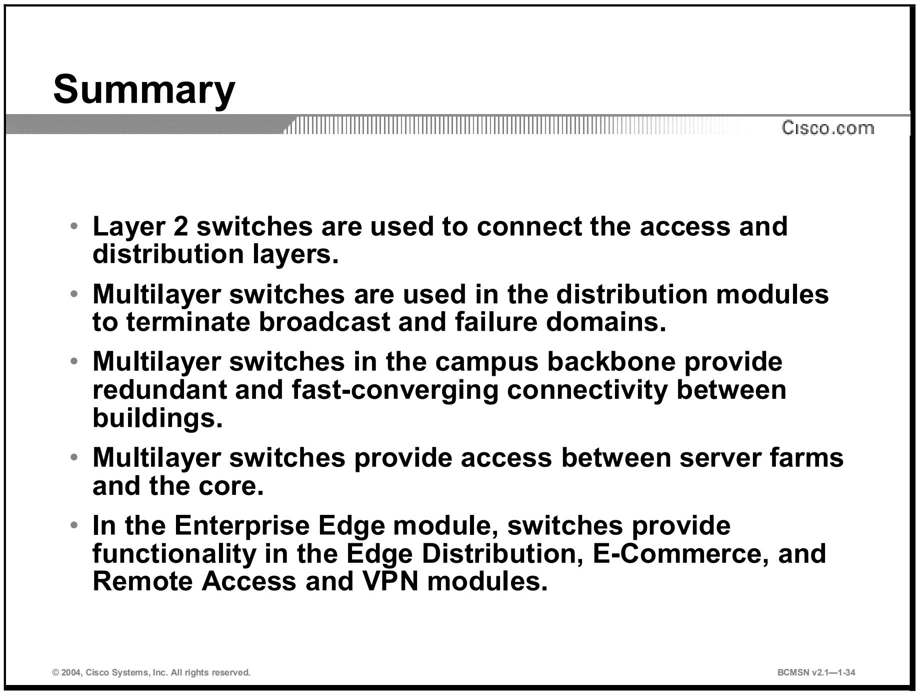

Summary

Cisco.com

- Layer 2 switches are used to connect the access and distribution layers.
- Multilayer switches are used in the distribution modules to terminate broadcast and failure domains.
- Multilayer switches in the campus backbone provide redundant and fast-converging connectivity between buildings.
- Multilayer switches provide access between server farms and the core.
- In the Enterprise Edge module, switches provide functionality in the Edge Distribution, E-Commerce, and Remote Access and VPN modules.

BCMSN v2.1—1-34

References

For additional information, refer to this resource:

- "Gigabit Campus Network Design—Principles and Architecture" at http://www.cisco.com/warp/public/cc/so/neso/lnso/cpso/gcnd_wp.htm

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Select the correct roles for multilayer switches in the Building Access submodule. (Choose two.)

A) connects end-user workstations to the Building Distribution submodule

B) provides load balancing within the Building Access submodule

C) switches in the Building Access submodule function at a lower cost

D) does not accept requests to the Campus Backbone and other submodules

Q2) Select the correct roles for multilayer switches in the Building Distribution submodule. (Choose two.)

A) switches in the Building Distribution submodule function at a lower cost

B) does not accept requests to the Campus Backbone and other submodules

C) frees the Building Access submodule from dedicating resources to routing and access control

D) provides aggregation of wiring closets and access control

Q3) Select the correct roles for multilayer switches in the Building Core submodule. (Choose two.)

A) provides redundant and fast-converging connectivity between buildings

B) does not accept requests to the other submodules

C) switches in the Building Core submodule function at a lower cost

D) routes and switches traffic as fast as possible from one module to another

Q4) Select the correct role for multilayer switches in the Server Farm module. (Choose two.)

A) connected to two different switches, enabling full redundancy and load sharing

B) provides access between the Server Farm module and the Building Access submodule

C) cross-connected with access layer switches

D) enables high reliability and availability of all servers

Q5) Select the correct roles for multilayer switches in the Edge Distribution module. (Choose two.)

A) performs the distribution function between Building Access and Building Core layers

B) plays a role in the server and storage aspects of an e-commerce solution

C) initiates VPN connections to remote sites through the Internet Connectivity module

D) provides Layer 3 connectivity for devices

Quiz Answer Key

Q1) A, B

Relates to: Identifying the Role of Multilayer Switches in the Building Access Submodule

Q2) A, C

Relates to: Identifying the Role of Multilayer Switches in the Building Distribution Submodule

Q3) A, D

Relates to: Identifying the Role of Multilayer Switches in the Building Core Submodule

Q4) A, D

Relates to: Identifying the Role of Multilayer Switches in the Server Farm Module

Q5) B, C

Relates to: Identifying the Role of Multilayer Switches in the Edge Distribution Module

Using Data-Link Technologies in an Enterprise Composite Network Model

Overview

The Enterprise Composite Network model makes clear delineations around specific operations that occur within a network. To support those defined operations, there must be an understanding about how to interconnect components within and between the functional areas and modules.

Relevance

To build a cost-effective and efficient multilayer switched network, it is critical to understand how a network design uses the various technologies available for a specific set of features and functionality.

Objectives

Upon completing this lesson, you will be able to:

- Select the best technology to interconnect modules in the Enterprise Composite Network model
- Select the best technology to interconnect the Building Access layer in the Enterprise Composite Network model
- Select the best technology to interconnect the Building Distribution layer in the Enterprise Composite Network model
- Select the best technology to interconnect the Building Core layer in the Enterprise Composite Network model

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Using Data Links to Interconnect Modules
- Using Data Links to Interconnect the Building Access Layer
- Using Data Links to Interconnect the Building Distribution Layer
- Using Data Links to Interconnect the Building Core Layer
- Summary
- Quiz

Using Data Links to Interconnect Modules

This topic identifies the best data link layer technology used to interconnect modules in the Enterprise Composite Network model.

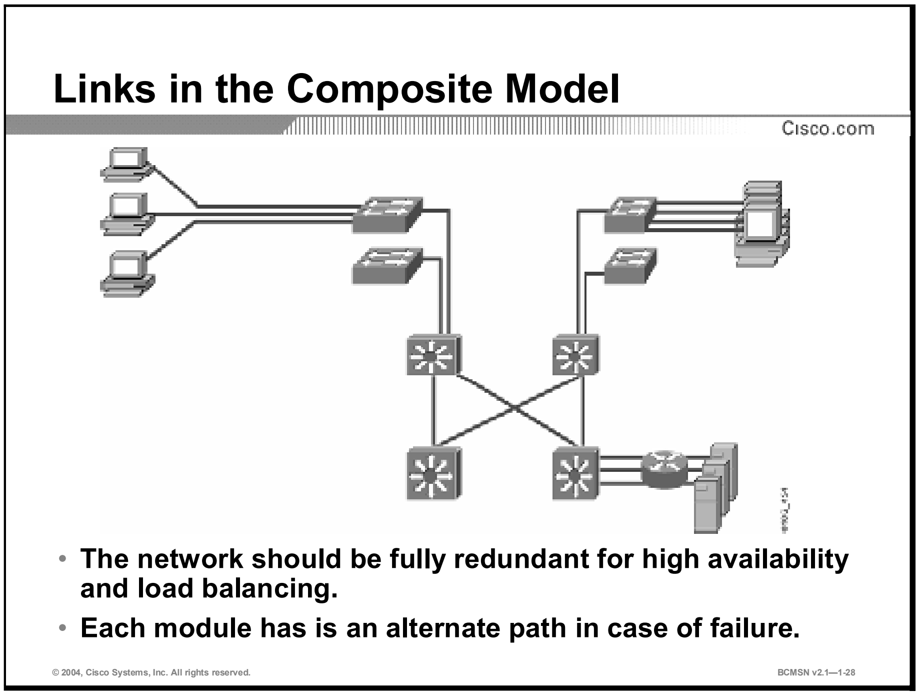

Links in the Composite Model

Cisco.com

- The network should be fully redundant for high availability and load balancing.
- Each module has is an alternate path in case of failure.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-28

The Enterprise Composite Network model applies some consistent concepts for providing a scalable and reliable network. One of those concepts is the idea that the network should be fully redundant for high availability and load balancing. This means that between each module there is an alternate path in case of failure. A Layer 2 spanning tree provides the mechanism by which redundant links are used for data communications. At Layer 3, the routing protocol ensures that redundant links are used.

Using multilayer switches allows for high-performance links and full duplex operations at a comparatively low cost per port. Ethernet supports 10 megabits of data per second, Fast Ethernet supports 100 megabits of data per second, Gigabit Ethernet supports up to 1000 megabits of data per second, and 10 Gigabit Ethernet supports up to 10,000 megabits of data per second. Links between modules should always be full duplex and capable of supporting the total aggregate amount of traffic coming from a particular module and the rest of the network.

The availability of powerful, affordable personal computers and workstations has driven the requirement for speed and availability in a campus network. In addition to existing applications, a new generation of multimedia, imaging, and database products can easily overwhelm a network.

Using Data Links to Interconnect the Building Access Layer

This topic identifies the best data link layer technology used to interconnect the Building Access layer in the Enterprise Composite Network model.

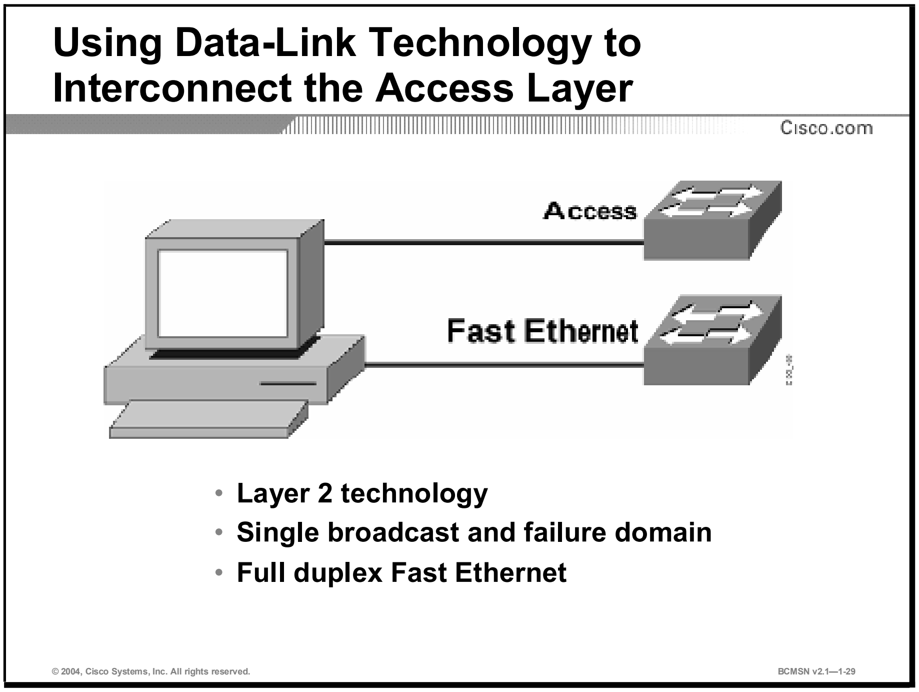

Using Data-Link Technology to Interconnect the Access Layer

Cisco.com

Access
Fast Ethernet

- Layer 2 technology
- Single broadcast and failure domain
- Full duplex Fast Ethernet

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-29

The Building Access submodule is usually a Layer 2-switched network segment and, as such, will be a single broadcast and failure domain. This is where segmentation is implemented to provide isolation from the negative effects of these domains.

Typical Building Access submodule connectivity requires transmission rates of 100 mbps (Fast Ethernet) and above. While the actual transmission rate at the Building Access submodule is important, even more vital is full duplex operation. Full duplex links from the end-user nodes into the switched network eliminate the possibility of collisions, allowing both a node and a switch to transmit and receive at the same time.

Another consideration at the Building Access submodule is in which module or functional area the access links exist. Since each module is comprised of a Building Access and Building Distribution submodule, there will be different considerations for the module-specific submodules.

Server farm access links should be multihomed into the servers themselves between different Building Access switches. Many server-grade Network Interface Cards (NICs) and operating systems allow for this type of high availability and even load balancing. The transmission rate on the NIC should be set at the maximum available rate on links between the server node and the access switch If possible, the server node should have Gigabit Ethernet. Server resources support many end users. When there is a one-to-many ratio in terms of network and resource access, the links should be greatly increased in bandwidth and reliability toward the resource to minimize the effects of any link failure to the server farm.

Conversely, end-user workstations do not typically warrant highly reliable links because a failure would only affect a single user. However, Building Access layer switches would require redundant links to reduce the impact of link failures for multiple users. Because there are multiple end-user workstations on a single access switch, Building Access layer switches should reflect a redundant and highly reliable configuration. A typical Building Access layer switch configuration would put dual uplinks into two separate switches located in the Building Distribution submodule. Because most end-user nodes typically run Fast Ethernet to the access switch, the uplinks from the Building Access submodule to the Building Distribution submodule should be Gigabit EtherChannel or Fast EtherChannel.

Using Data Links to Interconnect the Building Distribution Layer

This topic identifies the data link layer technology used to interconnect the Building Distribution layer in the Enterprise Composite Network model.

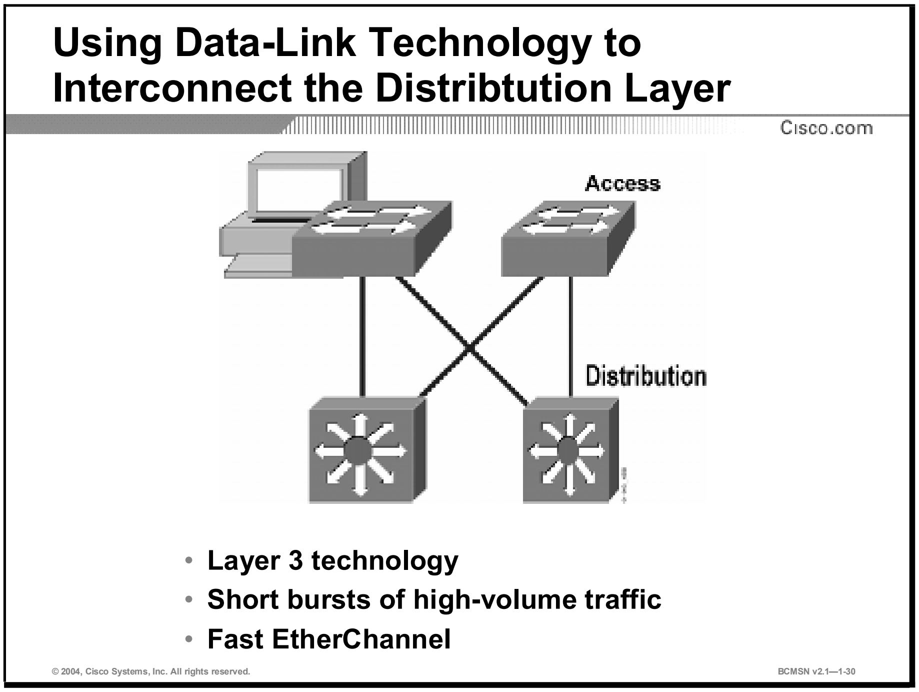

Using Data-Link Technology to Interconnect the Distribtution Layer

Cisco.com

- Layer 3 technology
- Short bursts of high-volume traffic
- Fast EtherChannel

BCMSN v2.1—1-30

Consistency, modularity, and a deterministic design are the keys to the Enterprise Composite Network model’s ability to enable scalability and reliability for enterprise networks.

At the Building Distribution submodule, the same theme is repeated. Two Building Distribution switches provide a network path for the Building Access submodule. These Building Distribution switches are dual-homed into the Campus Backbone submodule. The uplinks from the Building Distribution submodule into the Campus Backbone submodule must be of a bandwidth that is capable of handling the aggregate traffic from all supported Building Access submodules.

The Building Distribution submodule functions as a termination point for the Layer 2-switched environment of the Building Access submodules. Placing Layer 3 routing functionality at the Building Distribution submodules restricts broadcast and failure domains to the Building Access submodule.

The Building Distribution submodule is also responsible for any QoS and routing decisions that are to be made. Access control and optimal path determination are two key features of the Building Distribution submodule.

Using Data Links to Interconnect the Building Core Layer

This topic identifies the data link layer technology used to interconnect the core layer in the Enterprise Composite Network model.

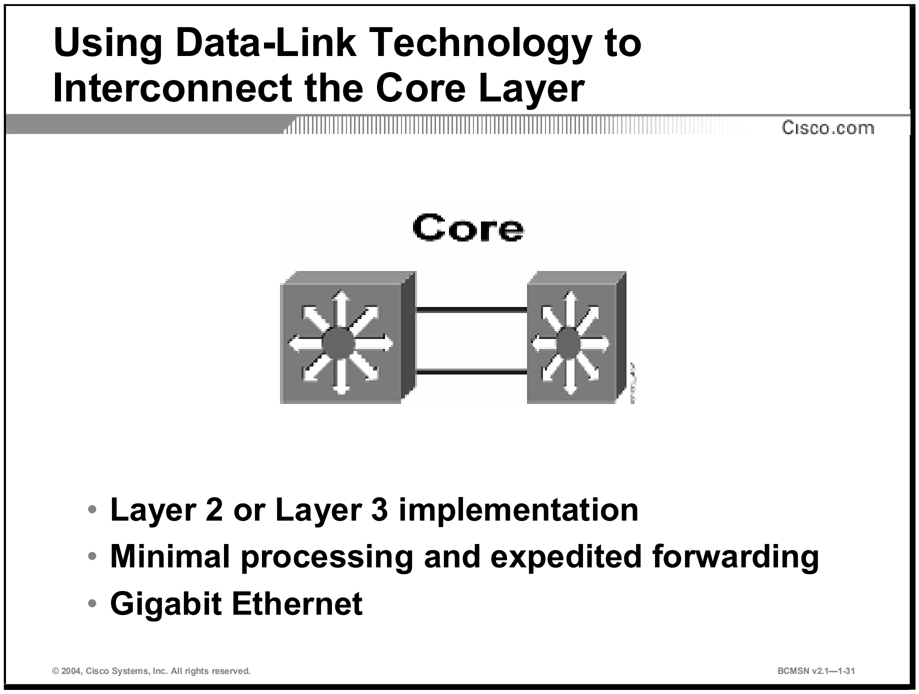

Using Data-Link Technology to Interconnect the Core Layer

Cisco.com

- Layer 2 or Layer 3 implementation
- Minimal processing and expedited forwarding
- Gigabit Ethernet

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-31

The Campus Backbone submodule is the aggregation point for all network communication within an enterprise network. The purpose of the Campus Backbone submodule is high-speed switching and forwarding between source and destination modules and submodules. The Campus Backbone submodule only handles those packet manipulations that are required for forwarding and queuing.

The Campus Backbone submodule can be either Layer 2 or Layer 3 implementations. Because this submodule ensures minimal processing and expedited forwarding, all links must have a bandwidth capable of accommodating the aggregate network traffic.

Summary

This topic summarizes the key points discussed in this lesson.

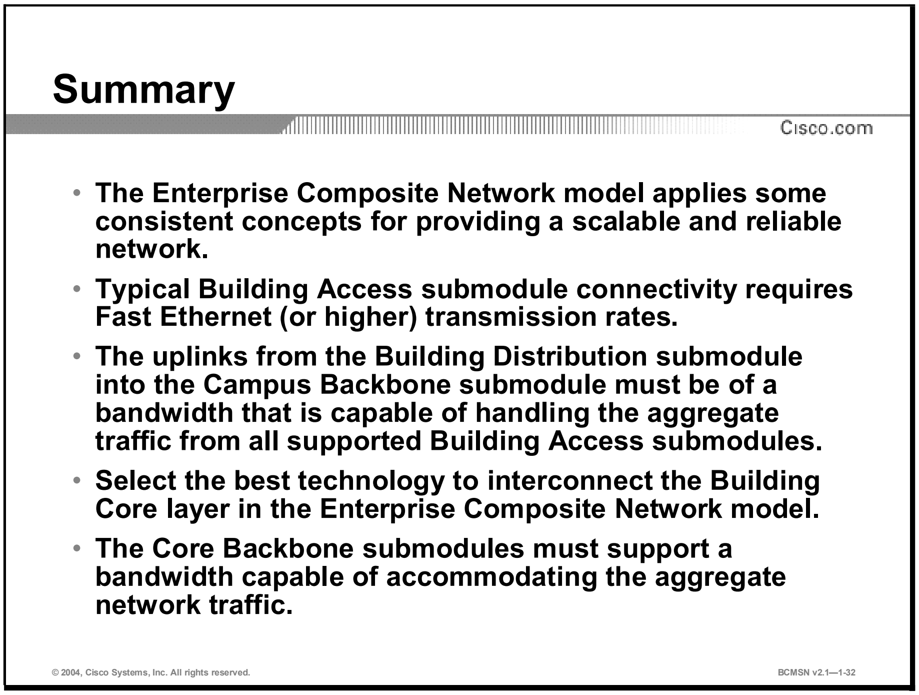

Summary

Cisco.com

- The Enterprise Composite Network model applies some consistent concepts for providing a scalable and reliable network.
- Typical Building Access submodule connectivity requires Fast Ethernet (or higher) transmission rates.
- The uplinks from the Building Distribution submodule into the Campus Backbone submodule must be of a bandwidth that is capable of handling the aggregate traffic from all supported Building Access submodules.
- Select the best technology to interconnect the Building Core layer in the Enterprise Composite Network model.
- The Core Backbone submodules must support a bandwidth capable of accommodating the aggregate network traffic.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—1-32

- The Enterprise Composite Network model applies some consistent concepts for providing a scalable and reliable network.
- Typical Building Access submodule connectivity requires Fast Ethernet (or higher) transmission rates.
- The uplinks from the Building Distribution submodule into the Campus Backbone submodule must be of a bandwidth that is capable of handling the aggregate traffic from all supported Building Access submodules.
- The Core Backbone submodules must support a bandwidth capable of accommodating the aggregate network traffic.

References

For additional information, refer to this resource:

- "Gigabit Campus Network Design—Principles and Architecture" at http://www.cisco.com/warp/public/cc/so/neso/lnso/cpso/gcnd_wp.htm

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 1-1: Getting Started with Catalyst Switches

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Select the best data link layer technologies to interconnect the modules in the Enterprise Composite Network model. (Choose two.)

A) Provides one path between modules to speed up performance.

B) Links between modules should be half duplex.

C) Network should be fully redundant.

D) Links between modules should be full duplex.

Q2) Select the best data link layer technology to interconnect the Building Access layer in the Enterprise Composite Network model. (Choose two.)

A) Links in the Server Farm submodule should be multihomed into the servers.

B) Submodule connectivity requires transmission rates of 10 mbps and above.

C) Does not require redundant links.

D) Provides access to many users.

Q3) Select the best data link layer technology to interconnect the Building Distribution layer in the Enterprise Composite Network model. (Choose two.)

A) Two Building Distribution switches provide a path into the network for the Building Access submodule.

B) The Building Distribution submodule acts as a termination point for the Building Access submodules in a Layer 2-switched environment.

C) The Building Distribution submodule is not responsible for any QoS and routing decisions.

D) The uplinks from the Building Distribution submodule into the Campus Backbone submodule can be of any bandwidth.

Q4) Select the best data link layer technology to interconnect the Building Core layer in the Enterprise Composite Network model. (Choose two.)

A) Ethernet

B) Gigabit Ethernet

C) Fast Ethernet

D) Gigabit EtherChannel

Quiz Answer Key

Q1) C, D

Relates to: Using Data Links to Interconnect Modules

Q2) A, D

Relates to: Using Data Links to Interconnect the Building Access Layer

Q3) A, B

Relates to: Using Data Links to Interconnect the Building Distribution Layer

Q4) B, D

Relates to: Using Data Links to Interconnect the Building Core Layer

Lesson Assessments

Overview

Use the lesson assessments here to test what you learned in this module. The correct answers and solutions are found in the Lesson Assessment Answer Key.

Outline

This section includes these assessments:

- Quiz 1-1: Using Switching and Routing in Ethernet Networks
- Quiz 1-2: Addressing Common Network Problems
- Quiz 1-3: Using Multilayer Switches in the Campus Infrastructure Module
- Quiz 1-4: Using Data-Link Technologies in an Enterprise Composite Network Model

Quiz 1-1: Using Switching and Routing in Ethernet Networks

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Match communication features to the correct communication type
- Select the issues that can occur in an unlayered network
- Select the correct features of a Layer 2 switch
- Select the issues that can occur in a Layer 2 network
- Select the correct features of routers
- Select the correct issues that can occur in a routed network
- Select the correct features of multilayer switches
- Identify the correct issue that can occur in a poorly designed network

Quiz

Answer these questions:

Q1) In your network, you have large broadcast and failure domains. What can be the result of this type of network? (Choose two.)

A) Boundaries and member devices are not always clearly defined.

B) Boundaries and member devices are always clearly defined.

C) The potential for miswiring increases.

D) The potential for miswiring decreases.

Q2) In your network, you are having throughput problems and Layer 3 bottlenecks. What would solve this problem—the introduction of a switch or a router?

A) router

B) switch

Q3) In your network, you find that most if not all of your network resources are consumed. With which communication type is this problem most likely to occur?

A) unicast

B) multicast

C) broadcast

Q4) In your network, you are having a lot of collision domains and a lot of hosts per segment. Introducing which device would solve this problem?

A) Layer 2 network

B) routed network

C) multilayer network

D) unlayered network

Q5) You are looking to upgrade your network, but finances are very tight. What type of network would you most likely NOT want to implement?

A) Layer 2 network

B) routed network

C) multilayer network

D) unlayered network

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Quiz 1-2: Addressing Common Network Problems

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Match a layer of a hierarchical network model with its correct function
- Select the correct features for the Enterprise Composite Network model
- Match the correct feature with the appropriate functional area
- Match the correct functional area with the appropriate feature
- Match the correct submodule with the appropriate function
- Select the correct Cisco components that support voice and data in the Enterprise Composite Network model

Quiz

Answer these questions:

Q1) Match the layer function with the correct location on the hierarchical network model diagram.

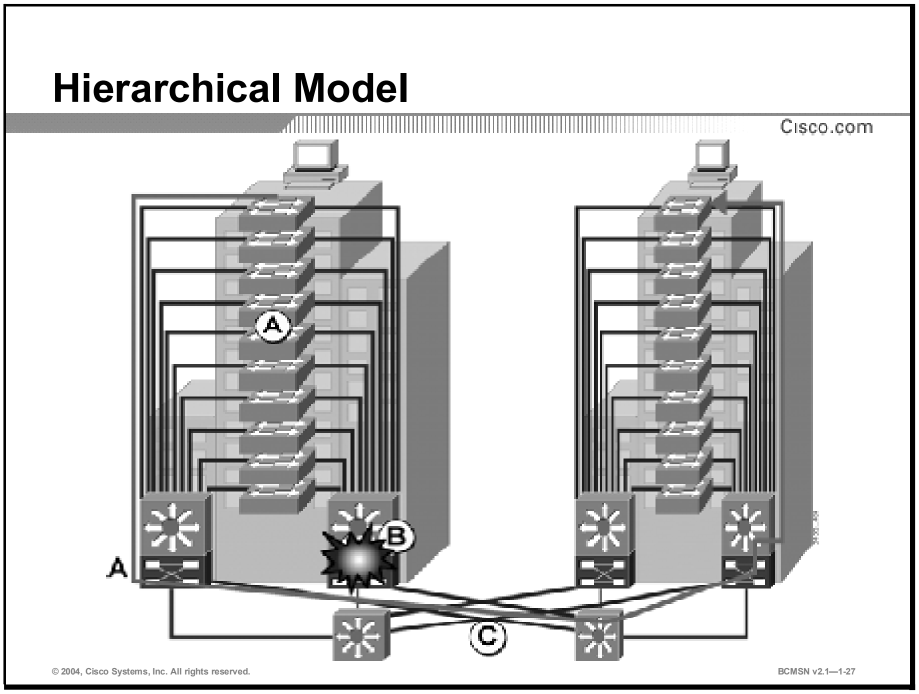

Hierarchical Model

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-27

_____ 1. policy-based connectivity

_____ 2. high-speed switching

_____ 3. local and remote workgroup access

Q2) Match the functional areas with the correct location on the Enterprise Composite Network model diagram.

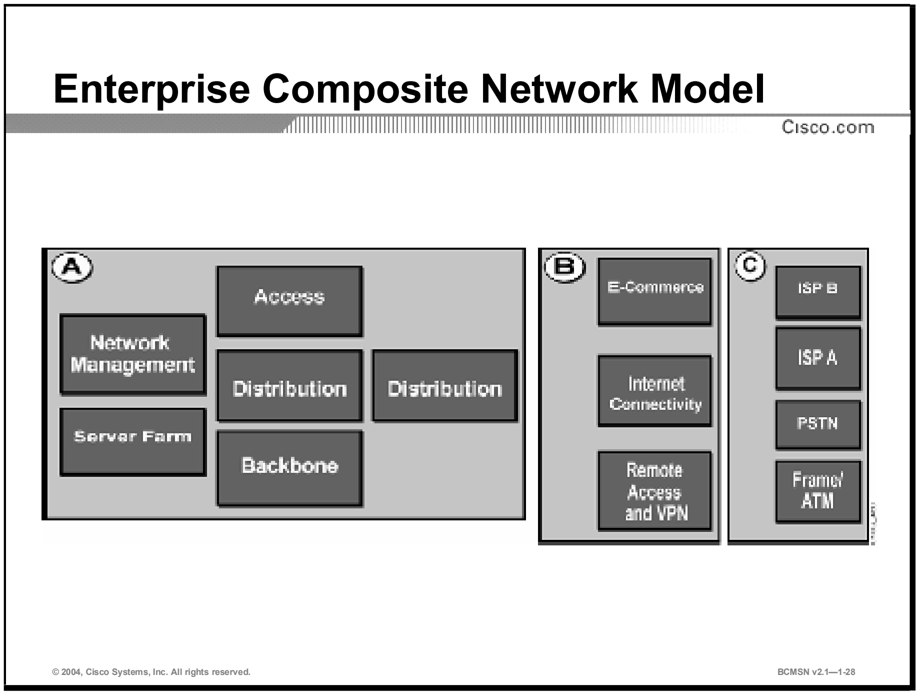

Enterprise Composite Network Model

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-28

_____ 1. contains network elements required for independent operation within a single campus

_____ 2. enables communications with other networks using WAN technologies

_____ 3. aggregates connectivity from the various elements at the edge of the enterprise network

Q3) Match the modules with the correct location on the Enterprise Campus infrastructure diagram, based on each area's function.

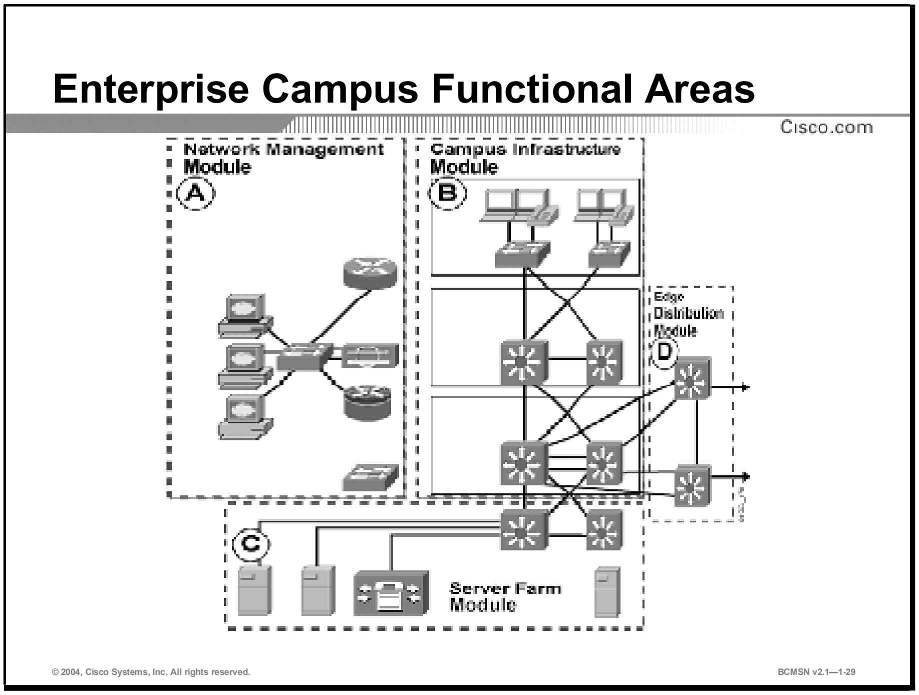

Enterprise Campus Functional Areas

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-29

_____ 1. contains e-mail and corporate servers providing application and print services

_____ 2. performs system logging and authentication

_____ 3. connects users within a campus with the Server Farm and Edge Distribution modules

_____ 4. routes traffic into the Campus Backbone submodule

Q4) Match the submodules with the correct location on the Campus Infrastructure module diagram, based on each submodule's function.

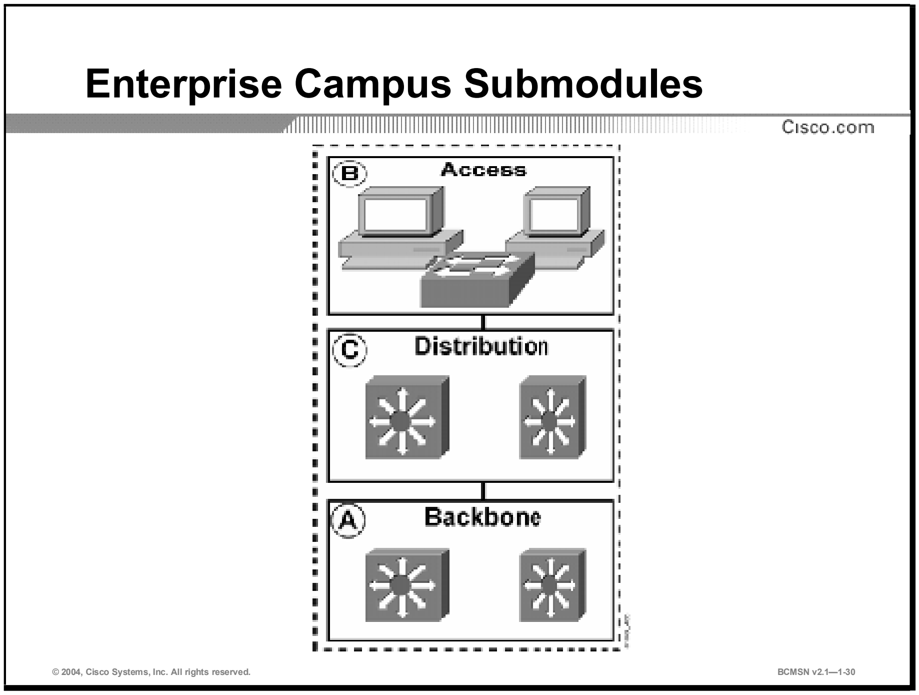

Enterprise Campus Submodules

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-30

_____ 1. performs routing, quality of service, and access control

_____ 2. provides redundant and fast-converging connectivity

_____ 3. contains end-user workstations, IP phones, and Layer 2 access switches

Q5) What Cisco AVVID component allows enterprises to make decisions not only about the business but also about networks and the technologies and applications that run on them?

A) vertical solutions

B) network solutions

C) network infrastructure

D) intelligent network services

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Quiz 1-3: Using Multilayer Switches in the Campus Infrastructure Module

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Select the correct role of multilayer switches in the Building Access submodule
- Select the correct role of multilayer switches in the Building Distribution submodule
- Select the correct role of multilayer switches in the Building Core submodule
- Select the correct role of multilayer switches in the Server Farm module
- Select the correct role of multilayer switches in the Edge Distribution module

Quiz

Answer this question:

Q1) Match the correct multilayer switch function with the appropriate submodule or module in which it would be performed.

_____ 1. Building Access submodule

_____ 2. Building Distribution submodule

_____ 3. Building Core submodule

_____ 4. Server Farm module

_____ 5. Edge Distribution module

A) supports a single broadcast or failure domain

B) routes and switches traffic between all modules and submodules

C) supports dual-homed access to application, file, print, e-mail, and DNS services

D) provides switches between the edge router and the rest of the module

E) terminates broadcast and failure domains, keeping them a manageable size

F) uses switches for high throughput functions

G) imposes routing, QoS, and security features

H) provides access between servers and storage devices by means of additional servers

Scoring

You have successfully completed the quiz for this lesson when you have correctly matched each multilayer switch function with the appropriate submodule or module.

Quiz 1-4: Using Data-Link Technologies in an Enterprise Composite Network Model

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Select the best technology to interconnect modules in the Enterprise Composite Network model
- Select the best technology to interconnect the Building Access layer in the Enterprise Composite Network model
- Select the best technology to interconnect the Building Distribution layer in the Enterprise Composite Network model
- Select the best technology to interconnect the Building Core layer in the Enterprise Composite Network model

Quiz

Answer this question:

Q1) Match the correct data-link technology with the appropriate location on the Enterprise Composite Network model.

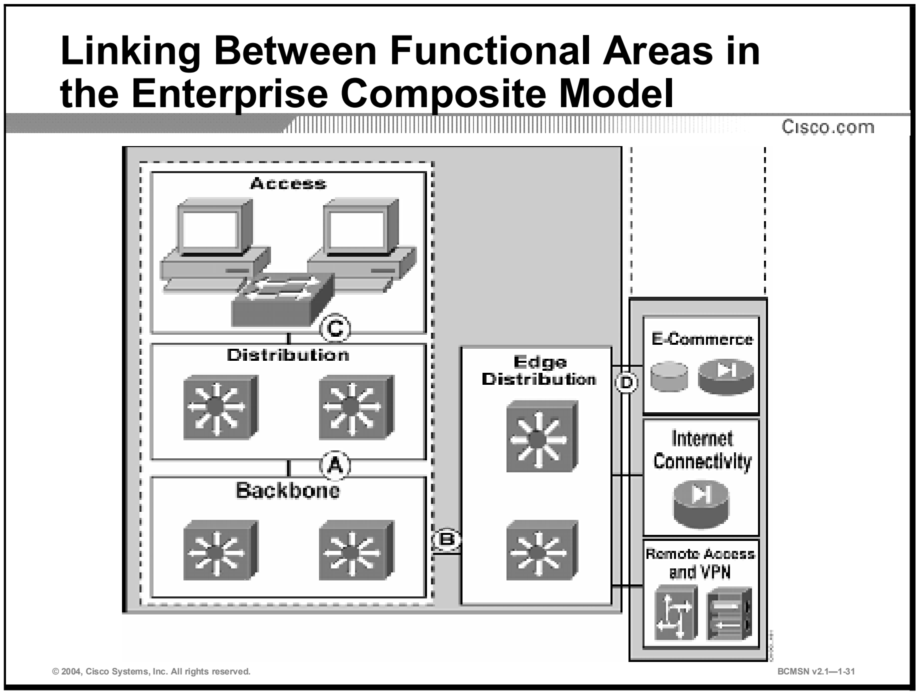

Linking Between Functional Areas in the Enterprise Composite Model

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—1-31

_____ 1. Layer 2, Fast Ethernet

_____ 2. Layer 3, gigabit, high-volume traffic

_____ 3. Layer 3, smaller bandwidth, Fast Ethernet

_____ 4. Layer 3, could handle short bursts of high-volume traffic

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Lesson Assessment Answer Key

Quiz 1-1: Using Switching and Routing in Ethernet Networks

Q1) A, C

Q2) B

Q3) B

Q4) A

Q5) B

Quiz 1-2: Addressing Common Network Problems

Q1) 1=B, 2=A, 3=C

Q2) 1=A, 2=C, 3=B

Q3) 1=C, 2=A, 3=B, 4=D

Q4) 1=C, 2=A, 3=B

Q5) B

Quiz 1-3: Using Multilayer Switches in the Campus Infrastructure Model

Q1) 1=A, 2=E, G; 3=B, F; 4=C, H; 5=D

Quiz 1-4: Using Data-Link Layer Technologies in an Enterprise Composite Network Model

Q1) 1=A, 2=C, 3=D, 4=B

Module 2

Configuring VLANs and VTP

Overview

A VLAN is a group of end stations with a common set of requirements, independent of their physical location. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a router.

Upon completing this module, you will be able to:

- Configure, verify, delete, and troubleshoot VLANs
- Resolve a problem on how to support multiple VLANS between two switches
- Apply the different trunking protocols to the Enterprise Composite Network model
- Configure, verify, and troubleshoot ISL and 802.1Q trunking links
- Propagate VLAN information in an Enterprise Composite Network

Outline

The module contains these components:

- Implementing VLANs
- Supporting Multiple VLANs Between Two Switches
- Defining Trunking Protocols
- Configuring Trunking Protocols
- Maintaining VLAN Consistency Across the Network
- Lesson Assessments

Implementing VLANs

Overview

Cisco Systems provides VLAN-capable solutions across its suite of internetworking switches and routers. Not only do VLANs solve many of the immediate problems associated with administrative changes, they also provide scalability, interoperability, and increased dedicated throughput.

Relevance

VLANs are an important aspect of switched networks. Configuring VLANs in switched networks improves overall performance.

Objectives

Upon completing this lesson, you will be able to:

- Select the characteristics that apply to VLANs
- Select the benefits that apply to VLANs
- Correctly identify the characteristics that apply to end-to-end and local VLANs
- Correctly select the conditions that no longer exist as a result of applying the Enterprise Composite Network model to a network design
- Identify the steps in the correct order necessary to create a generic Ethernet VLAN
- Correctly identify the steps necessary to create an Ethernet VLAN in global mode and the steps necessary to create an Ethernet VLAN in database mode
- Identify the steps in the correct order necessary to assign a port to a VLAN
- Identify which command is used to verify a VLAN configuration
- Identify which commands are used to verify a VLAN port configuration
- Select the correct set of commands that delete a VLAN in both global and database mode
- Select the recommended approach to troubleshoot VLANs to resolve a specific problem

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)
- Basic knowledge of the components that make up the Enterprise Composite Network model
- Basic knowledge of switch operations
- Basic knowledge of router operations
- Basic knowledge of how data is routed between devices

Outline

This lesson includes these topics:

- Overview
- Identifying the Features of a VLAN
- Identifying the Benefits of a VLAN
- Comparing VLAN Implementations
- Implementing VLANs in the Enterprise Composite Network Model
- Creating an Ethernet VLAN
- Creating a VLAN in Global and Database Mode
- Assigning Ports to a VLAN
- Verifying a VLAN Configuration
- Verifying a VLAN Port Configuration
- Deleting a VLAN
- Troubleshooting VLAN Operations
- Summary
- Quiz

Identifying the Features of a VLAN

This topic identifies the features of a VLAN.

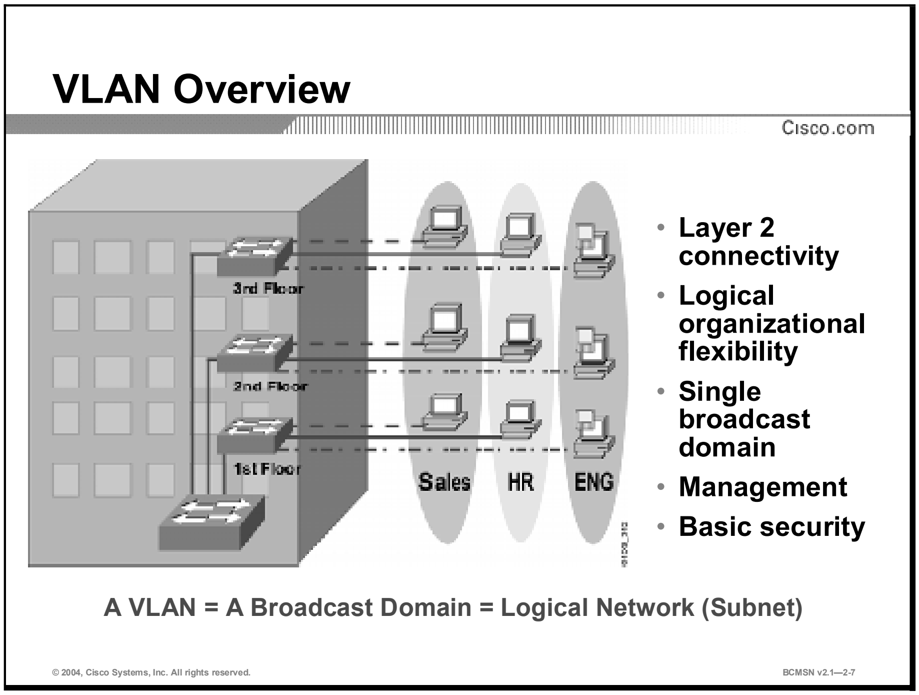

VLAN Overview

Cisco.com

- Layer 2 connectivity
- Logical organizational flexibility
- Single broadcast domain
- Management
- Basic security

A VLAN = A Broadcast Domain = Logical Network (Subnet)

BCMSN v2.1—2-7

A VLAN is a logical grouping of switch ports connecting to end-user workstations, servers, router interfaces, nodes, and even other LANs, with no regard to physical location.

A VLAN has these characteristics:

- All devices in a VLAN are members of the same broadcast domain. If a station transmits a broadcast, all other members of the VLAN will receive the broadcast. The broadcast is filtered from all ports or devices that are not members of the same VLAN.
- A VLAN is a logical subnet or segment made up of defined members. A VLAN is slightly different than a physical subnet. A physical subnet consists of devices on a physical cable segment. A logical subnet, or VLAN, consists of devices that have been configured as members of a VLAN. These devices can exist anywhere in the network. Just as you must have a router to communicate between physical subnets, you must also have a router to communicate between logical subnets, or VLANs.
- VLAN membership is always associated with a switch port.
- The most common type of VLAN is a local VLAN, often a wiring closet.
- End-to-end VLANs are defined throughout the entire network. An end-to-end VLAN may span several wiring closets or even several buildings. End-to-end VLANs are usually associated with a workgroup, such as a department or project team.

Identifying the Benefits of a VLAN

This topic identifies the benefits of a VLAN.

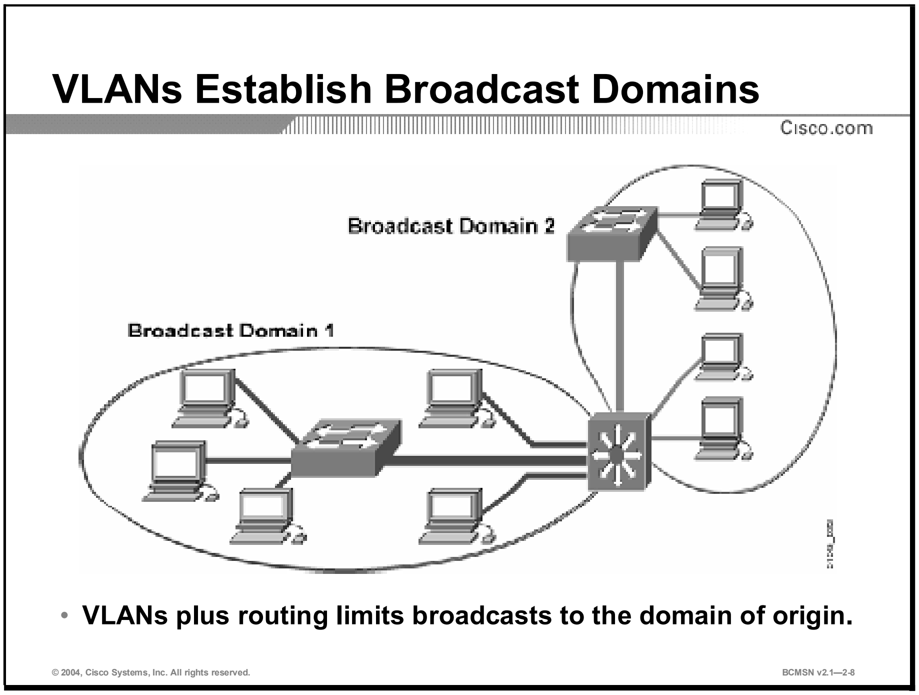

VLANs Establish Broadcast Domains

Cisco.com

- VLANs plus routing limits broadcasts to the domain of origin.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-8

VLANs solve many of the issues that arise from a switched campus network, including the following:

- **Efficient bandwidth utilization:** A VLAN solves the scalability problems found in large flat networks by dividing the network into smaller broadcast domains or subnets. All broadcast traffic is contained within the VLAN. For a packet to get to a different VLAN, it must be routed.
- **Security:** VLANs provide a very basic level of security by allowing you to segregate frames that contain sensitive or critical information from unauthorized users on separate VLANs.
- **Active redundant paths:** In combination with Spanning Tree Protocol (STP) tuning, multiple VLANs can be used as a method to distribute traffic across redundant paths that would otherwise be inactive.
- **Isolation of failure domains:** One of the most important reasons to implement VLANs is to reduce the impact of network problems. In a flat network, a faulty device, internetworking loop, or a broadcast-intensive application could potentially have an impact on the entire network to the point of a total failure. One of the most effective measures against such network failures is to properly segment the network with a router between segments. The router effectively prevents problems from being propagated to other segments or VLANs. This isolates the problem to a limited number of devices on a single VLAN.

Internally, the Catalyst switch implements VLANs by limiting data flooding or by forwarding to ports in the same VLAN. A Catalyst switch can retransmit a frame only to ports that belong to the same VLAN.

Normally, a port carries traffic only for the single VLAN to which it belongs. For a VLAN to span across multiple switches, a specialized link called a “trunk” is required to connect two switches. A trunk can carry traffic for multiple VLANs.

Comparing VLAN Implementations

This topic identifies those characteristics that apply to end-to-end and local VLANS.

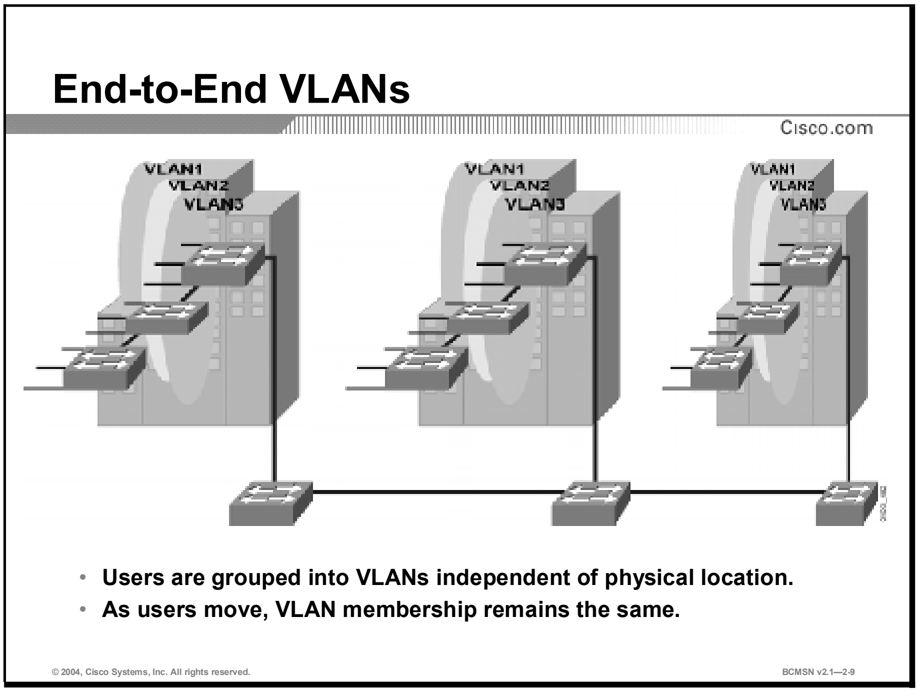

End-to-End VLANs

Cisco.com

- **Users are grouped into VLANs independent of physical location.**
- **As users move, VLAN membership remains the same.**

BCMSN v2.1—2-9

An end-to-end VLAN spans the entire switched network, while a local VLAN is restricted to a single switch.

An end-to-end VLAN network comprises these characteristics:

- Users are grouped into VLANs independent of physical location.
- As a user moves around the campus, VLAN membership of that user typically remains the same.
- Each VLAN has a common set of security and resource requirements for all members.

Starting in the wiring closet, 100 Mbps-dedicated Ethernet ports are provisioned for each user. Because users can be anywhere in the network, switches will be required to be aware of all VLANs and will receive flooded traffic even if they do not currently have any active ports in a particular VLAN.

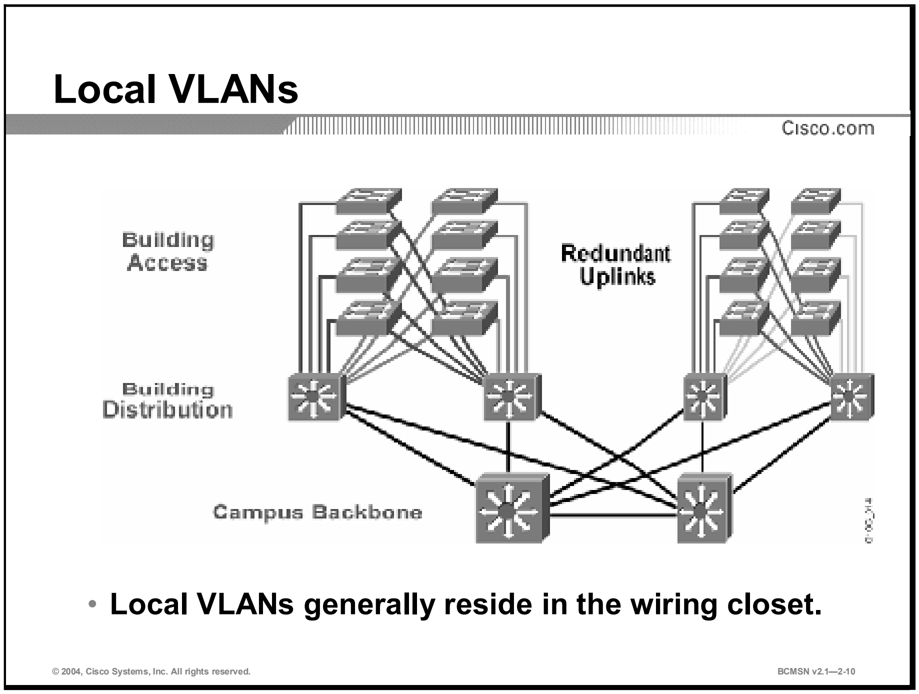

Local VLANs

Cisco.com

- **Local VLANs generally reside in the wiring closet.**

BCMSN v2.1—2-10

As corporations have moved to centralize their resources, end-to-end VLANs have become more difficult to maintain. Users might use many different resources, including many that are no longer in their VLAN.

VLANs are often created within physical boundaries rather than commonality boundaries. A local VLAN forces the user to cross a Layer 3 device to reach many of the resources. This design allows the network to provide for a deterministic, consistent method of accessing resources while maintaining the true performance and efficiency advantages of VLANs, including collision and broadcast domain segmentation.

Local VLANs, typically used in the Building Access submodule, are also easier to manage and conceptualize than VLANs spanning different areas of the network. A typical VLAN organization is to configure the minimum number of VLANs on a single access switch within a wiring closet, rather than having VLANs from multiple departments all configured on the same switch.

A single VLAN should not extend beyond the Building Distribution submodule. The local VLAN structure provides access into the network and Layer 3 connectivity. This structure allows users to move from one VLAN to another, without involving network administrators, and provides users with the same level of performance regardless of their location.

Troubleshooting local VLANs contained within a single area is much easier than troubleshooting a VLAN and modules spanning an entire functional area in the end-to-end design.

Implementing VLANs in the Enterprise Composite Network Model

This topic identifies the problems that are solved as a result of applying the Enterprise Composite Network model.

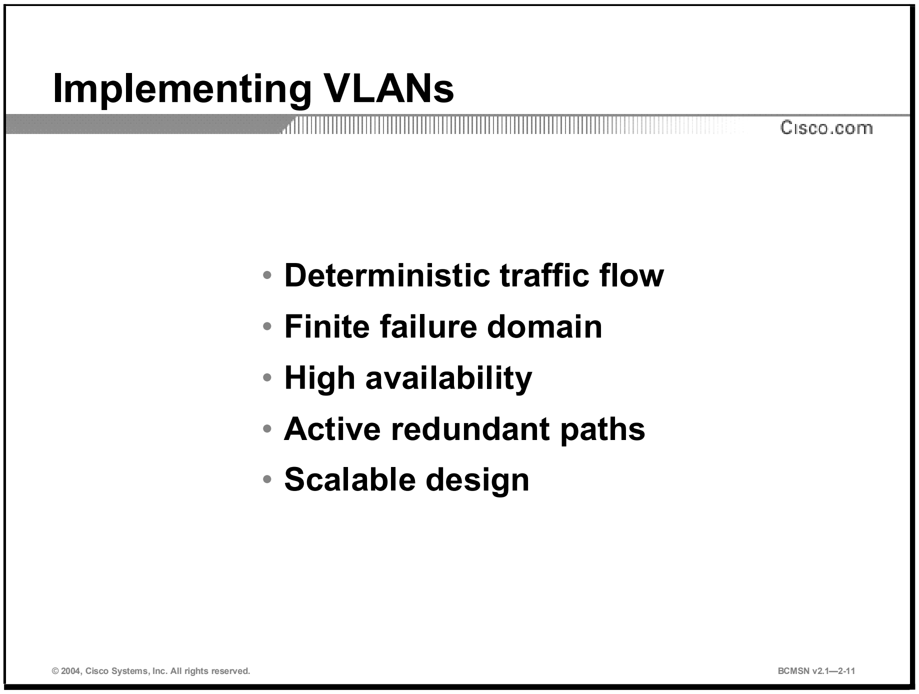

Implementing VLANs

Cisco.com

- Deterministic traffic flow
- Finite failure domain
- High availability
- Active redundant paths
- Scalable design

BCMSN v2.1—2-11

When you apply the Enterprise Composite Network model to a network, problems that can occur at Layer 2 are mitigated and the need for numerous or extensive VLANs is decreased. The Enterprise Composite Network model inherently provides these benefits to the network:

- **Deterministic traffic flow:** The simple layout and configuration gives a predictable traffic path. In the event of a failure, which cannot be mitigated by the redundancy features of the model, the network administrator can quickly isolate and rectify the problem.
- **Finite failure domain:** Because of the small size of the VLAN, any failures at Layer 2 that would affect the network are isolated to a subset of users.
- **High availability:** Redundancy has been designed into the model, rather than as an afterthought. This means that network availability is kept high, while maximizing the utilization of network equipment by sharing the traffic load across all the network resources.
- **Active redundant paths:** Instead of having inactive links in the network blocked by the Spanning Tree Protocol (STP), the model uses VLANs plus spanning-tree tuning to have forwarding for some VLANs on each link.
- **Scalable design:** The model uses a block approach, in which capacity can be added incrementally without requiring a new design effort. This design is sometimes referred to as a “cookie cutter” approach.

Creating an Ethernet VLAN

This topic identifies the steps to create an Ethernet VLAN.

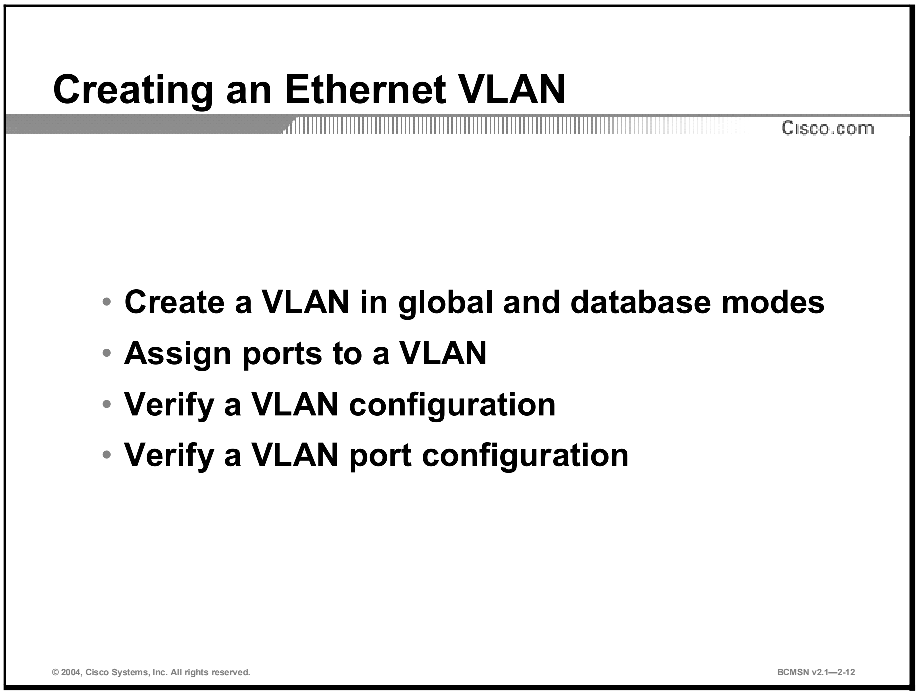

Creating an Ethernet VLAN

Cisco.com

- Create a VLAN in global and database modes
- Assign ports to a VLAN
- Verify a VLAN configuration
- Verify a VLAN port configuration

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-12

Here are the steps to create an Ethernet VLAN:

- Create the VLAN
- Assign the ports
- Verify the VLAN configuration
- Verify the VLAN port configuration

And, when needed:

- Delete a VLAN in global mode
- Delete a VLAN in database mode

Creating a VLAN in Global and Database Mode

This topic identifies the steps and the commands needed to create a VLAN in both global and database modes.

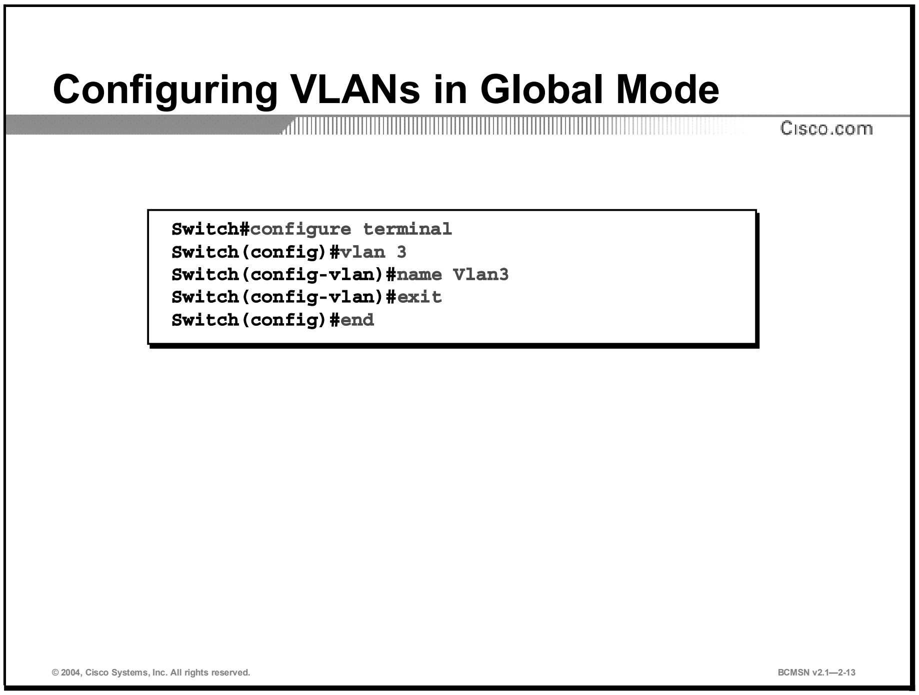

Configuring VLANs in Global Mode

Cisco.com

```
Switch#configure terminal
Switch(config)#vlan 3
Switch(config-vlan)#name Vlan3
Switch(config-vlan)#exit
Switch(config)#end
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-13

To create a new VLAN in global configuration mode, follow these steps:

Step	Action	Notes
1.	Enter global configuration mode. `Switch#`**`configure terminal`**	
2.	Create a new VLAN with a particular ID number. `Switch(config)#`**`vlan`** *`vlan_id`* `[`*`vlan_name`*`]`	Enter a VLAN ID and enter VLAN configuration mode. Enter a new VLAN ID to create a VLAN, or enter an existing VLAN ID to modify a VLAN. Valid values are 1-4094. (The precise range depends on the specific Catalyst switch, with 1 being the default.) Optionally, you can identify the VLAN with a name for management purposes.
3.	Name a VLAN. `Switch(config-vlan)#`**`name`** *`vlan_name`*	Name the VLAN.
4.	Exit VLAN configuration mode. `Switch(config-vlan)#`**`exit`**	After you have returned to privileged EXEC mode, the prompt will change back to `Switch#`.

Example: Creating a VLAN in Global Configuration Mode

This example shows how to create VLAN3 in global configuration mode:

```
Switch#configure terminal
Switch(config)#VLAN 3
Switch(config-vlan)#exit
```

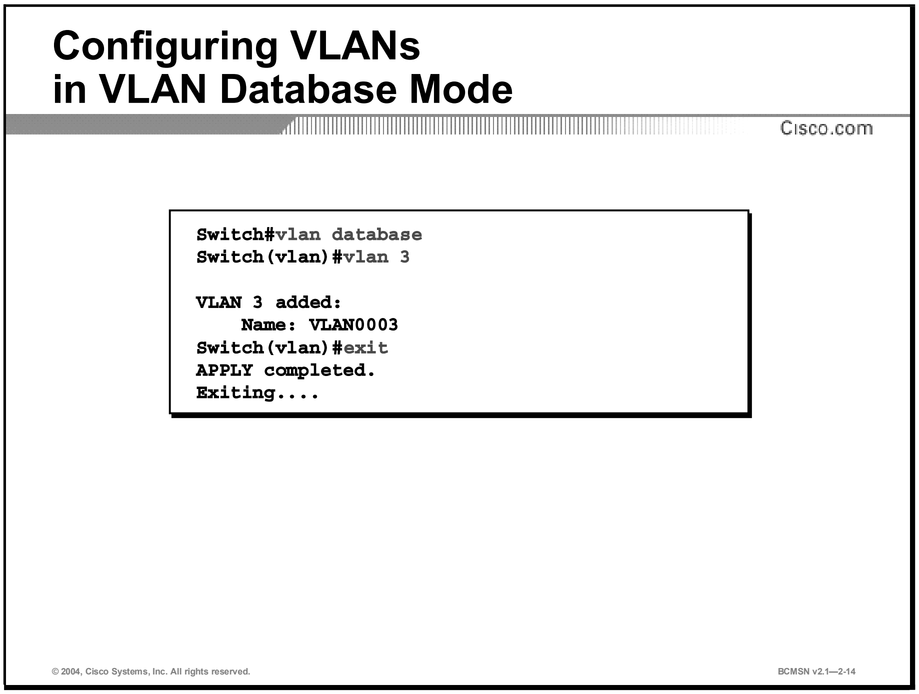

Configuring VLANs in VLAN Database Mode

Cisco.com

```
Switch#vlan database
Switch(vlan)#vlan 3

VLAN 3 added:
    Name: VLAN0003
Switch(vlan)#exit
APPLY completed.
Exiting....
```

BCMSN v2.1—2-14

Use the **vlan database** privileged EXEC command to enter VLAN configuration mode. From this mode, you can add, delete, and modify VLAN configurations for normal-range VLANs.

Note The VLAN configuration mode is different from other modes because it is session-oriented. When you add, delete, or modify VLAN parameters, the changes are not applied until you exit the session by entering the **apply** or **exit** command. When the changes are applied, the VLAN Trunk Protocol (VTP) configuration version is incremented. You can also *not* apply the changes to the VTP database by entering **abort**.

You can use the VLAN database configuration commands to configure VLANs 1 to 1005. To configure extended-range VLANs (VLAN IDs 1006 to 4094), use the **vlan** global configuration command to enter VLAN configuration mode. You can also configure VLAN IDs 1 to 1005 by using the **vlan** global configuration command.

To create a new VLAN in VLAN database configuration mode, follow these steps:

Step	Action	Notes
1.	Enter VLAN configuration mode. `Switch#`**`VLAN database`**	
2.	Create a new VLAN with a particular ID number. `Switch(vlan#`**`VLAN`** *`VLAN_id`*	Create a VLAN with the specified ID number.
3.	Exit VLAN database configuration mode. `Switch(vlan)#`**`exit`**	After you have returned to privileged EXEC mode, the prompt will change back to **`switch#.`**

Example: Creating a VLAN in VLAN Database Mode

This example shows how to create VLAN3 in VLAN database mode:

```
Switch#VLAN database
Switch(vlan)#VLAN 3
VLAN 3 added:
    Name: VLAN0003
Switch(vlan)#exit
APPLY completed.
Exiting....
Switch#
```

Note Cisco recommends using global configuration mode to define VLANs.

Assigning Ports to a VLAN

This topic identifies the steps involved in assigning a port to a particular VLAN.

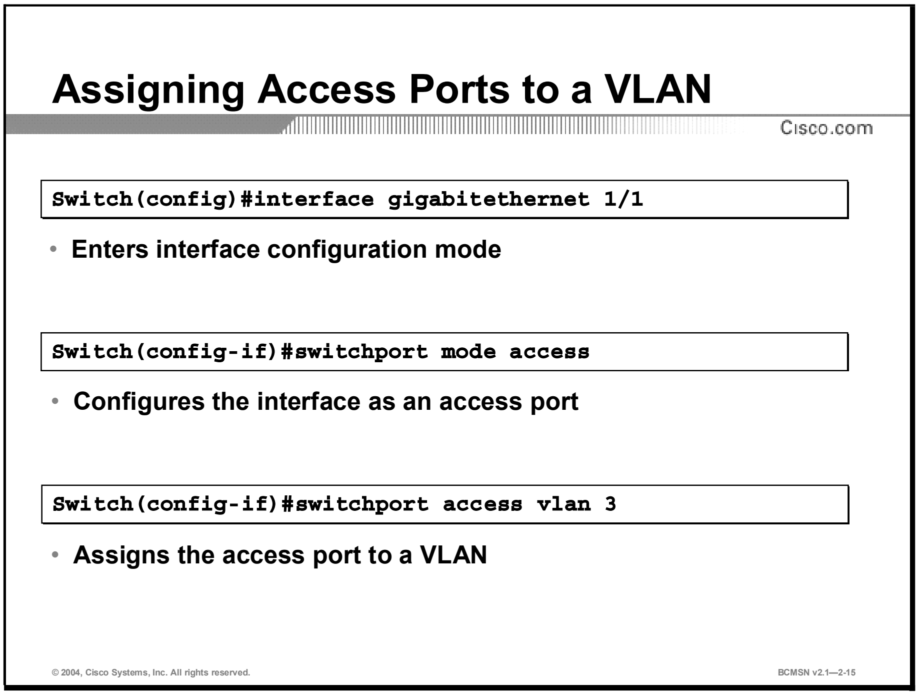

Assigning Access Ports to a VLAN

Cisco.com

```
Switch(config)#interface gigabitethernet 1/1
```

- Enters interface configuration mode

```
Switch(config-if)#switchport mode access
```

- Configures the interface as an access port

```
Switch(config-if)#switchport access vlan 3
```

- Assigns the access port to a VLAN

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-15

After a VLAN has been defined, switch ports need to be assigned to it. To assign an access switch port to a previously created VLAN, follow these steps:

Step	Action	Notes
1.	From global configuration mode, enter configuration mode for the particular port you want to add to the VLAN. `Switch(config)#`**`interface {fastethernet	gigabitethernet}`** *`slot/port`*	
2.	Specify the port as an access port. `Switch(config-if)#`**`switchport mode access`**	Use the **switchport** interface configuration command with no parameters to put an interface that is in Layer 3 mode into Layer 2 mode for Layer 2 configuration. Use the **no switchport** command (without parameters) to set the interface to the routed-interface status and to erase all Layer 2 configurations. You must use this command before assigning an IP address to a routed port.
3.	Place the port in a particular VLAN. `Switch(config-if)#`**`switchport access VLAN`** *`VLAN_id`*	If the VLAN you specify does not exist, the port will not become operational until you create the VLAN.

Note If an interface is to be configured as a Layer 3 interface, you must first enter the **switchport** command with no keywords to configure the interface as a Layer 2 port. Then you can enter additional **switchport** commands with keywords.

Example: Assigning an Access Port to a VLAN

This example shows how to configure the Fast Ethernet interface 5/6 as an access port in VLAN 200:

```
Switch#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)#interface fastethernet 5/6
Switch(config-if)#switchport mode access
Switch(config-if)#switchport access VLAN 200
Switch(config-if)#end
Switch#exit
```

Verifying a VLAN Configuration

This topic identifies the commands used to verify the VLAN configuration.

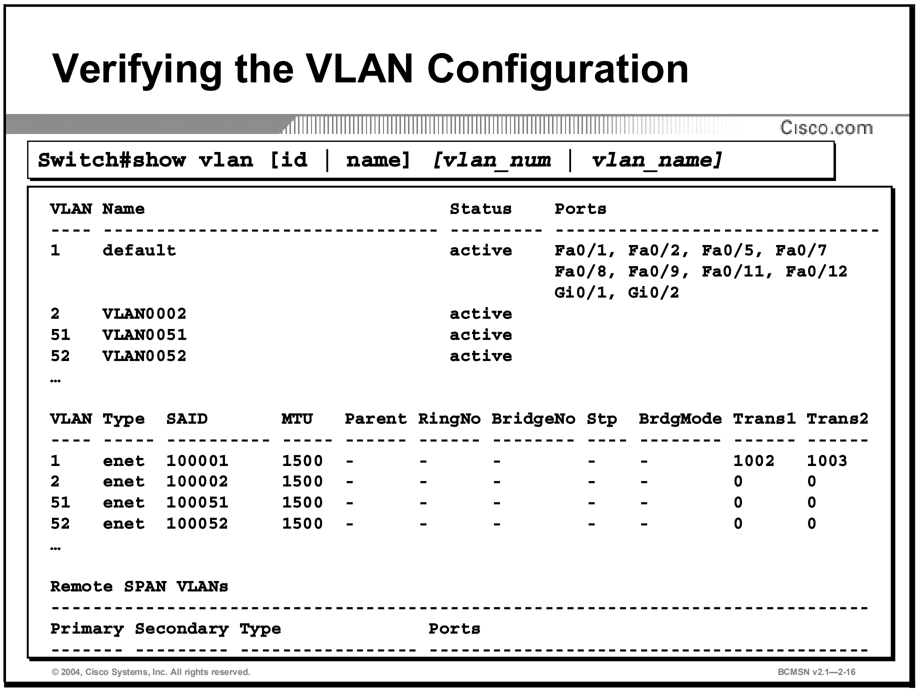

Verifying the VLAN Configuration

```
Switch#show vlan [id | name] [vlan_num | vlan_name]
```

```
VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/5, Fa0/7
                                                Fa0/8, Fa0/9, Fa0/11, Fa0/12
                                                Gi0/1, Gi0/2
2    VLAN0002                         active
51   VLAN0051                         active
52   VLAN0052                         active
…

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
---- ----- ---------- ----- ------ ------ -------- ---- -------- ------ ------
1    enet  100001     1500  -      -      -        -    -        1002   1003
2    enet  100002     1500  -      -      -        -    -        0      0
51   enet  100051     1500  -      -      -        -    -        0      0
52   enet  100052     1500  -      -      -        -    -        0      0
…

Remote SPAN VLANs
------------------------------------------------------------------------------
Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------------
```

BCMSN v2.1—2-16

The **show vlan** command from privileged EXEC mode displays information about a particular VLAN. The fields in the **show vlan** command output are as follows:

Field	Description
VLAN	VLAN number
Name	Name, if configured, of the VLAN
Status	Status of the VLAN (active or suspend)
Ports	Ports that belong to the VLAN
Type	Media type of the VLAN
SAID	Security association ID value for the VLAN
MTU	Maximum transmission unit size for the VLAN
Parent	Parent VLAN, if one exists
RingNo	Ring number for the VLAN, if applicable
BrdgNo	Bridge number for the VLAN, if applicable
Stp	Spanning Tree Protocol type used on the VLAN
BrdgMode	Bridging mode for this VLAN
Trans1	Translation bridge 1
Trans2	Translation bridge 2
AREHops	Maximum number of hops for all-routes explorer frames
STEHops	Maximum number of hops for spanning tree explorer frames

Example: Displaying Information about a VLAN by Number

This example shows how to display information about a VLAN identified by number:

```
Switch#show VLAN id 3

VLAN Name                             Status    Ports

---- -------------------------------- --------- --------------
-----------------

3    VLAN0003                         active

VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp
BrdgMode Trans1 Trans2

---- ----- ---------- ----- ------ ------ -------- ---- ------
-- ------ ------

3    enet  100003     1500  -      -      -        -    -
0      0

Primary Secondary Type              Interfaces

------- --------- ----------------- --------------------------
-----------------

Switch#
```

Example: Displaying Information about a VLAN by Name

This example shows how to display information about a VLAN identified by name:

```
Switch#show VLAN name VLAN0003

VLAN Name                             Status    Ports

---- -------------------------------- --------- --------------
-------

3    VLAN0003                         active


VLAN Type  SAID       MTU   Parent RingNo BridgeNo Stp  Trans1
Trans2

---- ----- ---------- ----- ------ ------ -------- ---- ------
------

3    enet  100003     1500  -      -      -        -    0
0

Switch#
```

Verifying a VLAN Port Configuration

This topic identifies the commands used to verify a VLAN port configuration.

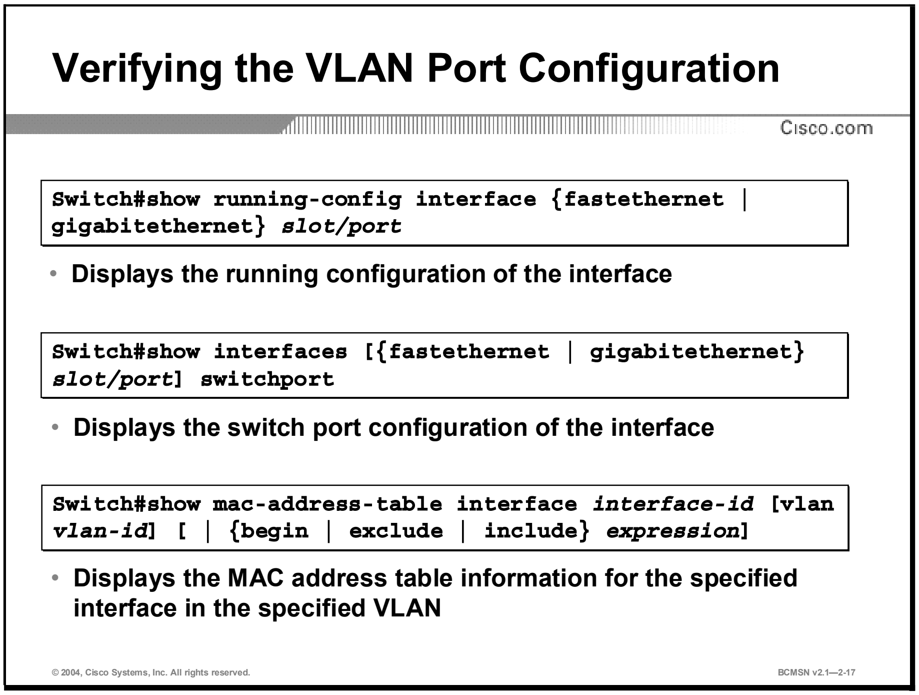

Verifying the VLAN Port Configuration

Cisco.com

```
Switch#show running-config interface {fastethernet |
gigabitethernet} slot/port
```

- Displays the running configuration of the interface

```
Switch#show interfaces [{fastethernet | gigabitethernet}
slot/port] switchport
```

- Displays the switch port configuration of the interface

```
Switch#show mac-address-table interface interface-id [vlan
vlan-id] [ | {begin | exclude | include} expression]
```

- Displays the MAC address table information for the specified interface in the specified VLAN

BCMSN v2.1—2-17

You can display information about a particular interface in general with the **show running-config** command, or the specific switchport information with the **show interfaces** command. The **show interfaces** *interface* **switchport** command displays switchport characteristics, including traffic suppression levels set on the interface.

The **show mac-address-table interface** command displays the MAC address table information for the specified interface in the specified VLAN.

Example: Displaying Information About an Interface

This example shows how to display information about a particular interface:

```
Switch# show running-config interface fastethernet 5/6
Building configuration...
!
Current configuration :33 bytes
interface FastEthernet 5/6
 switchport access vlan 200
 switchport mode access
end
```

Example: Displaying Detailed Switch Port Information

This example shows how to display detailed switch port information about an interface:

```
Switch#show interface gigabitEthernet 0/1 switchport
Name: Gi0/1
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001

Protected: false
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled

Broadcast Suppression Level: 100
Multicast Suppression Level: 100
Unicast Suppression Level: 100
```

Example: Displaying MAC Address Table Information

This example shows how to display MAC address table information for a specific interface in VLAN1:

```
Switch#show mac-address-table interface gigabitEthernet 0/1
VLAN 1
          Mac Address Table
------------------------------------------


VLAN    Mac Address       Type       Ports
----    -----------       ----       -----
   1    0008.2199.2bc1    DYNAMIC    Gi0/1
Total Mac Addresses for this criterion: 1
```

Deleting a VLAN

This topic identifies how to delete a VLAN in both global and database modes.

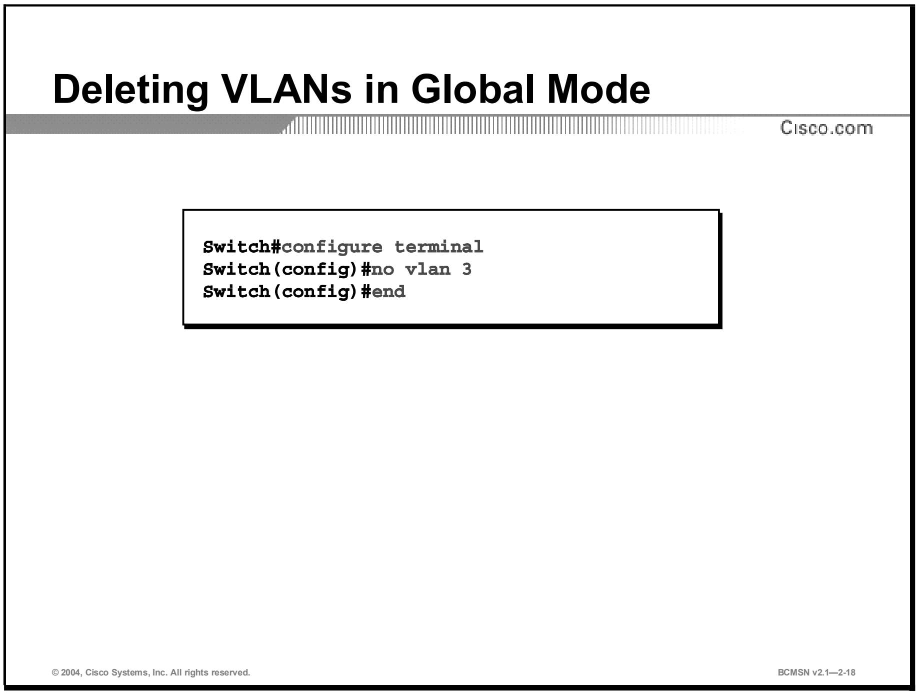

Deleting VLANs in Global Mode

Cisco.com

```
Switch#configure terminal
Switch(config)#no vlan 3
Switch(config)#end
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-18

To delete a VLAN in global configuration mode, follow these steps:

Step	Action	Notes
1.	Enter global configuration mode. `Switch#`**`configure terminal`**	
2.	Delete the VLAN with a particular ID number. `Switch(config)#`**`no vlan`** *`VLAN_id`*	Delete the VLAN with the specified ID number.
3.	Exit configuration mode. `Switch(config)#`**`end`**	After you have returned to privileged EXEC mode, the prompt will change back to `Switch#.`

Caution Deleting a VLAN will cause those switchports still assigned to be members of an **inactive** VLAN. This means that each switchport is isolated and will not be usable until reassigned to a different VLAN.

Example: Deleting a VLAN in Global Configuration Mode

This example shows how to delete VLAN3 in global configuration mode:

```
Switch#configure terminal
Switch(config)#no VLAN 3
Switch(config)#end
```

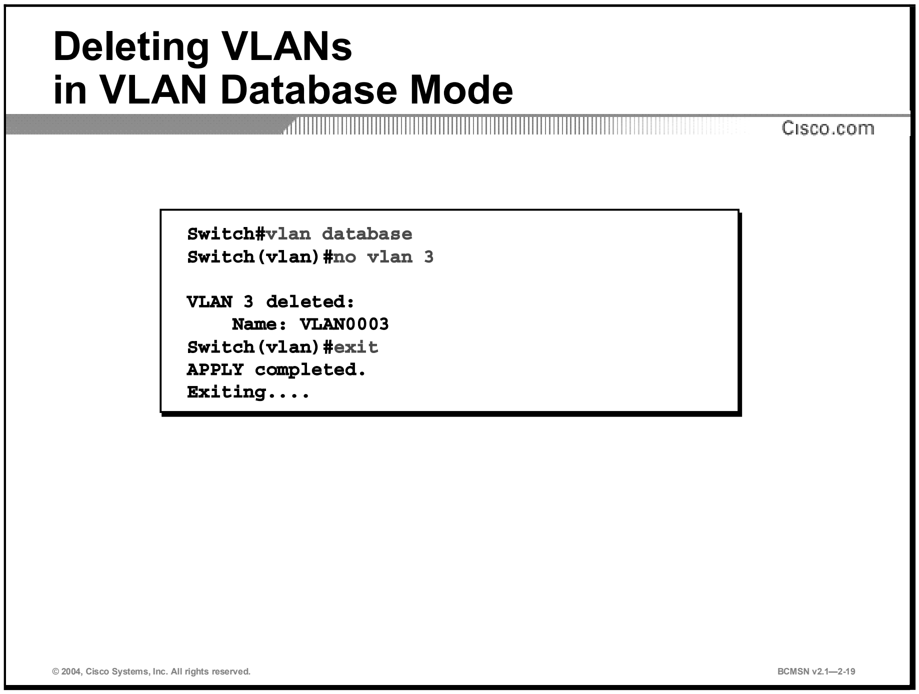

Deleting VLANs in VLAN Database Mode

Cisco.com

```
Switch#vlan database
Switch(vlan)#no vlan 3

VLAN 3 deleted:
    Name: VLAN0003
Switch(vlan)#exit
APPLY completed.
Exiting....
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-19

To delete an existing VLAN in VLAN database configuration mode, follow the steps:

Step	Action	Notes
1.	Enter VLAN configuration mode. `switch#`**`VLAN database`**	
2.	Delete a VLAN with a particular ID number. `switch(vlan)#`**`no VLAN`** *`VLAN_id`*	Delete the VLAN with the specified ID number.
3.	Exit VLAN database configuration mode. `switch(vlan)#`**`exit`**	After you have returned to privileged EXEC mode, the prompt will change back to `switch#`.

Troubleshooting VLAN Operations

This topic discusses troubleshooting VLANs.

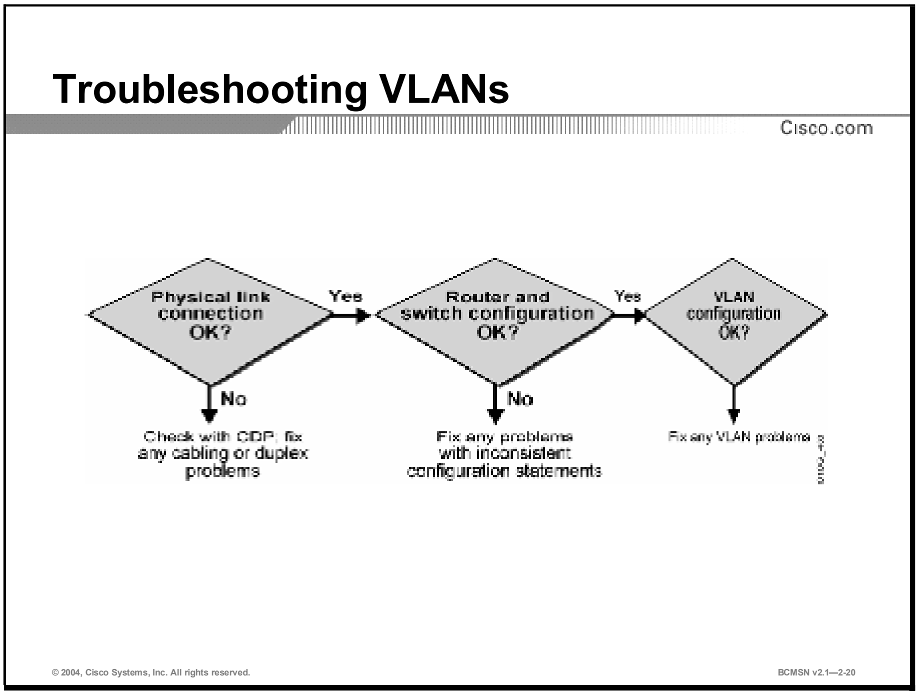

Troubleshooting VLANs

Cisco.com

Physical link connection OK?

Yes

Router and switch configuration OK?

Yes

VLAN configuration OK?

No

Check with CDP; fix any cabling or duplex problems

No

Fix any problems with inconsistent configuration statements

Fix any VLAN problems

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-20

To troubleshoot VLANs, you should check the physical connections, switch configuration, and VLAN configuration. Sometimes high-level VLAN problems can occur with a switch. The symptoms of the problem and the possible action plans may help you identify and solve the problems.

When faced with poor throughput problems, check to see what type of errors exist. There could be a bad adapter card. Combinations of frame check sequence (FCS) and alignment errors and runts generally point to a duplex mismatch; the usual culprit is the autonegotiation between devices or a mismatched setting between the two sides of a link. Use the Cisco Discovery Protocol (CDP) to report a half-duplex or full-duplex mismatch error. Consider these questions:

- Is the problem on the local or remote side of the link? Remember, a minimum number of switch ports are involved in a link.
- What path is the packet taking?

If you see that the number of collisions in output from a **show interface** command is increasing rapidly, the problem may be an overloaded link.

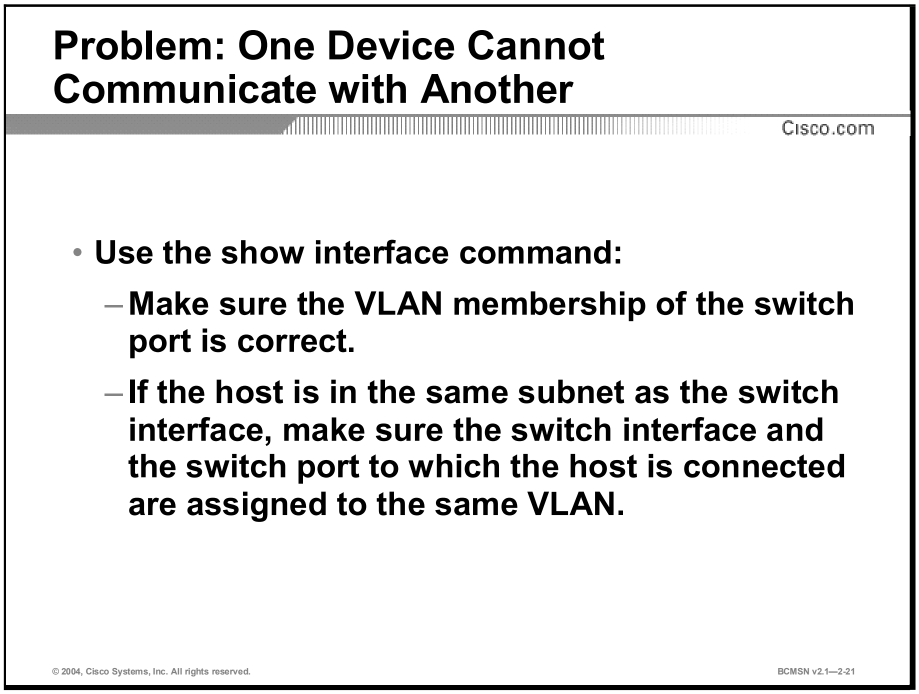

Problem: One Device Cannot Communicate with Another

Cisco.com

- Use the show interface command:
  - Make sure the VLAN membership of the switch port is correct.
  - If the host is in the same subnet as the switch interface, make sure the switch interface and the switch port to which the host is connected are assigned to the same VLAN.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-21

Problem: One device cannot communicate with another device.

Suggested solutions to the problem:

- Make sure the VLAN membership of the switch port is correct by using the **show interface switchport** command. Make sure the assigned VLAN is **active.**
- If the host is in the same subnet as the switch interface, make sure the switch interface and the switch port to which the host is connected are assigned to the same VLAN. Use the **show interface** command.

Summary

This topic summarizes the key points discussed in this lesson.

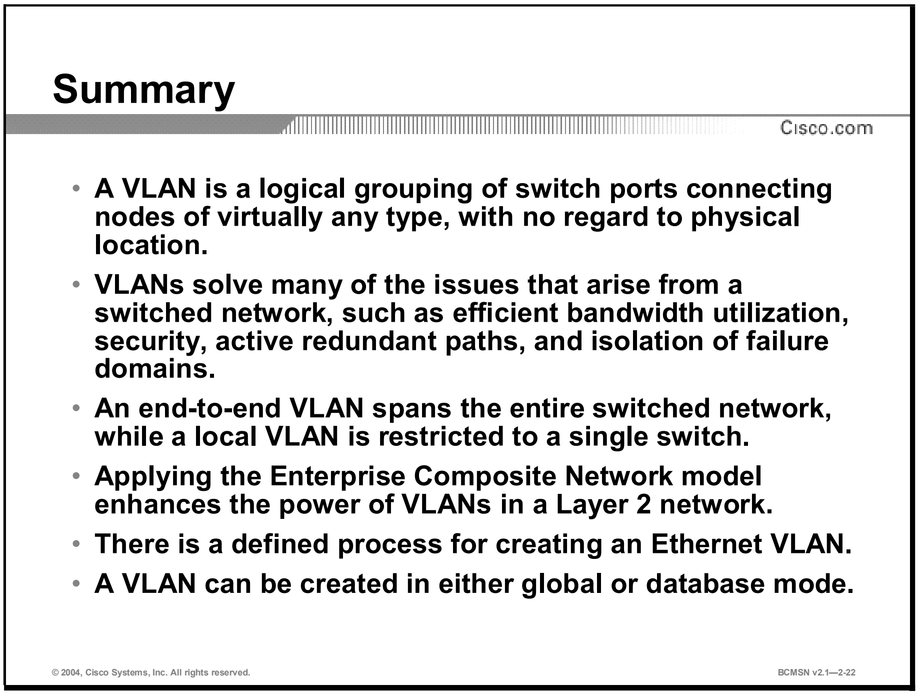

Summary

Cisco.com

- A VLAN is a logical grouping of switch ports connecting nodes of virtually any type, with no regard to physical location.
- VLANs solve many of the issues that arise from a switched network, such as efficient bandwidth utilization, security, active redundant paths, and isolation of failure domains.
- An end-to-end VLAN spans the entire switched network, while a local VLAN is restricted to a single switch.
- Applying the Enterprise Composite Network model enhances the power of VLANs in a Layer 2 network.
- There is a defined process for creating an Ethernet VLAN.
- A VLAN can be created in either global or database mode.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—2-22

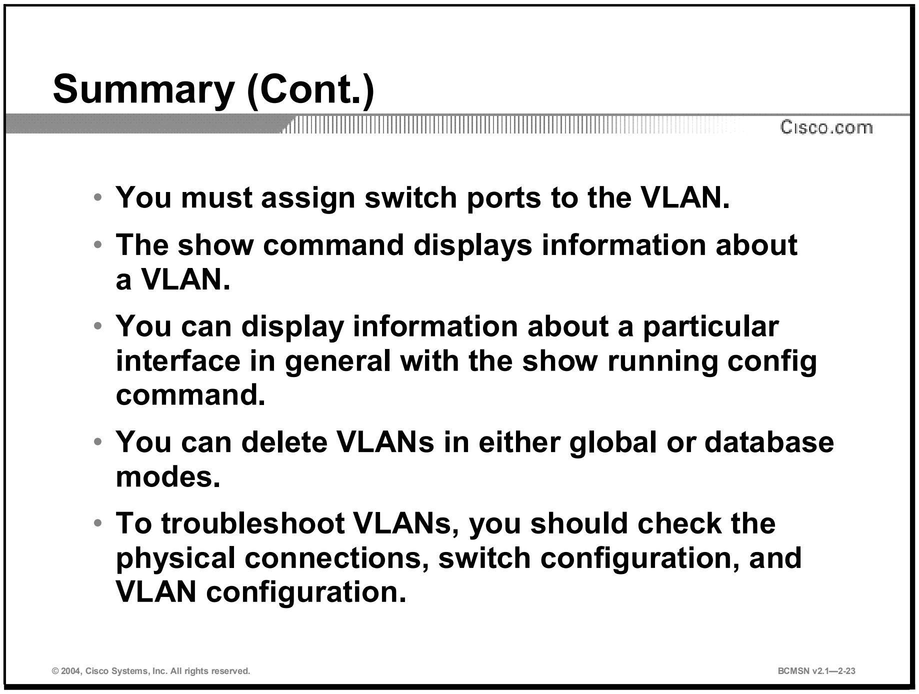

Summary (Cont.)

Cisco.com

- You must assign switch ports to the VLAN.
- The show command displays information about a VLAN.
- You can display information about a particular interface in general with the show running config command.
- You can delete VLANs in either global or database modes.
- To troubleshoot VLANs, you should check the physical connections, switch configuration, and VLAN configuration.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—2-23

References

For additional information, refer to these resources:

- Creating Ethernet VLANs on Catalyst Switches at http://www.cisco.com
- How To Configure InterVLAN Routing on Layer 3 Switches at http://www.cisco.com

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which two network characteristics apply to VLANs? (Choose two.)

A) All devices in a VLAN are members of the same broadcast domain.

B) The most common type of VLAN is a protocol VLAN.

C) VLAN membership is most commonly based on a switch port number.

D) A VLAN is the same as a physical subnet.

Q2) Which two statements identify the network benefits provided by VLANs? (Choose two.)

A) VLANs divide the network into larger broadcast domains or subnets.

B) VLANs reduce the impact of network problems.

C) VLANs help to isolate problem employees.

D) VLANs can transmit frames to all ports in all VLANs.

E) VLANs allow you to segregate frames that contain sensitive or critical information.

Q3) Match the type of VLAN with the correct VLAN characteristics.

_____ 1. end-to-end VLANs

_____ 2. local VLANs

A) It is grouped independent of physical location.

B) The network provides for a deterministic, consistent method of accessing resources.

C) Each VLAN has a common set of security and resource requirements for all members.

D) The user crosses a Layer 3 device to reach many of the resources.

E) Users can move from one VLAN to another without involving network administrators.

F) VLAN membership typically remains the same as a user moves around the campus.

Q4) Select those VLAN conditions that no longer exist as a result of applying the Enterprise Composite Network model. (Choose two.)

A) Failures at Layer 2 are isolated to the individual VLAN.

B) Redundancy is designed into the model.

C) There is a predictable traffic plan.

D) Capacity can be added with relatively small design effort.

Q5) Number the steps to create an Ethernet VLAN in the correct order.

_____ 1. Step 1

_____ 2. Step 2

_____ 3. Step 3

_____ 4. Step 4

A) Verify a VLAN configuration.

B) Verify a VLAN port configuration.

C) Assign ports to a VLAN.

D) Create a VLAN.

Q6) Match the steps with the appropriate VLAN mode.

_____ 1. global mode

_____ 2. database mode

_____ 3. both

A) Enter global configuration mode.

B) Create a new VLAN with a particular ID number.

C) Name a VLAN.

D) Enter VLAN configuration mode.

E) Exit VLAN data configuration mode.

F) Exit VLAN-configuration mode; Exit global configuration mode.

Q7) Number the steps to assign ports to a VLAN in the correct order.

_____ 1. Step 1

_____ 2. Step 2

_____ 3. Step 3

A) Specify the port as an access port.

B) Place the port in a particular VLAN.

C) From global config mode, enter the interface config mode for the particular port you want to add to the VLAN.

Q8) Select the command used to verify a VLAN configuration.

A) **show VLAN**

B) **VLAN database**

C) **config terminal**

D) **#VLAN**

Q9) Select the command that is used to verify a VLAN port configuration.

A) switch#**show running-config interface {fastethernet | gigabitethernet} slot/port**

B) show **VLAN id 3**

C) switchport **access VLAN 200**

D) switch#**show interfaces switchport**

Q10) Identify which sets of commands delete a VLAN in database mode and which sets delete a VLAN in global mode.

_____ 1. database mode

_____ 2. global mode

A) switch#**configure terminal**

B) switch#**vlan database**

C) switch**(VLAN)#no VLAN** ***VLAN_id***

D) switch**(config)no VLAN** ***VLAN_id***

E) switch**(config)#end**

F) switch**(VLAN)#exit**

Q11) What might cause two devices not to be able to communicate with each other?

A) The two switch ports are assigned to different subnets.

B) The two switch ports are assigned to different VLANs.

C) The two switch ports are assigned to different IP addresses.

D) The two switch ports are assigned the same subnet mask.

Quiz Answer Key

Q1) A, C

Relates to: Identifying the Features of a VLAN

Q2) B, E

Relates to: Identifying the Benefits of a VLAN

Q3) 1=A, C, E; 2=B, D, E

Relates to: Comparing VLAN Implementations

Q4) B, C

Relates to: Implementing VLANs in the Enterprise Composite Network Model

Q5) 1=D, 2=C, 3=A, 4=B

Relates to: Creating an Ethernet VLAN

Q6) 1=A, C, F; 2=D, E; 3=B

Relates to: Creating a VLAN in Global and Database Mode

Q7) 1=C, 2=A, 3=B

Relates to: Assigning Ports to a VLAN

Q8) A

Relates to: Verifying a VLAN Configuration

Q9) A

Relates to: Verifying a VLAN Port Configuration

Q10) 1=B, C, F; 2=A, D, E

Relates to: Deleting a VLAN

Q11) B

Relates to: Troubleshooting VLAN Operations

Supporting Multiple VLANs Between Two Switches

Overview

Multiple VLANs are supported between switches through the use of VLAN trunks. A trunk is a Layer 2 link between switches running a specialized trunking protocol.

Relevance

In a switched network, users can be assigned to VLANs that span multiple connected switches. It is therefore necessary to understand how switches identify frames that belong to respective VLANs.

Objectives

Upon completing this lesson, you will be able to:

- Select the problems that occur when supporting multiple VLANs
- Select the correct features of trunking protocol
- Identify the benefits trunks provide in the Enterprise Composite Network model
- Correctly identify which functions belong to an access port, a trunk port, and a dynamic port
- Correctly match the features and functions with their appropriate trunking protocol
- Match the operation mode with the correct switch port DTP mode

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)
- Basic knowledge of the components of the Enterprise Composite Network model
- Basic knowledge of the characteristics of VLANs
- Basic knowledge of VLAN configuration

Outline

This lesson includes these topics:

- Overview
- Maintaining Specific VLAN Identification
- Identifying the Features of Trunking Protocols
- Identifying the Operational Mode of Switch Ports
- Identifying the Modes for Dynamic Trunking Protocol
- Summary
- Quiz

Maintaining Specific VLAN Identification

This topic identifies the problems that can occur when supporting multiple VLANs.

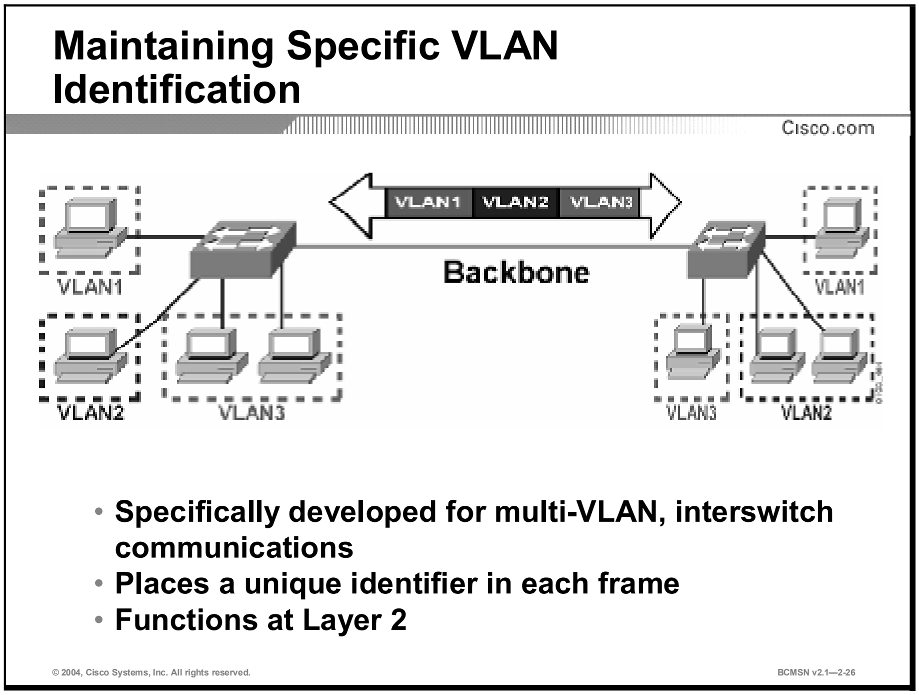

Maintaining Specific VLAN Identification

- Specifically developed for multi-VLAN, interswitch communications
- Places a unique identifier in each frame
- Functions at Layer 2

BCMSN v2.1—2-26

Multiple VLANs are supported between switches through the use of VLAN trunks. A trunk is a Layer 2 link between switches running a specialized trunking protocol. Trunks carry the traffic of multiple VLANs over physical links (multiplexing) and enable the extension of Layer 2 operations between switches.

In a switched network, users can be assigned to VLANs that span multiple connected switches. In this environment, an identification method by which switches identify frames belonging to their respective VLANs is necessary.

Frame identification uniquely assigns an ID, referred to as a VLAN ID (VID) to each frame. The VID is placed in the header of each frame as it is forwarded on a trunk link. Each switch examines this VID to determine the destination VLAN of the frame. Then the switch can make the appropriate decision to flood or forward the frame to ports in the VLAN.

When the frame exits an access link that is *not* by definition running a trunking protocol, the switch removes the identifier before transmitting the frame.

Identifying the Features of Trunking Protocols

This topic identifies the features of VLAN trunking protocols.

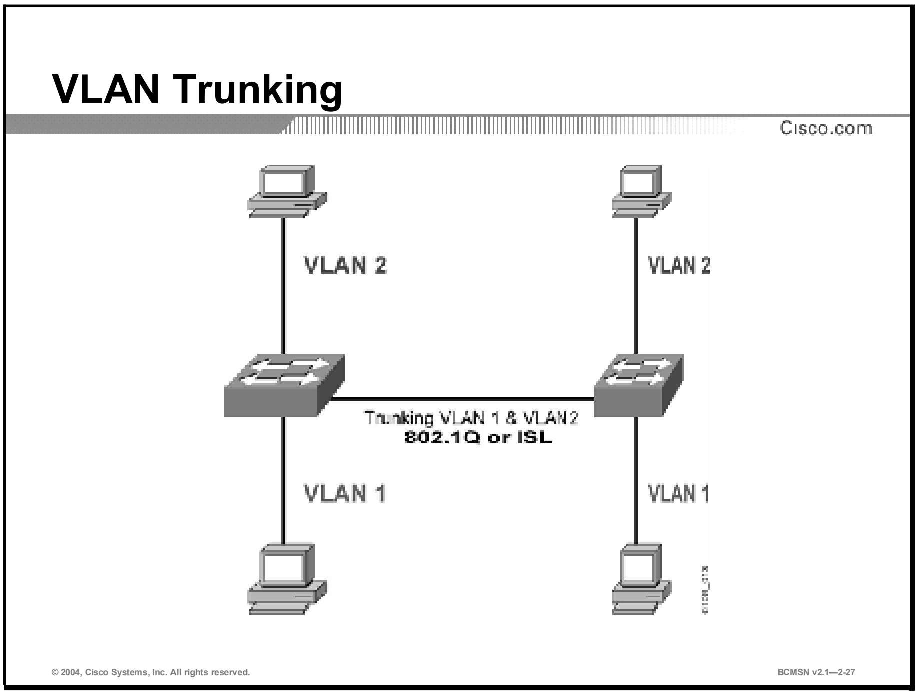

VLAN Trunking

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-27

Trunks transport VLAN information between connected Layer 2 devices. Each frame passed on a trunk is either encapsulated or tagged, depending on the protocol. The trunking protocol, which defines the type and method of trunking, can be Inter-Switch Link (ISL), which does encapsulation, or IEEE 802.1Q, which tags.

Cisco supports the following multiple trunking methodologies:

- **ISL:** A Cisco proprietary encapsulation protocol for interconnecting multiple switches. This protocol is supported in Catalyst switches as well as routers.
- **IEEE 802.1Q:** An IEEE standard method for identifying VLANs by inserting a VLAN identifier into the frame header. This process is called frame tagging.
- **LAN Emulation (LANE):** An IEEE standard method for transporting VLANs over Asynchronous Transfer Mode (ATM) networks.
- **IEEE 802.10:** A Cisco proprietary method of transporting VLAN information inside the standard 802.10 frame (FDDI). The VLAN information is written to the Security Association Identifier (SAID) portion of the 802.10 frame. This method is typically used to transport VLANs across FDDI backbones.

ISL and 802.1Q trunking protocols are discussed in more detail later in this module. The following table summarizes frame identification methods and encapsulation with their associated media types.

Frame Tagging and Encapsulation Methods

Identification Method	Encapsulation	Tagging (Insertion into Frame)	Media
ISL	Yes	No	Ethernet
802.1Q	No	Yes	Ethernet
LANE	No	No	ATM
802.10	No	No	FDDI

Identifying the Operational Mode of Switch Ports

This topic specifies the functions that belong to access, trunk, and dynamic ports.

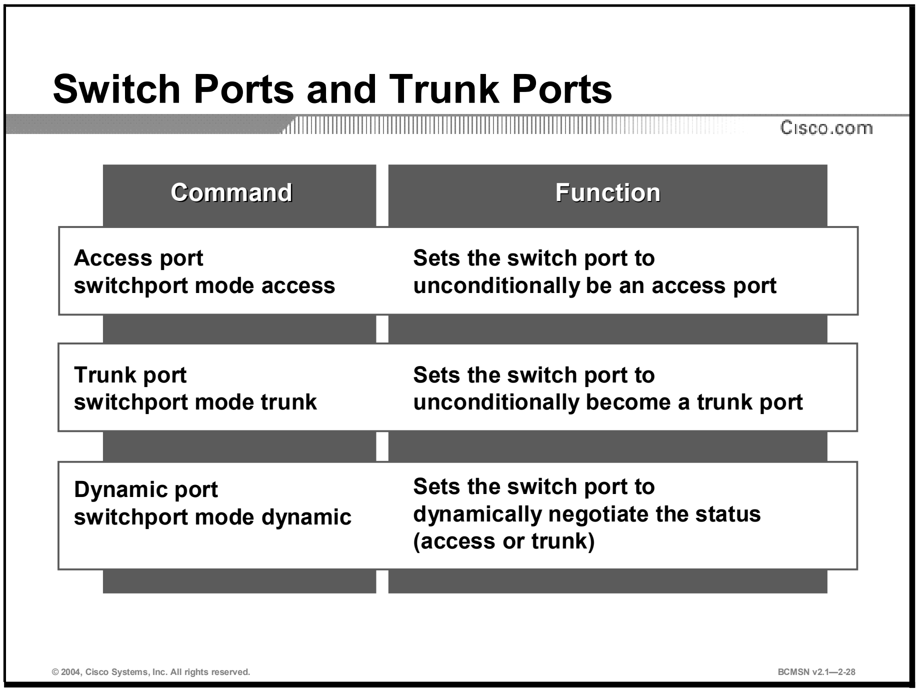

Switch Ports and Trunk Ports

Command	Function
Access port **switchport mode access**	**Sets the switch port to unconditionally be an access port**
Trunk port **switchport mode trunk**	**Sets the switch port to unconditionally become a trunk port**
Dynamic port **switchport mode dynamic**	**Sets the switch port to dynamically negotiate the status (access or trunk)**

Switch ports are Layer 2 interfaces on Catalyst switches. Routed ports are Layer 3 interfaces also on Catalyst switches. To establish a trunk link, you must first ensure that the interface has been enabled as a Layer 2 interface.

The two most commonly used operational modes for switch ports are as follows:

- **Access (for access ports):** Provides access to a VLAN
- **Trunk (for trunk ports):** Represents the endpoint mode required to support a trunk link

You can configure the switch port mode statically via explicit configuration or dynamically by specifying negotiation. Negotiation uses Dynamic Trunking Protocol (DTP) to negotiate whether a switch port will become a trunk port through a protocol exchange sequence. DTP can also be configured to negotiate the type of trunking protocol to be used; that is, ISL or 802.1Q.

Ports have the following two other modes:

- **Administrative mode:** This is the configuration parameter applied to the switch port.
- **Operational mode:** This is the current operation state of the switch port.

These modes can be different. For example, if DTP is unable to establish a trunk link through negotiation, the switch port remains as an access port.

The three basic switch port administrative modes are as follows:

- **Access port (switchport mode access):** Sets the switch port operational mode to unconditionally be an access port. Regardless of any other DTP configurations or switch port status, this switch port will never transition to a trunk port to become one endpoint of the trunk link; hence its operational mode is static. Access ports carry Layer 2 frames for a single VLAN or provide access for a network node into a VLAN. The Layer 2 frames are not tagged or encapsulated in any way. Use the switchport access command to configure a VLAN membership, if necessary. The default VLAN membership will be VLAN1.
- **Trunk port (switchport mode trunk):** Sets the switch port to unconditionally become a trunk port. Regardless of any other DTP configurations or switch port status, this switch port will always be one endpoint of the trunk link. Using the switchport trunk encapsulation command, the trunking encapsulation protocol must be manually configured before statically defining the switch port as a trunk port. If this is not done, the switchport mode trunk command will fail, and an error will be generated.
- **Dynamic port (switchport mode dynamic):** Sets the switch port to become a dynamic port that will participate in DTP negotiations. These negotiations determine if the switch port will be able to transition to a trunk port and, if configured to do so, what trunking protocol is used.

  Before a switch port can be configured as a dynamic port, DTP must be enabled with the **no switchport nonegotiate** command. Additional configuration is generally required to establish trunk links with DTP.

Identifying the Modes for Dynamic Trunking Protocol

This topic identifies the functions of the various modes for DTP.

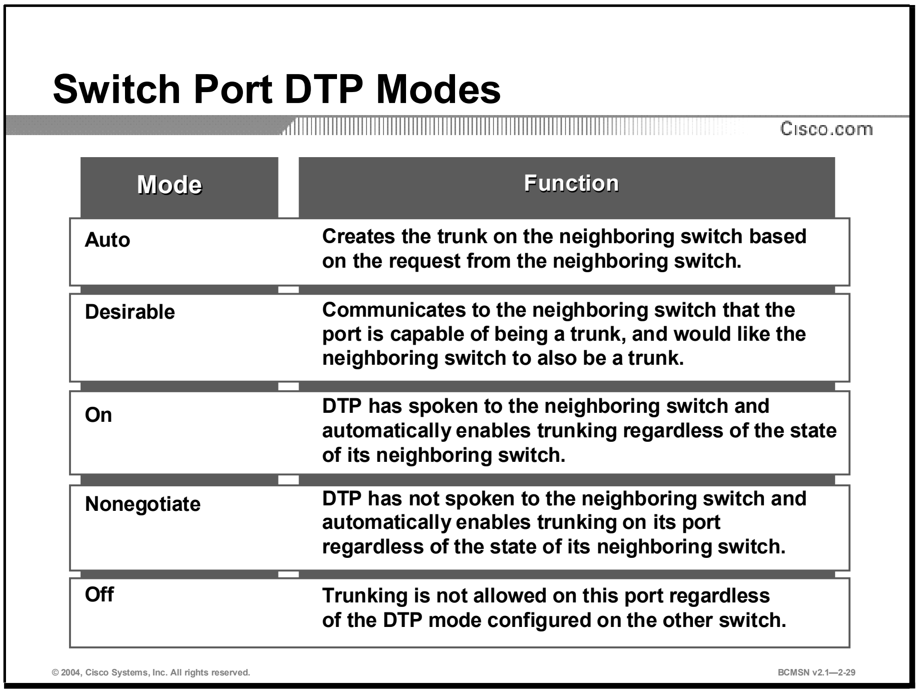

Switch Port DTP Modes

Cisco.com

Mode	Function
Auto	Creates the trunk on the neighboring switch based on the request from the neighboring switch.
Desirable	Communicates to the neighboring switch that the port is capable of being a trunk, and would like the neighboring switch to also be a trunk.
On	DTP has spoken to the neighboring switch and automatically enables trunking regardless of the state of its neighboring switch.
Nonegotiate	DTP has not spoken to the neighboring switch and automatically enables trunking on its port regardless of the state of its neighboring switch.
Off	Trunking is not allowed on this port regardless of the DTP mode configured on the other switch.

BCMSN v2.1—2-29

Catalyst switches run DTP, a Cisco proprietary protocol. DTP is a trunk-related protocol that negotiates the operational mode of directly connected switch ports to a trunk port, if possible. DTP also selects the trunking protocol if configured to do so.

Summary

This topic summarizes the key points discussed in this lesson.

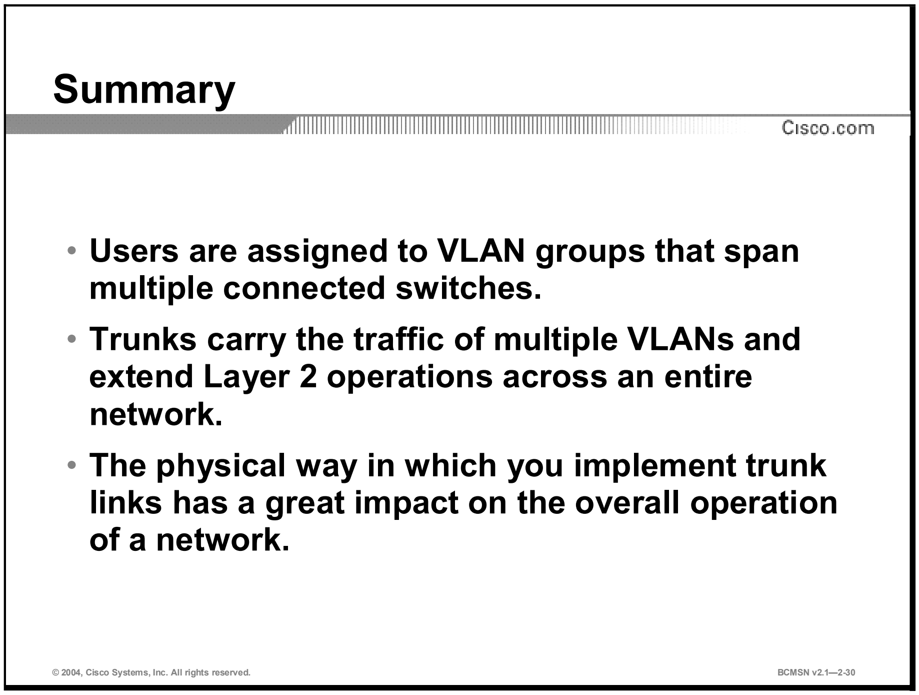

Summary

- Users are assigned to VLAN groups that span multiple connected switches.
- Trunks carry the traffic of multiple VLANs and extend Layer 2 operations across an entire network.
- The physical way in which you implement trunk links has a great impact on the overall operation of a network.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—2-30

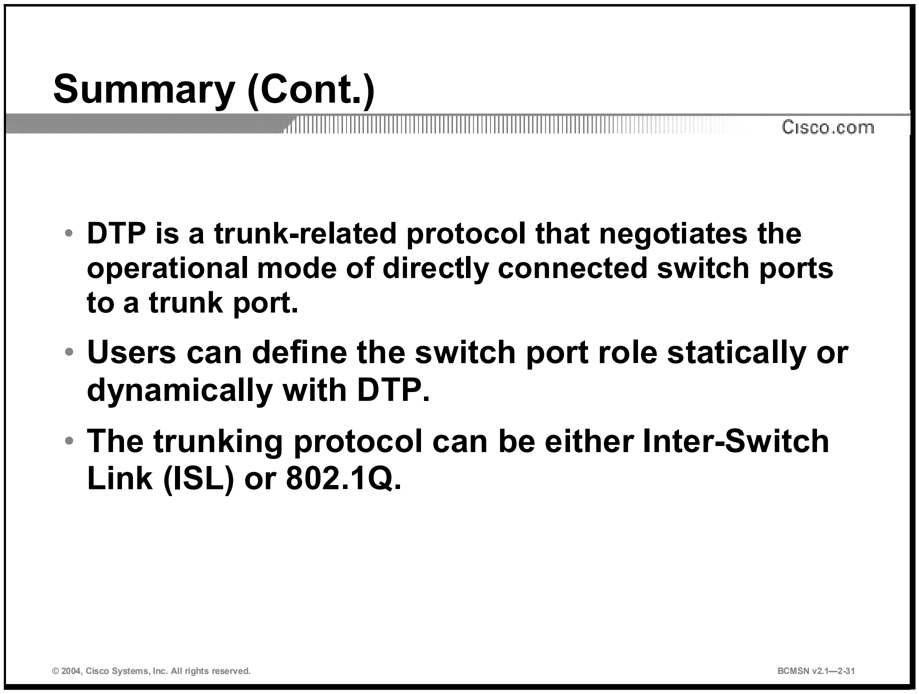

Summary (Cont.)

- DTP is a trunk-related protocol that negotiates the operational mode of directly connected switch ports to a trunk port.
- Users can define the switch port role statically or dynamically with DTP.
- The trunking protocol can be either Inter-Switch Link (ISL) or 802.1Q.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—2-31

References

For additional information, refer to these resources:

- The documentation that accompanied your Cisco Catalyst switch
- Standards Requirements documentation on VLAN tagging (IEEE 802.1Q and IEEE802.1ac)

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Select the problem that can occur when supporting multiple VLANs.

A) An incorrectly assigned VID can corrupt the frame.

B) If a VID is not assigned to a frame, the switch cannot determine where to forward the frame.

C) The link can become oversubscribed.

D) If a frame exits an access link that is not running a trunking protocol, the switch removes the VID before transmitting the frame.

E) Users can only be assigned to VLANs that are relatively small.

Q2) Select the features that best describe trunking protocols. (Choose two.)

A) Trunks carry the traffic of a single VLAN over a single physical link.

B) Trunks transport VLAN information between connected Layer 2 devices.

C) Frames passed on trunks are either static or dynamic.

D) Trunks can extend Layer 2 operations across an entire network.

Q3) Match the operation description with the appropriate switch port.

_____ 1. access port

_____ 2. trunk port

_____ 3. dynamic port

A) carries Layer 2 frames for a single VLAN

B) always the endpoint of the trunk link

C) participates in DTP negotiations

D) default operational mode of a switch port

E) default for the **switchport nonegotiate** command is **no switchport nonegotiate**

Q4) Match the appropriate command of the Dynamic Trunking Protocol to the operational mode of a switch port.

_____ 1. access

_____ 2. trunk

_____ 3. no nonegotiate

_____ 4. dynamic desirable

A) puts interface into permanent nontrunking mode and negotiates to convert the link into a trunk link

B) puts the interface into permanent trunking mode but prevents the interface from generating DTP frames

C) puts the interface into permanent trunking mode and negotiates to convert the link into a trunk link

D) makes the interface actively attempt to convert the link to a trunk link

Quiz Answer Key

Q1) C

Relates to: Maintaining Specific VLAN Identification

Q2) B, D

Relates to: Identifying the Features of Trunking Protocols

Q3) 1=A, D; 2=B, E; 3=C

Relates to: Identifying the Operational Mode of Switch Ports

Q4) 1=A, 2=C, 3=B, 4=D

Relates to: Identifying the Modes for Dynamic Trunking Protocol

Defining Trunking Protocols

Overview

A VLAN trunk is either a physical or a virtual point-to-point link primarily used to carry multiple VLANs and their traffic. A trunk is used to carry Layer 2 traffic and other information needed to maintain networks as a whole.

Relevance

VLAN trunks are necessary in any multilayer switched network and, therefore, require in-depth understanding to implement and troubleshoot appropriately.

Objectives

Upon completing this lesson, you will be able to:

- Match the features with the appropriate ISL or 802.1Q trunking protocols
- Identify the features and benefits that belong to the ISL protocol
- Identify the steps that belong to the ISL and Layer 2 encapsulation processes
- Select the features and benefits that belong to the 802.1Q trunking protocol
- Match the correct features with the 802.1Q tagging process and Layer 2 encapsulation protocol
- Select the issues that can occur if you misconfigure a native VLAN
- Match the correct VLAN range to its appropriate usage
- Identify how trunking protocols are used in the Enterprise Composite Network model

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)
- Basic knowledge of the components of the Enterprise Composite Network model
- Basic knowledge of VLANs
- Basic knowledge of the features and functions of VTP
- Basic knowledge of switch operation
- Basic knowledge of frame encapsulation

Outline

This lesson includes these topics:

- Overview
- Comparing ISL and 802.1Q Trunking Protocols
- Identifying the Features and Benefits of the ISL Protocol
- Comparing the ISL and Layer 2 Encapsulation Processes
- Identifying the Features and Benefits of 802.1Q Trunking
- Comparing 802.1Q Tagging Process and Layer 2 Encapsulation Protocol
- Supporting Native VLANs
- Mapping VLAN Ranges
- Using Trunking Protocols in the Enterprise Composite Network Model
- Summary
- Quiz

Comparing ISL and 802.1Q Trunking Protocols

This topic compares the features of ISL and 802.1Q trunking protocols.

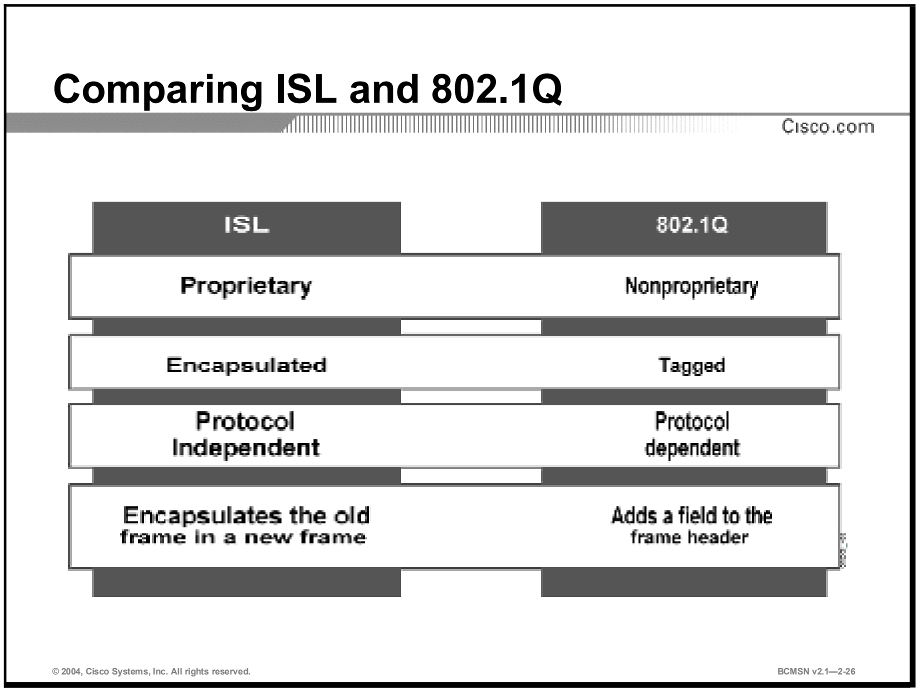

Comparing ISL and 802.1Q

Cisco.com

ISL	802.1Q
Proprietary	Nonproprietary
Encapsulated	Tagged
Protocol Independent	Protocol dependent
Encapsulates the old frame in a new frame	Adds a field to the frame header

BCMSN v2.1—2-26

Each data frame sent across a trunk is either encapsulated or tagged, depending on the trunking protocol configured. The trunking protocol defines the type and method of trunking. The protocol can be either Inter-Switch Link (ISL), which does encapsulation and is a Cisco proprietary protocol, or IEEE 802.1Q, which inserts a tag into the data frame and is a defined open standard. The purpose of encapsulating or tagging frames being sent on a trunk is to provide the switch that is receiving the frame specific information about the frame. The information that is provided can be the VID, which identifies from which VLAN the frame originated.

The 802.1Q trunking protocol provides the same functionality as ISL. 802.1Q is not proprietary and can be deployed in any standards-based network and on any standards-based Layer 2-capable device. 802.1Q supports VLANs on Ethernet technology, such as 802.3 Ethernet. 802.1Q is protocol-dependent. However, because it modifies the Layer 2 frame by inserting a tag between two specific fields of the frame, the 802.1Q protocol must be aware of the frame type to modify it.

ISL is protocol-independent. The encapsulated frame is capable of transporting data frames from other various media types, and it is a direct result of the true encapsulation method by which ISL operates. ISL does not modify the contents of the frame, it only appends to the beginning and end of the frame.

Identifying the Features and Benefits of the ISL Protocol

This topic identifies the features and benefits of the ISL protocol.

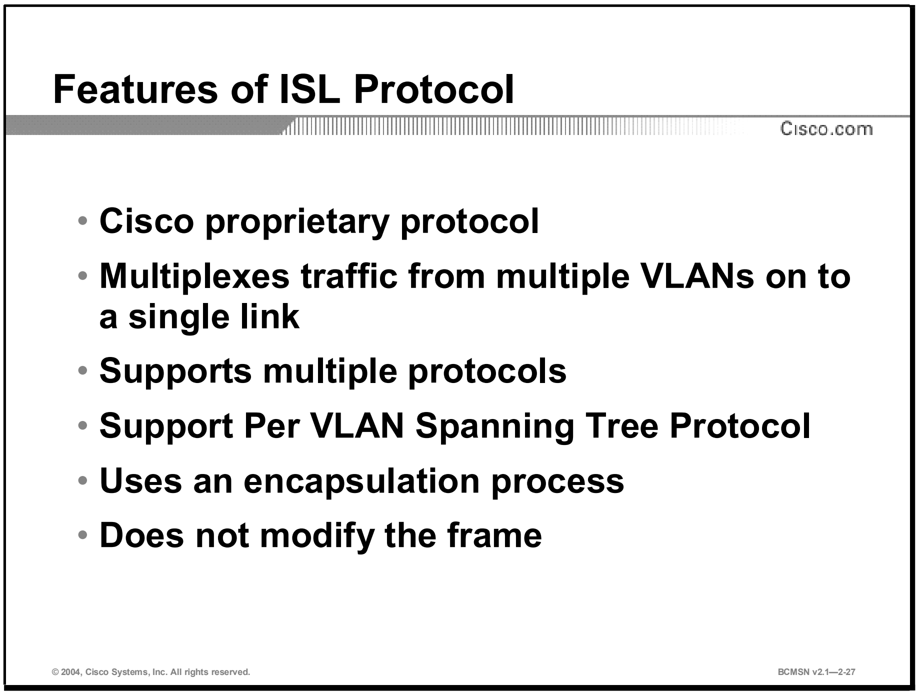

Features of ISL Protocol

- Cisco proprietary protocol
- Multiplexes traffic from multiple VLANs on to a single link
- Supports multiple protocols
- Support Per VLAN Spanning Tree Protocol
- Uses an encapsulation process
- Does not modify the frame

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—2-27

ISL is a Cisco proprietary protocol that maintains VLAN information used in Layer 2 operations. ISL is used for multiplexing traffic from multiple VLANs on a single link.

ISL has been implemented for connections between switches, routers, and network interface cards (NICs) used on nodes such as servers. To support the ISL trunking, each trunk link endpoint must be configured; otherwise, it must have negotiated for ISL and must be a physical or virtual point-to-point link.

Only ISL trunk ports can properly handle ISL-encapsulated frames. A non-ISL device that receives ISL-encapsulated Ethernet frames may consider the frame to be transmission errors, if the size of the header plus the data frame exceeds the maximum transmission unit (MTU) size. An ISL-encapsulated frame will be dropped if the receiving switch port is not an ISL trunk port. This is because the receiving device will not know how to process the encapsulated frame.

The following are some benefits of the ISL protocol:

- It is a Cisco proprietary protocol (predates standard by several years).
- It supports multiple protocols (Ethernet, Token Ring, FDDI).
- It supports 1000-user configurable VLANs.
- It supports Per VLAN Spanning Tree Protocol (STP).
- It has a large installed base.
- It uses an encapsulation process that leaves frames unmodified (more secure and transparent).
- It supports a point-to-point topology.

Comparing the ISL and Layer 2 Encapsulation Processes

This topic identifies the steps that belong to the ISL and Layer 2 encapsulation processes.

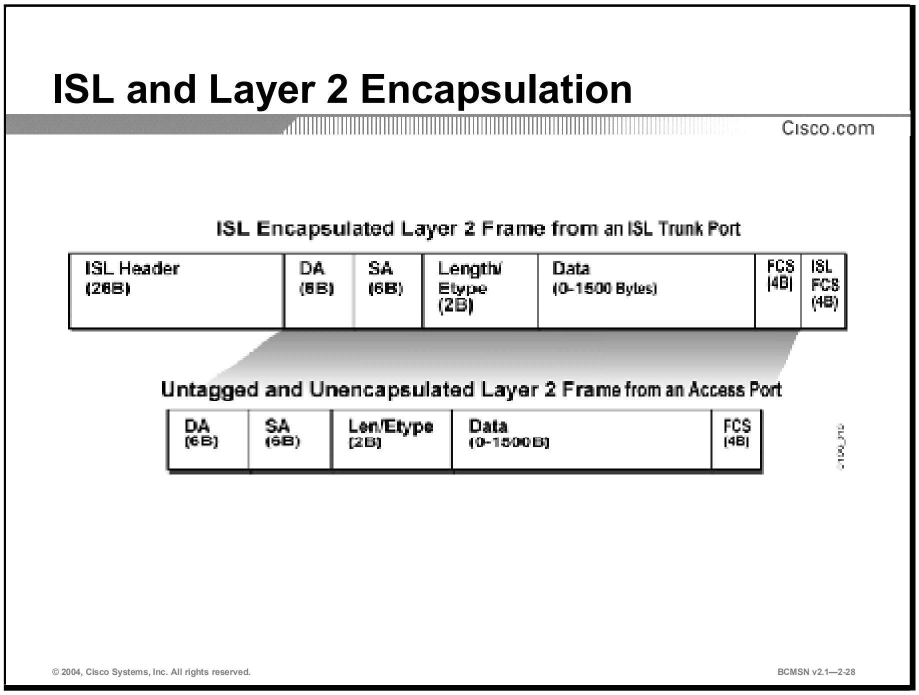

ISL and Layer 2 Encapsulation

Cisco.com

ISL Encapsulated Layer 2 Frame from an ISL Trunk Port

ISL Header (26B)	DA (6B)	SA (6B)	Length/ Etype (2B)	Data (0-1500 Bytes)	FCS (4B)	ISL FCS (4B)

Untagged and Unencapsulated Layer 2 Frame from an Access Port

DA (6B)	SA (6B)	Len/Etype (2B)	Data (0-1500B)	FCS (4B)

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-28

Switch ports configured as ISL trunk ports encapsulate each Layer 2 frame before sending it out the trunk port. With ISL, the original frame is encapsulated and an additional header is added before the frame is carried over a trunk link. At the receiving end, the header is removed and the packet is forwarded to the assigned VLAN. For each frame that is received on a trunk port, the switch recalculates the total frame cyclic redundancy check (CRC) and checks it against the CRC value in the FCS. If the CRC values do not match, the frame is discarded. If the values match, the switch discards the FCS, then de-encapsulates or strips away the header to read in the field information for processing.

The field information contains values that define various attributes of the original Layer 2 data within the encapsulated frame. This information is used for forwarding, media identification, and VLAN identification. The actual frame header and the fields within the frame depend upon the type of VLAN and media that connects to that VLAN. Ethernet is encapsulated with a 26-byte ISL header and a four-byte FCS. The 30-byte ISL encapsulation overhead is consistent among the different media types supported on Catalyst switches, but the overall size of the frame may vary depending on the media type and the supported MTUs of those media types.

Note This course only covers the Ethernet technology.

The ISL Ethernet frame header contains these information fields:

- **DA (destination address):** 40-bit destination address. This is a multicast address and is set at 0x01-00-0C-00-00 or 0x03-00-0c-00-00. The first 40 bits of the DA field signal the receiver that the packet is in ISL format.
- **TYPE:** Four-bit descriptor of the encapsulated frame types: Ethernet (0000), Token Ring (0001), FDDI (0010), and ATM (0011).
- **USER:** Four-bit descriptor used as the TYPE field extension or to define Ethernet priorities; it is a binary value from 0, the lowest priority, to 3, the highest priority. The default USER field value is "0000." For Ethernet frames, the USER field bits "0" and "1" indicate the priority of the packet as it passes through the switch.
- **SA (source address):** 48-bit source MAC address of the transmitting Catalyst switch port.
- **LEN (length):** 16-bit frame-length descriptor minus DA, TYPE, USER, SA, LEN, and CRC.
- **AAAA03:** Standard Subnetwork Access Protocol (SNAP) 802.2 logical link control (LLC) header.
- **HSA (high system availability):** First three bytes of the SA (manufacturer or unique organizational ID).
- **VLAN ID:** 15-bit VID. Only the lower 10 bits are used for 1024 VLANs.
- **BPDU (bridge protocol data unit):** One-bit descriptor identifying whether the frame is a spanning tree BPDU; also identifies if the encapsulated frame is a CDP frame.
- **INDX (index):** Indicates the port index of the source of the packet as it exits the switch. It is used for diagnostic purposes only and may be set to any value by other devices. It is a 16-bit value and is ignored in received packets.
- **ENCAP FRAME:** Encapsulated data packet, including its own CRC value, completely unmodified. The internal frame must have a CRC value that is valid when the ISL encapsulation fields are removed. A receiving switch may strip off the ISL encapsulation fields and use this ENCAP FRAME field as the frame is received (associating the appropriate VLAN and other values with the received frame as indicated for switching purposes).
- **FCS (frame check sequence):** Consists of four bytes. This sequence contains a 32-bit CRC value, which is created by the sending MAC and is recalculated by the receiving MAC to check for damaged frames. The FCS is generated over the DA, SA, LEN, TYPE, and Data fields. When an ISL header is attached, a new FCS is calculated over the entire ISL packet and added to the end of the frame.
- **RES (reserved):** 16-bit reserved field used for additional information, such as the FDDI frame control field.

When examining the frame formats for Layer 2 frames, notice that the ISL Layer 2 header comes before the other Layer 2 field information in the frame.

The ISL frame contains two FCS fields. The first FCS field is generated from the original transmitting device of the frame. The second FCS field is generated from the ISL trunk port. ISL encapsulates the frame without modifying its contents.

Identifying the Features and Benefits of 802.1Q Trunking

This topic identifies the features and benefits of the 802.1Q trunking protocol.

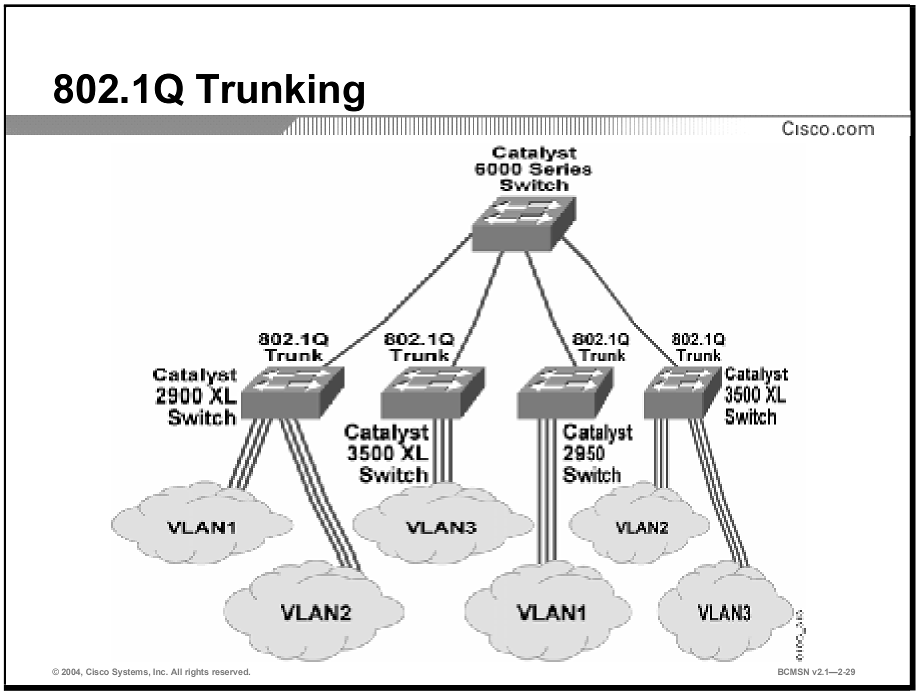

802.1Q Trunking

Cisco.com

BCMSN v2.1—2-29

The 802.1Q frame tagging that the Catalyst series of switches uses multiplexes traffic from multiple VLANs on a single link. To support the 802.1Q trunking protocol on a trunk link where multiple VLANs are multiplexed on that single link, you must configure each trunk link endpoint to support 802.1Q trunking. 802.1Q trunk links employ a tagging mechanism to carry frames for multiple VLANs, in which each frame is tagged to identify the VLAN to which the frame belongs.

The benefits of applying the 802.1Q protocol in a network are as follows:

- IEEE standard
- Tagging process
- Supports Ethernet and Token Ring
- Supports 4096 VLANs
- Supports Per VLAN Spanning Tree (PVST) and supports Multiple Spanning Tree (MST) and Rapid PVST, which are enhanced spanning tree protocols)
- Growing installed base
- Native VLAN untagged
- Requires NIC card to support 1522 MTU
- Point-to-multipoint topologies
- Extended quality of service (QoS) support

Comparing the 802.1Q Tagging Process and Layer 2 Encapsulation Protocol

This topic compares the features of the 802.1Q tagging process and the Layer 2 encapsulation protocol.

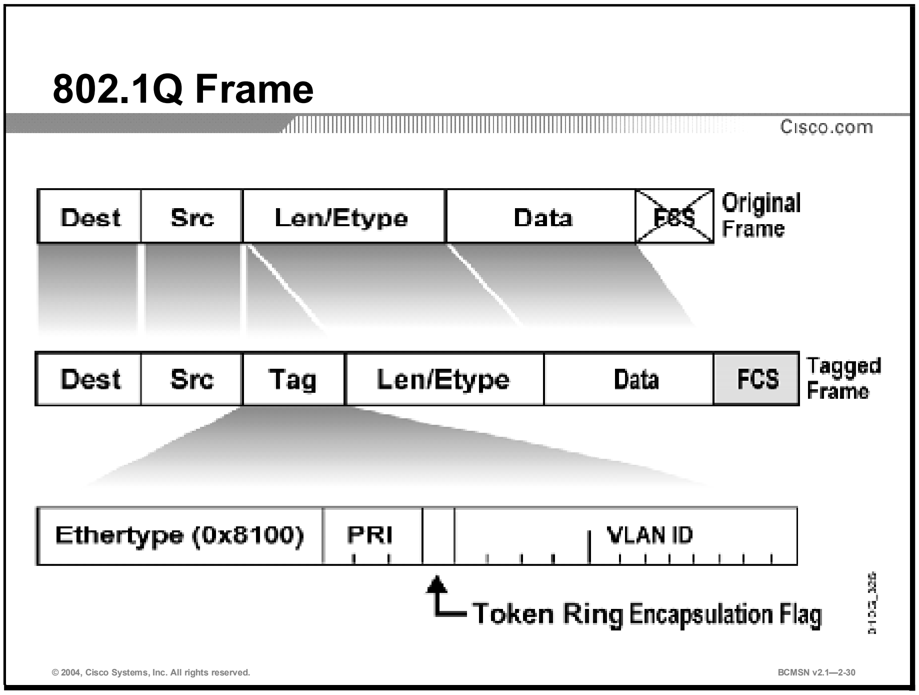

802.1Q Frame

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-30

The 802.1Q encapsulation adds a tag in a standard Layer 2 Ethernet data frame. 802.1Q uses an internal tagging mechanism, in which a tag is inserted within the data frame itself. 802.1Q predefined the three most significant bits in the 802.1Q tag to allow for prioritization of the Layer 2 frame. The switch then recalculates the checksum value for the entire tagged frame and inserts the new value in a new header. ISL, in comparison, does not modify the frame at all.

Compare the 802.1Q tagged frame to a standard Layer 2 Ethernet untagged frame. The first information in the header is the destination and source addresses. The header contains only a single frame checksum sequence field at the end. Layer 2 devices, except those with 802.1Q trunk ports, do not have the capability of deciphering the 802.1Q tag.If a non-802.1Q-enabled device or an access port receives a frame, the tagged portions of the frame are ignored, and the packet is switched at Layer 2 as if it was a standard Ethernet frame. This allows for the placement of Layer 2 intermediate devices, such as other switches or bridges, on the 802.1Q trunk link.

The Layer 2 frame header for 802.1Q and ISL includes the VID, which indicates the VLAN source and destination encapsulation or tagging. The default behavior of VLAN trunks is to permit all normal and extended range VLANs across the link. A best practice is to limit the trunk to only the intended VLANs. This practice reduces the possibility of loops and improves bandwidth utilization by restricting unwanted VLAN data traffic from the link.

Supporting Native VLANs

This topic describes the issues that can occur if you misconfigure a native VLAN.

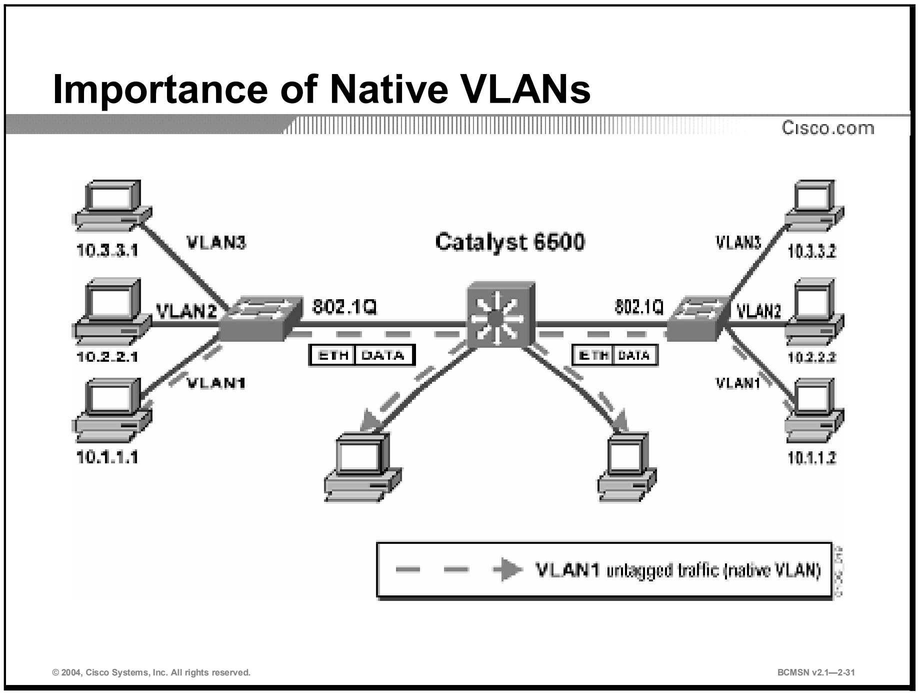

Importance of Native VLANs

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-31

802.1Q trunks define a native VLAN for frames that are not tagged by default. An 802.1Q native VLAN is defined as one of the following:

- The VLAN that a port is in when not in trunking operational mode
- The VLAN from which Layer 2 frames will be transmitted untagged on an 802.1Q trunk port
- The VLAN to which Layer 2 frames will be forwarded if received untagged on an 802.1Q trunk port

With an 802.1Q native VLAN, a switch can forward any Layer 2 frame received on a trunk port, whether tagged or not, to an intended VLAN. Any Layer 2 frames from a native VLAN are transmitted from the trunk port untagged. Compare 802.1Q to ISL, where any unencapsulated frames received on a trunk port are immediately dropped and all frames transmitted from a trunk port are encapsulated.

Note A frame with the frame tag is called a “baby giant.” For ISL, a baby giant is 1548 bytes; while for 802.1Q, a baby giant is 1522 bytes.

By default on Catalyst switches, all switch ports and native VLANs are initially assigned to VLAN1. The 802.1Q trunk ports connected to each other via physical or logical segments must all have the same native VLAN configured. A trunk port will support only one native VLAN. If you do not configure the same native VLAN on all switches and you use CDP, CDP will issue a “VLAN mismatch” error message to any active consoles. There are some specific cases where either CDP is turned off or cannot be transmitted through an intermediate Layer 2 device in the same manner that 802.1Q frames can. If the native VLAN is misconfigured for trunk ports on the same trunk link, Layer 2 loops can occur.

Each physical port has a parameter called a Port VLAN identifier (PVID). Every 802.1Q port is assigned a PVID value based on its native VID (default is VLAN1). All untagged frames are assigned to the PVID parameter. When a port receives a tagged frame, the tag is respected. If the frame is untagged, the value contained in the PVID is considered a tag. This allows the coexistence on the same Ethernet segment of VLAN-aware bridges or stations and VLAN-unaware bridges or stations.

When configuring 802.1Q on a Cisco switch, it is possible to define a different native VLAN or access VLAN to an access port from that defined for an operational trunk port.

Mapping VLAN Ranges

This topic maps each VLAN range to its correct usage.

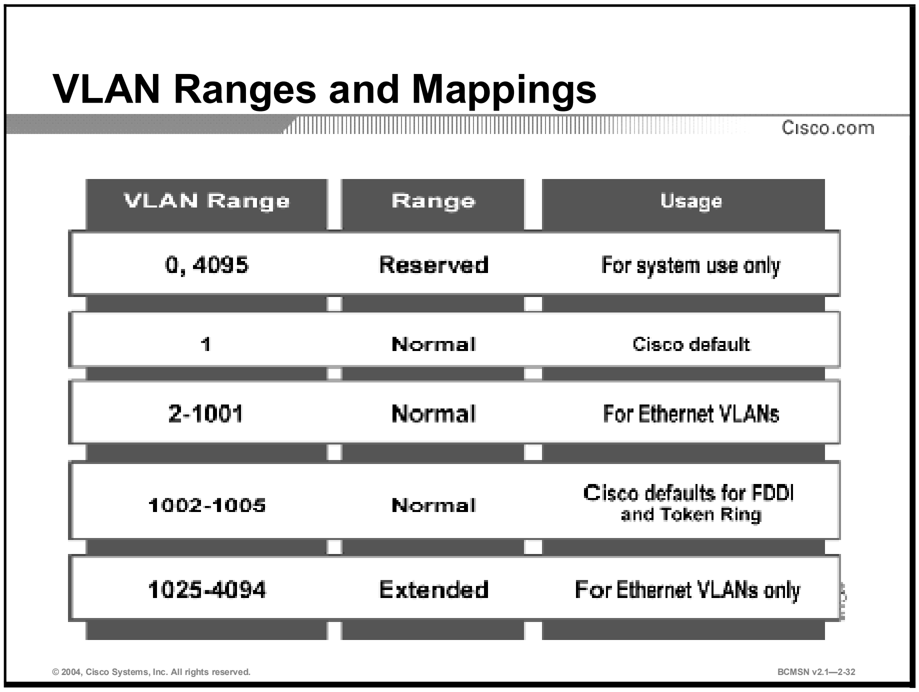

VLAN Ranges and Mappings

Cisco.com

VLAN Range	Range	Usage
0, 4095	Reserved	For system use only
1	Normal	Cisco default
2-1001	Normal	For Ethernet VLANs
1002-1005	Normal	Cisco defaults for FDDI and Token Ring
1025-4094	Extended	For Ethernet VLANs only

BCMSN v2.1—2-32

Each VLAN on the network must have a unique VID. The valid range of user-configurable ISL VLANs is 1 to 1001. The valid range of VLANs specified in the IEEE 802.1Q standard is 0 to 4094. The table describes the VLAN ranges and usage.

VLAN Ranges	Range	Usage	Propagated
0, 4095	Reserved	For system use only. You cannot see or use these VLANs.	—
1	Normal	Cisco default. You can use this VLAN, but you cannot delete it.	Yes
2-1001	Normal	For Ethernet VLANs. You can create, use, and delete these VLANs.	Yes
1002-1005	Normal	Cisco defaults for FDDI and Token Ring. You cannot delete VLANs 1002–1005.	Yes
1025-4094	Extended	For Ethernet VLANs only.	No

In a network environment with other vendor devices connected to Cisco switches through 802.1Q trunks, you must map 802.1Q VLAN numbers greater than 1000 to ISL VLAN numbers on the Cisco switches.

802.1Q VLANs in the range 1 to 1000 are automatically mapped to the corresponding ISL VLAN. 802.1Q VLAN numbers greater than 1000 must be mapped to an ISL VLAN to be recognized and forwarded by Cisco switches. As a best practice, assign extended VLANs beginning with 4094 and work downward.

These restrictions apply when mapping 802.1Q VLANs to ISL VLANs:

- You can configure up to eight 802.1Q-to-ISL VLAN mappings on the switch.
- You can only map 802.1Q VLANs to Ethertype ISL VLANs.
- Do not enter the native VLAN of any 802.1Q trunk in the mapping table to avoid overlapping numbers.
- When you map an 802.1Q VLAN to an ISL VLAN, traffic on the 802.1Q VLAN corresponding to the mapped ISL VLAN is blocked. For example, if you map 802.1Q VLAN 2000 to ISL VLAN 200, traffic on 802.1Q VLAN 200 is blocked.
- VLAN mappings are local to each switch. Make sure you configure the same VLAN mappings on all appropriate switches in the network.

Using Trunking Protocols in the Enterprise Composite Network Model

This topic identifies how trunks are used in the Enterprise Composite Network model.

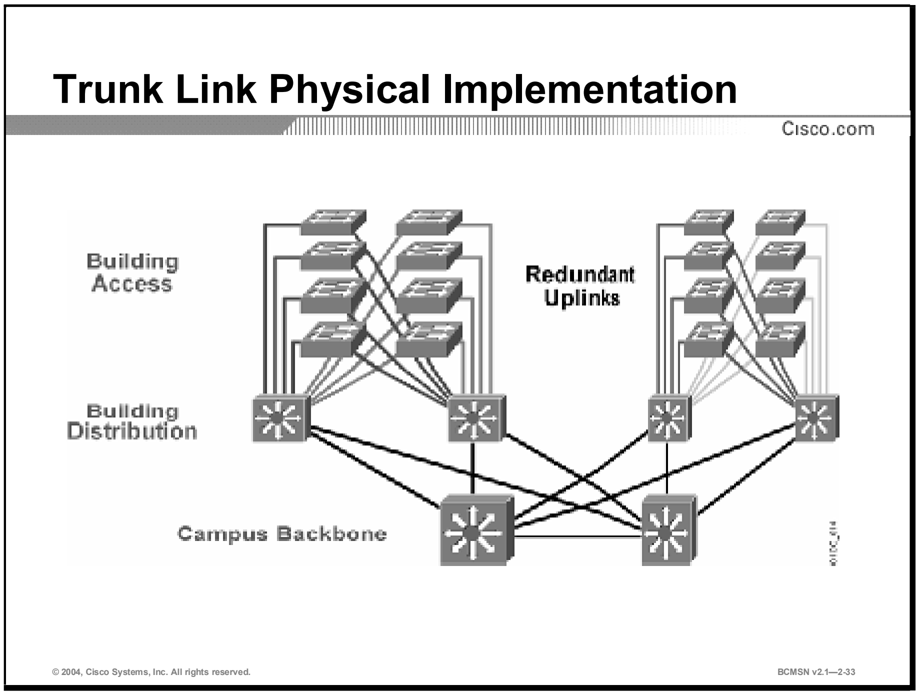

Trunk Link Physical Implementation

Cisco.com

BCMSN v2.1—2-33

VLAN trunks make it possible for the same VLAN to exist on multiple switches. Careful design and consideration should be taken when implementing VLAN trunks because while most switched networks require trunk links for connectivity, they can also propagate and add to overall network congestion.

The physical way in which you implement trunk links has a great impact on the overall operation of a network. Typically, several end-user nodes are configured with a selected number of access VLANs that are located on the single access-layer switch. End users are provided basic connectivity to their resources through the Building Distribution and Campus Backbone submodules and the trunk links that connect them.

Cisco recommends that you implement redundant links for each Building Access switch. The trunk links carry user VLAN traffic and VLAN information, such as VIDs.

Proper traffic aggregation enhances overall network performance and functionality. Redundant links from the Building Distribution switches to the Campus Backbone switches are recommended to provide multihomed redundancy.

In the figure, there is a switched network environment that is implementing trunk links to take advantage of VLANs and the isolation benefits they can provide. Note that each access switch consists of only one VLAN per switch and that each trunk link only carries traffic for that VLAN. End-user nodes on each VLAN are provided Layer 2 access into the network, which is transparent into the Building Access switch and across the trunk links. In the Building Distribution submodule, multilayer switches provide Layer 3 routing and connectivity into the rest of the network. VLANs are terminated at and do not extend beyond the Layer 3 boundary in the Building Distribution submodule.

Also, VLAN1 should be removed from the trunks to ensure that no user data propagates among the switches on VLAN1. While each Catalyst switch requires VLAN1 on the actual switch and it is not possible to delete, removing VLAN1 from the trunk link is possible and only restricts data frames from being carried across the link. The required management and control frames are still transmitted and received across the trunk link.

The default behavior of VLAN trunks is to permit all VLANs across a trunk link. To further address issues of bandwidth contention and consumption across a trunk link and on VLAN1, a network administrator can either permit or deny any single, range, or group of VLANs to traverse the trunk link. It is best practice to limit the trunk link to only the intended VLANs required for Layer 2 access and connectivity. This improves bandwidth utilization by restricting unwanted VLAN traffic from the link. In addition, explicitly permitting or denying VLANs to a specific trunk link creates a simple deterministic Layer 2 switched domain where there are few unknown variables to cause problems and complicate troubleshooting.

Because of the deterministic nature of the composite network model, DTP should not be required. To fully realize the advantages of a deterministic network, trunk links, encapsulation types, and access ports should be statically configured across specific links according to the network design and requirements.

The Campus Infrastructure module is hierarchical to accommodate the traffic from many Building Access switches into fewer Building Distribution switches. Switch performance capability should increase when fewer devices deal with the combined traffic coming from many devices.

Sometimes, the aggregate traffic of all end users can reach a volume that exceeds the links to the Building Distribution switches, creating performance issues. A solution is to segment the traffic among more Building Access submodules.

The IEEE 802.1Q/p standard presents three types of rules that provide for inherent architectural advantages over ISL.

- **Ingress rules (input of frames to the switch):** Provides rules relevant to the prioritization or classification of received frames belonging to a VLAN
- **Forwarding rules between ports:** Determines whether to filter or forward the frame
- **Egress rules (output of frames from the switch):** Determines if the frame must be sent tagged or untagged

Cisco is now migrating to use 802.1Q as the recommended trunking protocol because of the interoperability and compatibility between the Layer 2 prioritization and Layer 3 prioritization methods, and to support additional spanning tree functionality.

Summary

This topic summarizes the key points discussed in this lesson.

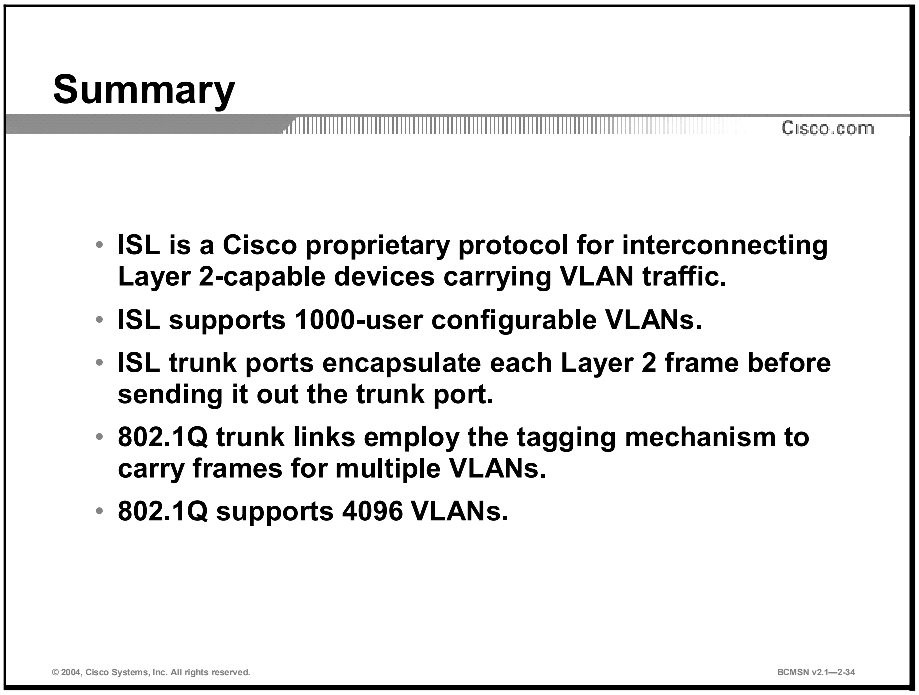

Summary

- ISL is a Cisco proprietary protocol for interconnecting Layer 2-capable devices carrying VLAN traffic.
- ISL supports 1000-user configurable VLANs.
- ISL trunk ports encapsulate each Layer 2 frame before sending it out the trunk port.
- 802.1Q trunk links employ the tagging mechanism to carry frames for multiple VLANs.
- 802.1Q supports 4096 VLANs.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—2-34

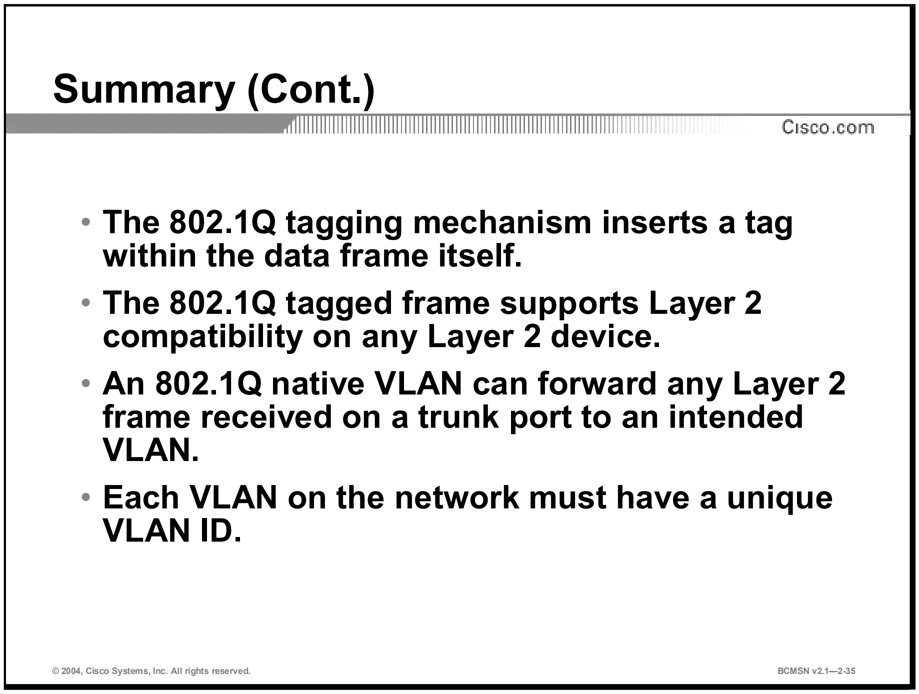

Summary (Cont.)

- The 802.1Q tagging mechanism inserts a tag within the data frame itself.
- The 802.1Q tagged frame supports Layer 2 compatibility on any Layer 2 device.
- An 802.1Q native VLAN can forward any Layer 2 frame received on a trunk port to an intended VLAN.
- Each VLAN on the network must have a unique VLAN ID.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—2-35

References

For additional information, refer to this resource:

- The documentation that accompanied your Cisco Catalyst switch

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Match the correct ISL and 802.1Q features with the appropriate protocol.

_____ 1. ISL

_____ 2. 802.1Q

A) performs encapsulation

B) supports Ethertype VLANs

C) inserts a tag into the data frame

D) is Cisco proprietary

E) is protocol-independent

F) modifies Layer 2 frames

Q2) Select the features that apply to the ISL protocol. (Choose two.)

A) It is nonproprietary.

B) It interconnects Layer 2-capable devices carrying VLAN traffic.

C) ISL trunk ports can handle any type of frame.

D) ISL frame encapsulation uses a mechanism for multiplexing traffic from multiple VLANs on a single link.

Q3) Match the encapsulation steps with the correct process.

_____ 1. ISL encapsulation

_____ 2. Layer 2 encapsulation process

A) encapsulates each Layer 2 frame before sending it out

B) has only one FCS field on the frame

C) original frame encapsulated and an additional header added

D) is media independent

Q4) Select the features that belong to the 802.1Q trunking protocol. (Choose two.)

A) Trunk links employ the tagging mechanism to carry frames for multiple VLANs.

B) It alters the Layer 2 frame structure.

C) It supports native VLAN encapsulated.

D) It encapsulates the Layer 2 frame.

Q5) Match the correct feature with the appropriate protocol.

_____ 1. 802.1Q tagging process

_____ 2. Layer 2 encapsulation protocol

A) supports Layer 2 compatibility on any Layer 2 device

B) cannot determine the Ethertype field

C) inserts a tag within the data frame

D) expects to see the DA and the SA upon receipts of a frame

Q6) What might happen if you misconfigure a native VLAN? (Choose two.)

A) A "VLAN mismatch" error message displays.

B) CDP is turned off.

C) Trunk ports all have different native VLANs.

D) All switch ports and native VLANs are initially assigned to VLAN1.

Q7) Match the correct VLAN range to its use.

_____ 1. 1025-4094

_____ 2. 2-1001

_____ 3. 1002-1005

_____ 4. 1

A) Cisco default. You can use this VLAN but cannot delete it.

B) For Ethernet VLANs. You can create, use, and delete these VLANs.

C) For Ethernet VLANs only.

D) Cisco default for FDDI. You cannot delete these VLANs.

Q8) Select two ways that trunking protocols are used in the Enterprise Composite Network model.

A) One end-user node is configured with a selected number of access VLANs.

B) The same VLAN can exist on multiple switches.

C) Each access switch has several VLANs per switch.

D) Trunk links carry VID.

Quiz Answer Key

Q1) 1=A, D, E; 2=C, B, F

Relates to: Comparing ISL and 802.1Q Trunking Protocols

Q2) B, D

Relates to: Identifying the Features and Benefits of the ISL Protocol

Q3) 1=A, C; 2=B, D

Relates to: Comparing the ISL and Layer 2 Encapsulation Processes

Q4) A, B

Relates to: Identifying the Features and Benefits of 802.1Q Trunking

Q5) 1=A, C; 2=D, B

Relates to: Comparing the 802.1Q Tagging Process and Layer 2 Encapsulation Protocol

Q6) A, B

Relates to: Supporting Native VLANs

Q7) 1=D, 2=B, 3=A, 4=C

Relates to: Mapping VLAN Ranges

Q8) B, D

Relates to: Using Trunking Protocols in the Enterprise Composite Network Model

Configuring Trunking Protocols

Overview

This lesson illustrates how to configure a VLAN trunk. Trunk ports enable connections between switches to carry traffic from more than one VLAN. If trunking is not enabled, the link connecting the two switches will only carry traffic from the VLAN that is configured on the port.

Relevance

Trunking is not required in very simple switched networks with only one VLAN (broadcast domain). However, trunks become necessary in a large multilayer switched networks where traffic between VLANs crosses multiple switches and links.

Objectives

Upon completing this lesson, you will be able to:

- Identify the steps and the correct order that apply to configuring both ISL or 802.1Q trunking protocols
- Correctly identify the commands used to configure an ISL trunk
- Correctly identify the commands used to verify the configuration of an ISL trunk
- Correctly identify the commands used to configure an 802.1Q trunk
- Correctly identify the commands used to verify an 802.1Q trunk configuration
- Identify and resolve a specific trunk link problem in the network

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of VLAN configuration
- Basic knowledge of trunking protocol
- Basic knowledge of switch configuration modes
- Basic knowledge of Ethernet link types
- Basic knowledge of switch port types

Outline

This lesson includes these topics:

- Overview
- Identifying the Procedure for Configuring Trunking Protocols
- Configuring an ISL Trunk
- Verifying the ISL Trunk Configuration
- Configuring an 802.1Q Trunk
- Verifying the 802.1Q Trunk Configuration
- Resolving Trunk Link Problems
- Summary
- Quiz

Identifying the Procedure for Configuring Trunking Protocols

This topic identifies the steps to configure ISL and 802.1Q trunking protocols.

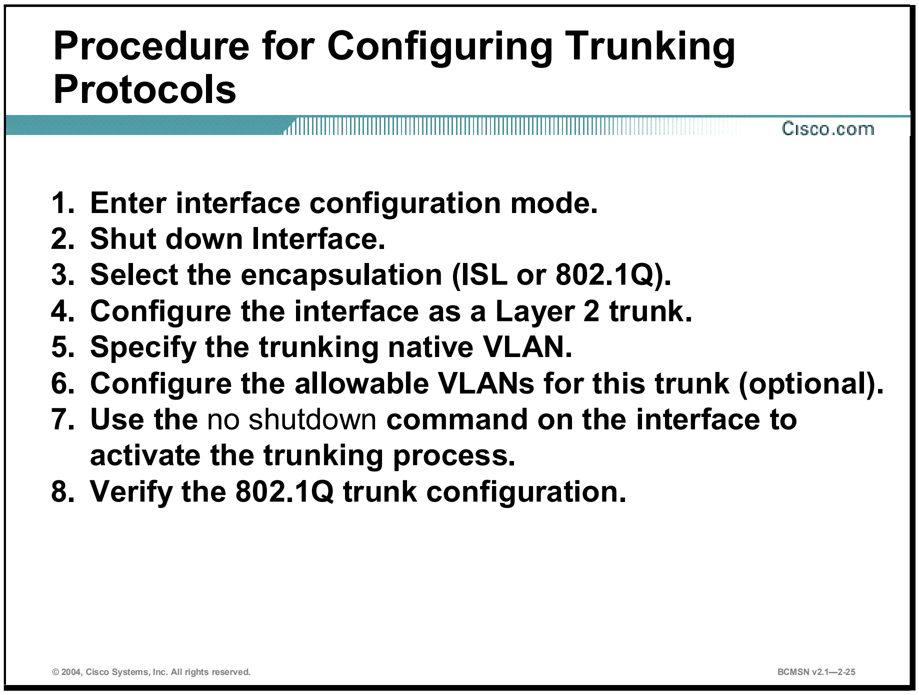

Procedure for Configuring Trunking Protocols

1. **Enter interface configuration mode.**
2. **Shut down Interface.**
3. **Select the encapsulation (ISL or 802.1Q).**
4. **Configure the interface as a Layer 2 trunk.**
5. **Specify the trunking native VLAN.**
6. **Configure the allowable VLANs for this trunk (optional).**
7. **Use the** no shutdown **command on the interface to activate the trunking process.**
8. **Verify the 802.1Q trunk configuration.**

BCMSN v2.1—2-25

Switch ports are configured for ISL trunking using Cisco IOS commands. To configure a switch port as an ISL or 802.1Q trunking port, follow these steps:

Step 1 Enter interface configuration mode.

Step 2 Shut down the interface to prevent the possibility of premature auto configuration taking place.

Step 3 Select the trunking encapsulation. Note some switches support only ISL *or* 802.1Q.

Step 4 Configure the interface as a Layer 2 trunk.

Step 5 Configure the trunking native VLAN number. This number *must* match at both ends of an 802.1Q trunk, but it is not required or significant on an ISL trunk.

Step 6 Configure the allowable VLANs for this trunk. This is necessary if VLANs are restricted to certain trunk links.

Step 7 Use the **no shutdown** command on the interface to activate the trunking process.

Step 8 Verify the ISL trunk configuration using **show IOS** commands.

Caution Ensure that the native VLAN for an 802.1q trunk is the same on both ends of the trunk link. If the native VLAN on one end of the trunk is different from the native VLAN on the other end, the traffic of the native VLANs on both sides cannot be transmitted correctly on the trunk.

Configuring an ISL Trunk

This topic identifies the commands used to configure an ISL trunk.

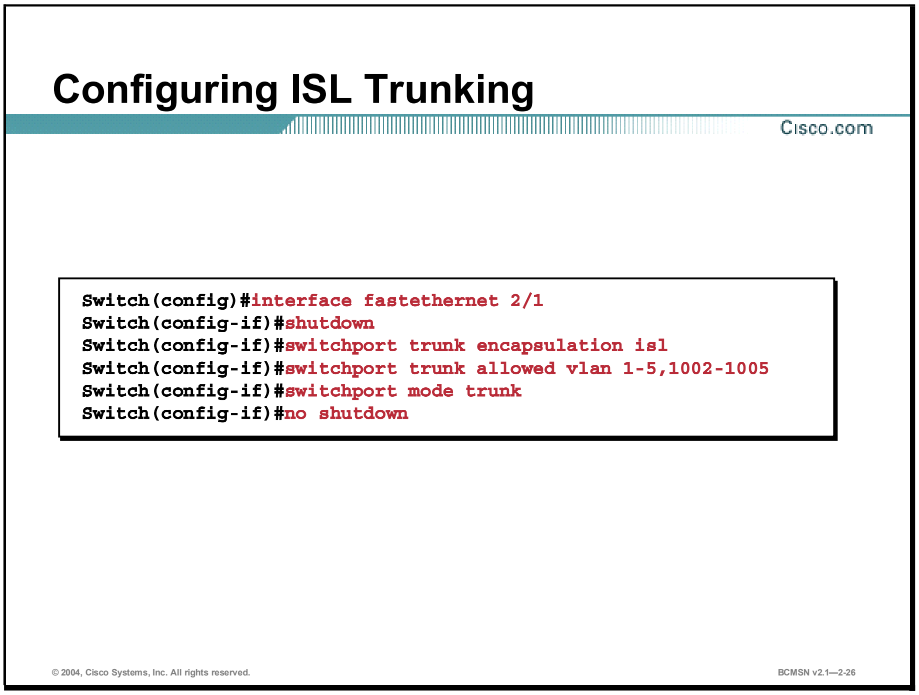

Configuring ISL Trunking

Cisco.com

```
Switch(config)#interface fastethernet 2/1
Switch(config-if)#shutdown
Switch(config-if)#switchport trunk encapsulation isl
Switch(config-if)#switchport trunk allowed vlan 1-5,1002-1005
Switch(config-if)#switchport mode trunk
Switch(config-if)#no shutdown
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-26

To configure a switch port as an ISL trunking port, follow these steps:

Step	Action	Notes
1.	Enter interface configuration mode. `Switch(config)#` **`interface {fastethernet	gigabitethernet}`** *`slot/port`*	Select the interface to configure.
2.	Select the encapsulation. (If multiple encapsulations are supported.) `Switch(config-if)#` **`switchport trunk encapsulation {isl	dot1q	negotiate}`**	This command is optional, unless you configure the port in switchport **trunk** mode. In that case, you must use this command with either the **isl** or **dot1q** argument. **negotiate** is the default. This command is only supported if the switch hardware supports both ISL and dot1q encapsulation.
3.	Configure the allowable VLANs for this trunk. `Switch(config-if)#`**`switchport trunk allowed vlan {add	except	all	remove}`** *`vlan_num1[,vlan_num[,vlan_num[,...]]`*	If not specified, all VLANs are allowed on the trunk. VLANS can be specifically allowed or disallowed.
4.	Configure the interface as a Layer 2 trunk. `Switch(config-if)#` **`switchport mode {dynamic {auto	desirable}	trunk}`**	The switchport mode of the directly connected interface helps determine if the link will perform trunking.

Note A good practice is to shut down an interface while configuring trunking attributes so that premature autonegotiation cannot occur.

When configuring the Layer 2 trunk not to use DTP, note the following syntax information:

- First, enter the **shutdown** command in the interface mode.
- Enter the **switchport trunk encapsulation** command.
- Enter the **switchport mode trunk** command.
- Enter the **switchport nonegotiate** command.
- Finally, enter the **no shutdown** command.

Example: Configuring a Port for ISL Trunking

This example shows how to configure a port for ISL trunking:

```
Switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#interface fastethernet 2/1
Switch(config-if)#shutdown
Switch(config-if)#switchport trunk encapsulation isl
Switch(config-if)#switchport mode trunk
Switch(config-if)#switchport nonegotiate
Switch(config-if)#no shutdown
Switch(config-if)#end
Switch#exit
```

Verifying the ISL Trunk Configuration

This topic identifies the commands used to verify the configuration of an ISL trunk.

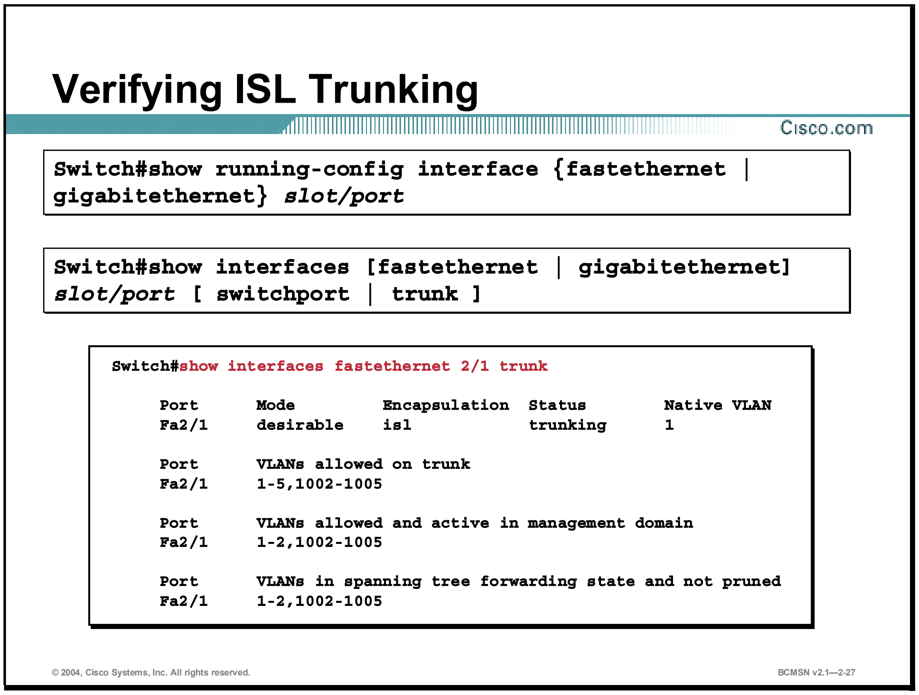

Verifying ISL Trunking

```
Switch#show running-config interface {fastethernet |
gigabitethernet} slot/port
```

```
Switch#show interfaces [fastethernet | gigabitethernet]
slot/port [ switchport | trunk ]
```

```
Switch#show interfaces fastethernet 2/1 trunk

     Port      Mode         Encapsulation  Status        Native VLAN
     Fa2/1     desirable    isl            trunking      1

     Port      VLANs allowed on trunk
     Fa2/1     1-5,1002-1005

     Port      VLANs allowed and active in management domain
     Fa2/1     1-2,1002-1005

     Port      VLANs in spanning tree forwarding state and not pruned
     Fa2/1     1-2,1002-1005
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-27

Use **show** commands to display port information, switch port information, or trunking information.

Configuring an 802.1Q Trunk

This topic identifies the commands used to configure an 802.1Q trunk.

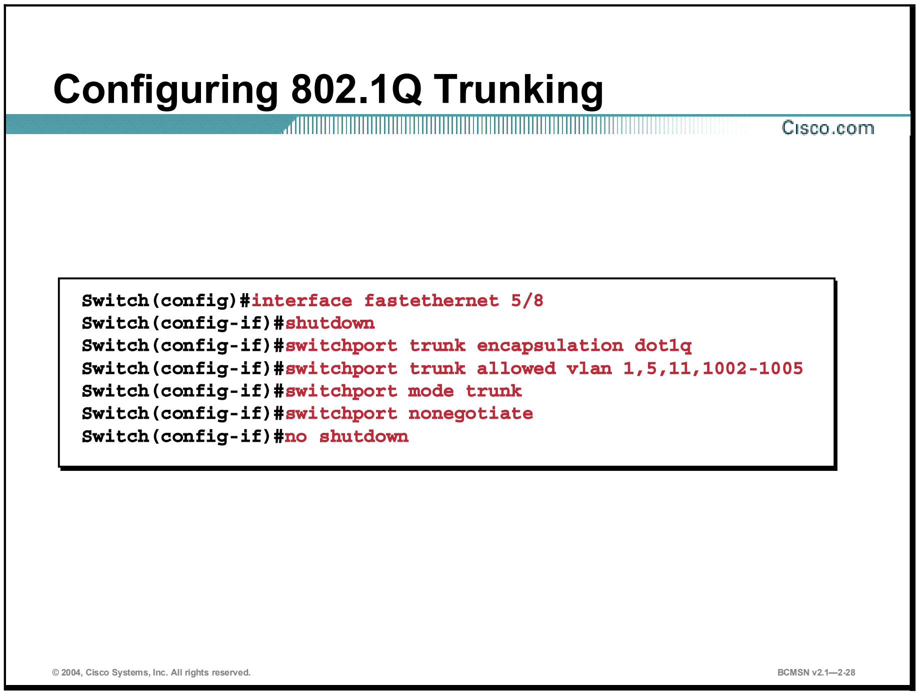

Configuring 802.1Q Trunking

Cisco.com

```
Switch(config)#interface fastethernet 5/8
Switch(config-if)#shutdown
Switch(config-if)#switchport trunk encapsulation dot1q
Switch(config-if)#switchport trunk allowed vlan 1,5,11,1002-1005
Switch(config-if)#switchport mode trunk
Switch(config-if)#switchport nonegotiate
Switch(config-if)#no shutdown
```

BCMSN v2.1—2-28

To configure a switch port as an 802.1Q trunking port, follow these steps:

Step	Action	Notes
1.	Enter interface configuration mode. Switch(config)#**interface {fastethernet | gigabitethernet}** *slot/port*	Select the interface to configure.
2.	Select the encapsulation. Switch(config-if)#**switchport trunk encapsulation {isl | dot1q | negotiate}**	This command is optional, unless you configure the port in switchport **trunk** mode, in which case you must use this command with either the **isl** or **dot1q** argument. **negotiate** is the default.
3.	Configure the interface as a Layer 2 trunk. Switch(config-if)#**switchport mode {dynamic {auto | desirable} | trunk}**	The switchport mode of the directly connected interface helps determine if the link will perform trunking.
4.	Specify the native VLAN. Switch(config-if)#**switchport trunk native vlan** *vlan_num*	The default is VLAN1.
5.	Configure the allowable VLANs for this trunk. Switch(config-if)#**switchport trunk allowed vlan {add | except | all | remove}** *vlan_num1[,vlan_num[,vlan_num[,...]]]*	If not specified, all VLANs are allowed on the trunk. VLANS can be specifically allowed or disallowed.

Example: Configuring a Port for 802.1Q Trunking

This example shows how to configure a port for 802.1Q trunking:

```
Switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#interface fastethernet 5/8
Switch(config-if)#shutdown
Switch(config-if)#switchport mode dynamic desirable
Switch(config-if)#switchport trunk encapsulation dot1q
Switch(config-if)#no shutdown
Switch(config-if)#end
Switch#exit
```

Verifying the 802.1Q Trunk Configuration

This topic identifies the commands used to verify the 802.1Q trunk configuration.

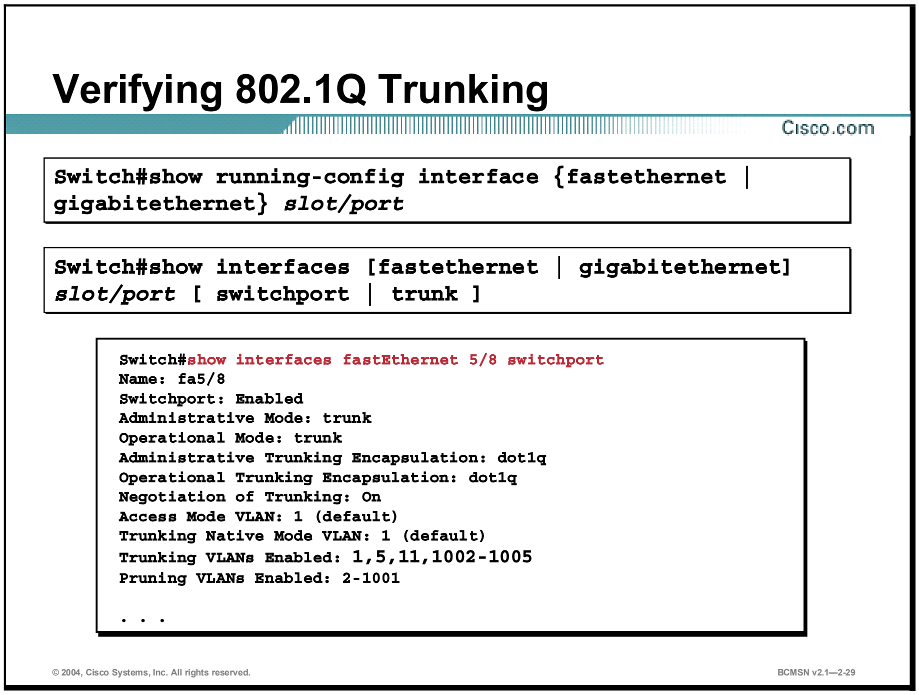

Verifying 802.1Q Trunking

```
Switch#show running-config interface {fastethernet |
gigabitethernet} slot/port
```

```
Switch#show interfaces [fastethernet | gigabitethernet]
slot/port [ switchport | trunk ]
```

```
Switch#show interfaces fastEthernet 5/8 switchport
Name: fa5/8
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: 1,5,11,1002-1005
Pruning VLANs Enabled: 2-1001

. . .
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-29

Use **show** commands to display port information, switch port information, or trunking information.

Example: Displaying Port Information for 802.1Q Trunking

This example shows how to display port information for 802.1Q trunking:

```
Switch#show running-config interface fastethernet 5/8
Building configuration...
Current configuration:
!
interface FastEthernet5/8
 switchport mode dynamic desirable
 switchport trunk encapsulation dot1q
end
```

Example: Displaying Switch Port Information for 802.1Q Trunking

This example shows how to display switch port information for 802.1Q trunking:

```
Switch#show interfaces fastethernet 5/8 switchport

Name: Fa5/8

Switchport: Enabled

Administrative Mode: dynamic desirable

Operational Mode: trunk

Administrative Trunking Encapsulation: negotiate

Operational Trunking Encapsulation: dot1q

Negotiation of Trunking: Enabled

Access Mode VLAN: 1 (default)

Trunking Native Mode VLAN: 1 (default)

Trunking VLANs Enabled: ALL

Pruning VLANs Enabled: 2-1001
```

Example: Displaying Trunk Information for 802.1Q Trunking

This example shows how to display trunk information for 802.1Q trunking. Notice that the encapsulation type is n-802.1q, showing that DTP negotiated the trunking protocol with the other switch.

```
Switch#show interfaces fastethernet 5/8 trunk


Port      Mode         Encapsulation  Status        Native
vlan

Fa5/8     desirable    n-802.1q       trunking      1


Port      Vlans allowed on trunk

Fa5/8 1-1005


Port      Vlans allowed and active in management domain

Fa5/8 1-6,10,20,50,100,152,200,300,303-305,349-
351,400,500,521,524,570,801-8

02,850,917,999,1002-1005


Port      Vlans in spanning tree forwarding state and not
pruned

Fa5/8 1-6,10,20,50,100,152,200,300,303-305,349-
351,400,500,521,524,570,801-8

02,850,917,999,1002-1005


Switch#
```

Resolving Trunk Link Problems

This topic describes how to resolve a specific trunk link problem.

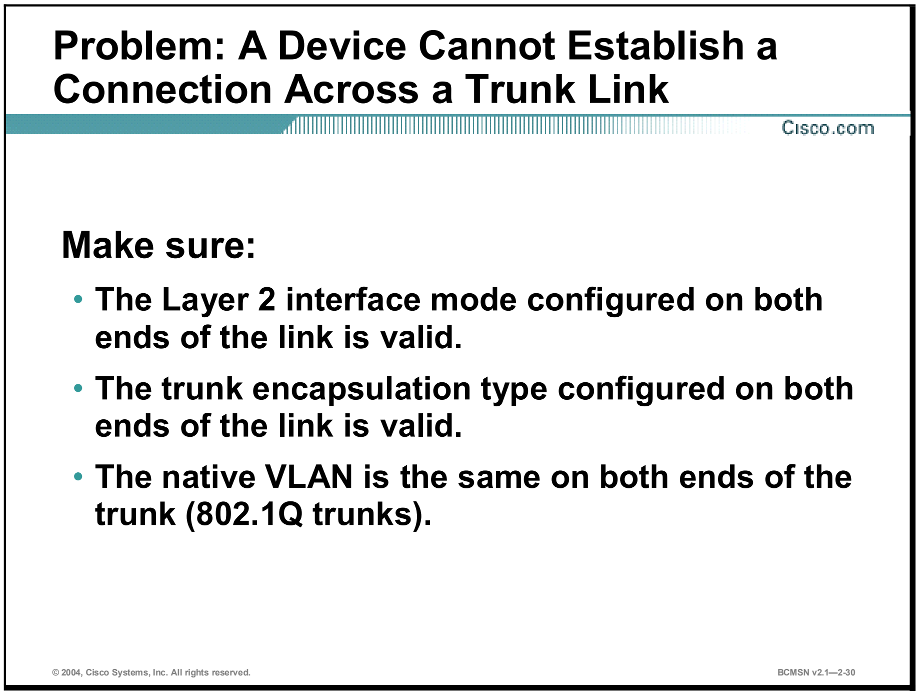

Problem: A Device Cannot Establish a Connection Across a Trunk Link

Make sure:

- The Layer 2 interface mode configured on both ends of the link is valid.
- The trunk encapsulation type configured on both ends of the link is valid.
- The native VLAN is the same on both ends of the trunk (802.1Q trunks).

BCMSN v2.1—2-30

If a problem exists with a trunking link, make sure the interface modes, encapsulation types, and native VLANs are correct on both sides of the link.

Problem: A device cannot establish a connection across a trunk link.

Suggested solutions to the problem are as follows:

- Make sure the Layer 2 interface mode configured on both ends of the link is valid. The interface mode should be trunk, dynamic, or nonegotiate. Use **show** commands to verify the configuration.
- Use **show** commands to verify the trunk encapsulation type configured on both ends of the link is valid and compatible.
- On IEEE 802.1Q trunks, use **show** commands to verify the native VLAN is the same on both ends of the trunk.

Summary

This topic summarizes the key points discussed in this lesson.

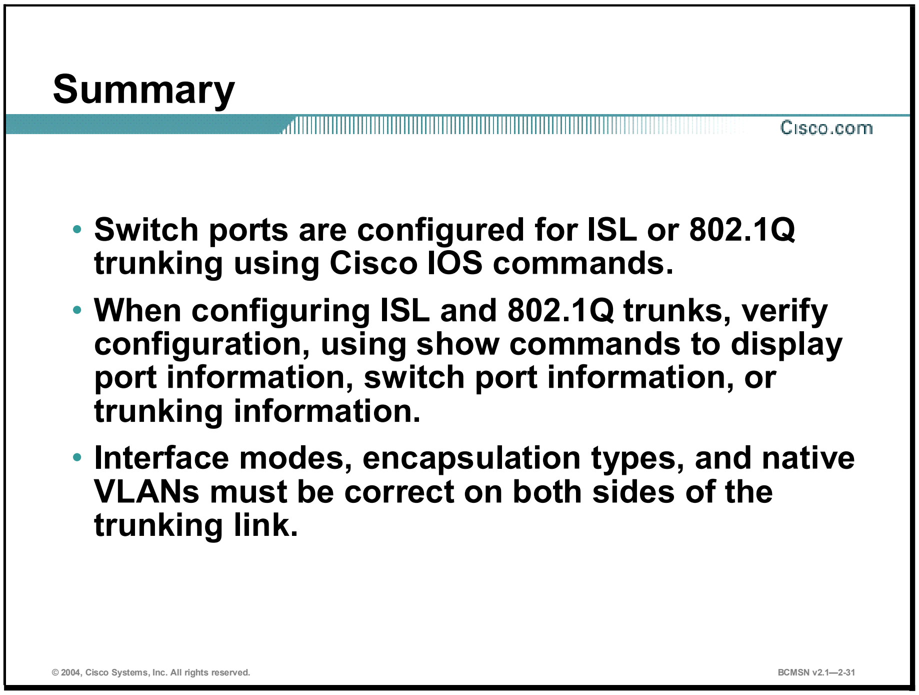

Summary

Cisco.com

- Switch ports are configured for ISL or 802.1Q trunking using Cisco IOS commands.
- When configuring ISL and 802.1Q trunks, verify configuration, using show commands to display port information, switch port information, or trunking information.
- Interface modes, encapsulation types, and native VLANs must be correct on both sides of the trunking link.

BCMSN v2.1—2-31

References

For additional information, refer to this resource:

- The documentation that accompanied your Cisco Catalyst switch

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Match the appropriate trunking protocol configuration step with the correct order.

_____ 1. Step 1

_____ 2. Step 2

_____ 3. Step 3

_____ 4. Step 4

_____ 5. Step 5

_____ 6. Step 6

_____ 7. Step 7

_____ 8. Step 8

A) Specify the trunking native VLAN.

B) Enter interface configuration mode.

C) Configure the interface as a Layer 2 trunk.

D) Verify the 802.1Q trunk configuration.

E) Select the encapsulation (ISL or 802.1Q).

F) Use the **no shutdown** command on the interface to activate the trunking process.

G) Configure the allowable VLANs for this trunk (optional).

H) Shut down the interface to eliminate the possibility of premature auto configuration taking place.

Q2) Select the command to configure the allowable VLANs for a specific trunk.

A) switch(config)#**switchport trunk allowed vlan {add | except | all | remove}**

B) switch(config-if)#**switchport trunk vlan {add | except | all | remove}**

C) switch(config-if)#**switchport trunk allowed vlan {add | except | all | remove}**

D) switch(config-no)#**switchport trunk allowed vlan {add | except | all | remove}**

Q3) Select the command that displays port information to verify the ISL trunk configuration.

A) Switch#**show running-config interface {fastethernet | gigabitethernet}** *slot/port*

B) Switch(config)#**interface {fastethernet | gigabitethernet}** *slot/port*

C) Switch(config-if)#**switchport mode dynamic desirable**

D) Switch#**show interfaces fastethernet 5/8 trunk**

Q4) Identify the correct order for the commands to configure an 802.1Q trunk.

_____ 1. Step 1

_____ 2. Step 2

_____ 3. Step 3

_____ 4. Step 4

_____ 5. Step 5

A) Switch(config-if)#**switchport trunk allowed vlan {add | except | all | remove}**

B) Switch(config-if)#**switchport trunk encapsulation {isl | dot1q | negotiate}**

C) Switch(config)#**interface {fastethernet | gigabitethernet} slot/port**

D) Switch(config-if)#**switchport trunk native vlan vlan_num**

E) Switch(config-if)#**switchport mode {dynamic {auto | desirable} | trunk}**

Q5) What command do you use to display port information to verify the 802.1Q trunk configuration?

A) Switch#**show running-config interface fastethernet slot/port**

B) Switch#**configure terminal**

C) Switch(config)#**interface {fastethernet | gigabitethernet} slot/port**

D) Switch#**show running-config interface {fastethernet | gigabitethernet} slot/port**

Q6) If the Layer 2 interface mode on one side of a link is set to dynamic auto, a trunk will be established if the directly connected interface is configured with either of which two interface modes? (Choose two.)

A) trunk

B) access

C) nonegotiate

D) dynamic auto

E) dynamic desirable

Quiz Answer Key

Q1) 1=B; 2=H, 3=E, 4=C, 5=A, 6=G, 7=F, 8=D

Relates to: Identifying the Procedure for Configuring Trunking Protocols

Q2) C

Relates to: Configuring an ISL Trunk

Q3) A

Relates to: Verifying the ISL Trunk Configuration

Q4) 1=C, 2=B, 3=E, 4=D, 5=A

Relates to: Configuring an 802.1Q Trunk

Q5) A

Relates to: Verifying the 802.1Q Trunk Configuration

Q6) A, E

Relates to: Resolving Trunk Link Problems

Maintaining VLAN Consistency Across the Network

Overview

VTP is used to distribute and synchronize information about VLANs configured throughout a switched network. VTP reduces the manual configuration needed in the network.

Relevance

In a large switched network, VTP allows you to manage VLAN implementation.

Objectives

Upon completing this lesson, you will be able to:

- Correctly identify the features and functions that apply to VTP
- Match the features that apply to the correct VTP mode
- Correctly list those steps on how VTP shares VLAN information in a management domain, in order of occurrence
- Correctly identify which functions are used by VTP to control traffic within a management domain
- Match the correct feature to the appropriate VTP version
- Select the guidelines used when configuring a VTP management domain
- Correctly match the configuration steps to the appropriate VTP mode
- Correctly identify the state of the VTP configuration
- Select the guidelines most likely to be used to troubleshoot a VTP configuration problem
- Identify the benefits most likely to occur as a result of applying VTP within an Enterprise Composite Network model

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)
- Basic knowledge of the components of the Enterprise Composite Network model
- Basic knowledge of VLAN operations
- Basic knowledge of VLAN configuration
- Basic knowledge of switch operations
- Basic knowledge of trunking protocols

Outline

This lesson includes these topics:

- Overview
- Defining VTP
- Identifying VTP Modes
- Sharing VLAN Attributes in a Management Domain
- Controlling VLAN Traffic on a Trunk
- Distinguishing Between VTP Versions 1 and 2
- Identifying the Procedure for Configuring VTP
- Creating a VTP Management Domain
- Verifying the VTP Configuration
- Selecting a Troubleshooting Approach for VTP
- Identifying How VTP Is Used in the Enterprise Composite Network Model
- Summary
- Quiz

Defining VTP

This topic identifies the features and functions that apply to VTP.

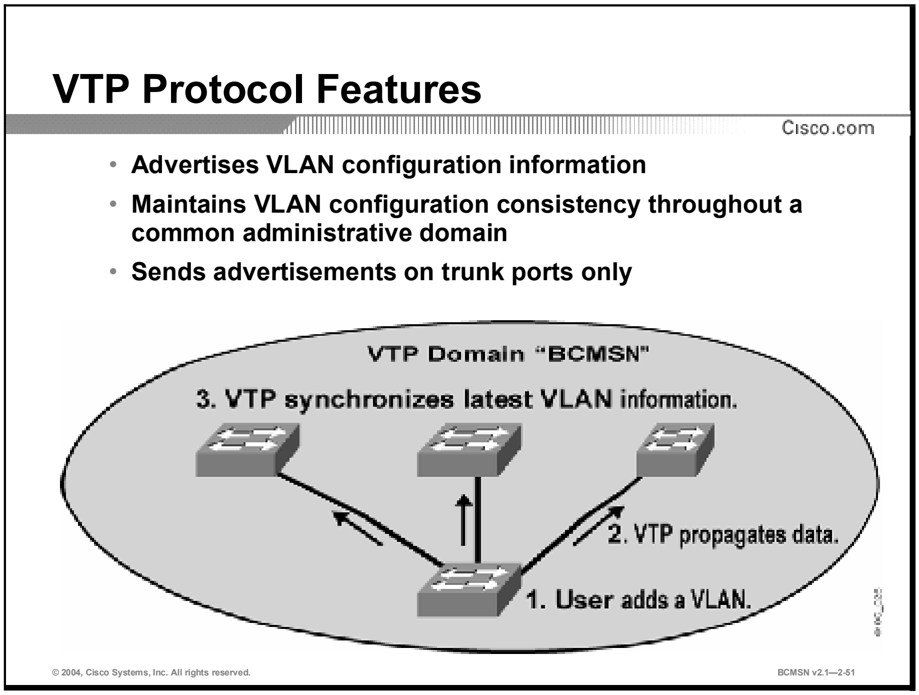

VTP Protocol Features

Cisco.com

- **Advertises VLAN configuration information**
- **Maintains VLAN configuration consistency throughout a common administrative domain**
- **Sends advertisements on trunk ports only**

BCMSN v2.1—2-51

VTP is a protocol used to distribute and synchronize information about VLANs configured throughout a switched network. VTP minimizes misconfigurations and configuration inconsistencies that can result in a number of problems, such as duplicate VLAN names, incorrect VLAN-type specifications, and security violations.

VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the additions, deletions, and name changes of VLANs within a VTP domain.

A VTP domain is one switch or several interconnected switches sharing the same VTP environment. You can only configure a switch to be in one VTP domain.

By default, a Catalyst switch is in the no-management-domain state until it receives an advertisement for a domain over a trunk link, or until you configure a management domain. Configurations made to a single VTP server are propagated across links to all connected switches in the network by taking these steps:

Step 1 An administrator adds a new VLAN definition.

Step 2 VTP propagates the VLAN information to all switches in the VTP domain.

Step 3 Each switch synchronizes its configuration to incorporate the new VLAN data.

Identifying VTP Modes

This topic identifies the features associated with the three VTP modes.

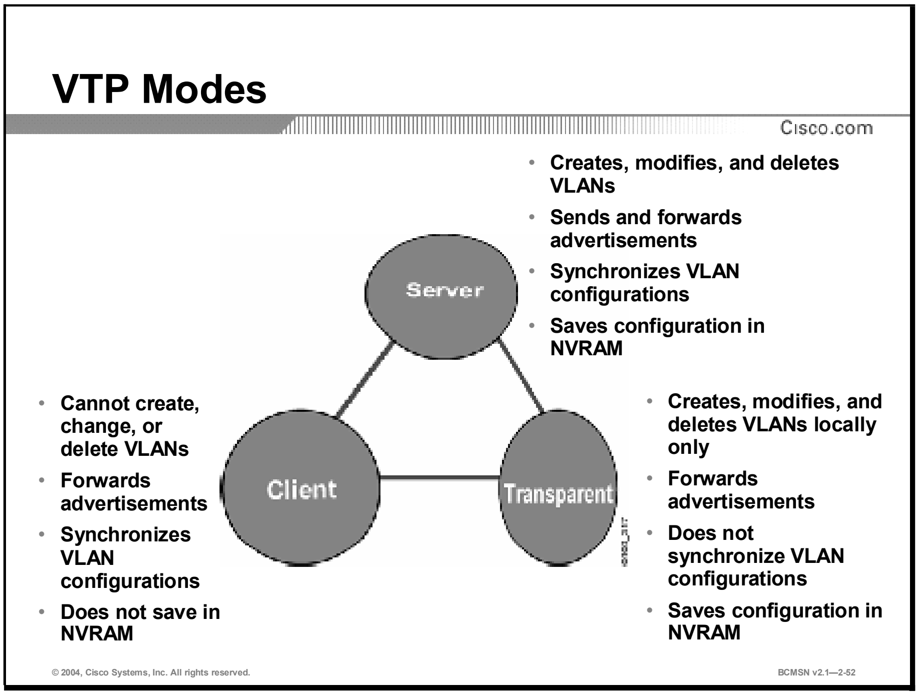

VTP Modes

- Creates, modifies, and deletes VLANs
- Sends and forwards advertisements
- Synchronizes VLAN configurations
- Saves configuration in NVRAM

- Cannot create, change, or delete VLANs
- Forwards advertisements
- Synchronizes VLAN configurations
- Does not save in NVRAM

- Creates, modifies, and deletes VLANs locally only
- Forwards advertisements
- Does not synchronize VLAN configurations
- Saves configuration in NVRAM

BCMSN v2.1—2-52

VTP operates in one of three modes: server mode, client mode, or transparent mode. The default VTP mode is server mode, but VLANs are not propagated over the network until a management domain name is specified or learned.

You can complete different tasks depending on the VTP operation mode. VTP messages are transmitted out all trunk connections.

The table describes the features of the VTP client, server, and transparent modes.

VTP Mode	Features
Client	■ Cannot create, change, or delete VLANs on command line interface (CLI) ■ Forwards advertisements to other switches ■ Synchronizes VLAN configuration with latest information received from other switches in the management domain ■ Does not save VLAN configuration in nonvolatile random-access memory (NVRAM)
Server	■ Create, modifies, and deletes VLANs ■ Sends and forwards advertisements to other switches ■ Synchronizes VLAN configuration with latest information received from other switches in the management domain ■ Saves VLAN configuration in NVRAM

VTP Mode	Features
Transparent	■ Creates, deletes, and modifies VLANs only on the local switch ■ Forwards VTP advertisements received from other switches in the same management domain ■ Does not synchronize its VLAN configuration with information received from other switches in the management domain ■ Saves VLAN configuration in NVRAM

Sharing VLAN Attributes in a Management Domain

This topic identifies the steps on how VTP distributes VLAN attributes throughout a VTP domain.

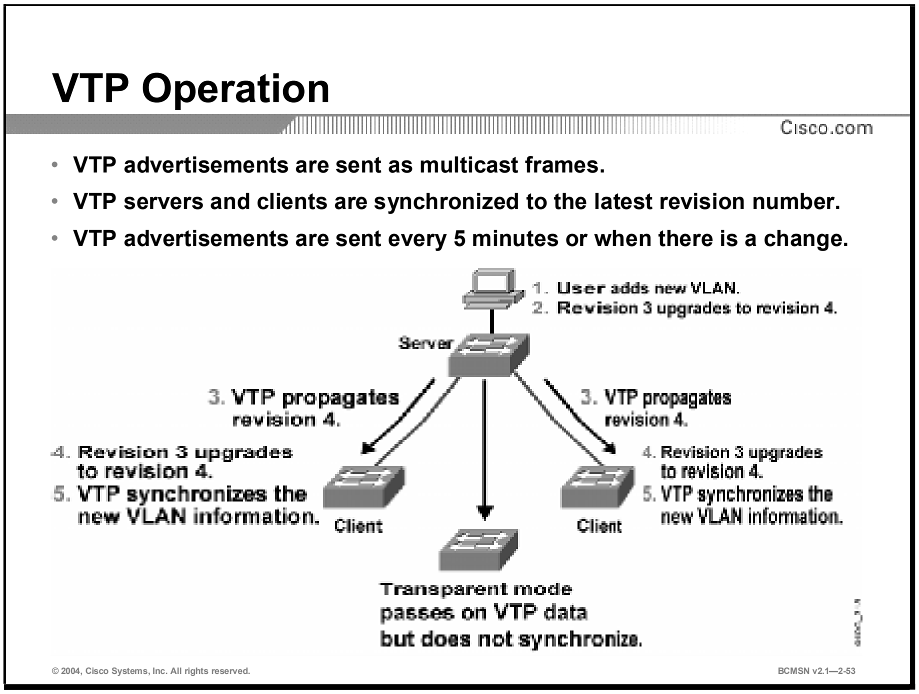

VTP Operation

Cisco.com

- VTP advertisements are sent as multicast frames.
- VTP servers and clients are synchronized to the latest revision number.
- VTP advertisements are sent every 5 minutes or when there is a change.

1. User adds new VLAN.
2. Revision 3 upgrades to revision 4.
Server
3. VTP propagates revision 4.
3. VTP propagates revision 4.
4. Revision 3 upgrades to revision 4.
5. VTP synchronizes the new VLAN information.
Client
Client
4. Revision 3 upgrades to revision 4.
5. VTP synchronizes the new VLAN information.
Transparent mode passes on VTP data but does not synchronize.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-53

VTP advertisements are flooded throughout the management domain. VTP advertisements are sent every five minutes or whenever there is a change in VLAN configurations, and they are transmitted over the management VLAN (default VLAN1) using a multicast frame.

Note VTP propagates VLAN configuration information, not switch port configuration information.

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

- First, the management domain name and password in the advertisement must match those configured in the local switch.
- Next, if the configuration revision number indicates that the message was created after the configuration currently in use and the switch is a VTP server or client, the switch incorporates the advertised VLAN information.

One of the most critical components of VTP is the configuration revision number. Each time a VTP server modifies its VLAN information, it increments the configuration revision number by one. It then sends out a VTP advertisement with the new configuration revision number. If the configuration revision number being advertised is higher than the number stored on the other switches in the VTP domain, they will overwrite their VLAN configurations with the new information being advertised.

A VTP transparent switch does not participate in VTP, in that it does not advertise its VLAN configuration or synchronize its VLAN database on receipt of a VTP advertisement. Because of this, the configuration revision number in VTP transparent mode is always 0.

Note The overwrite process would mean that if the VTP server deleted all VLANs and advertised with a higher revision number, the client devices in the VTP domain would also delete their VLANs.

Controlling VLAN Traffic on a Trunk

This topic identifies the functions used by VTP to control traffic within a management domain.

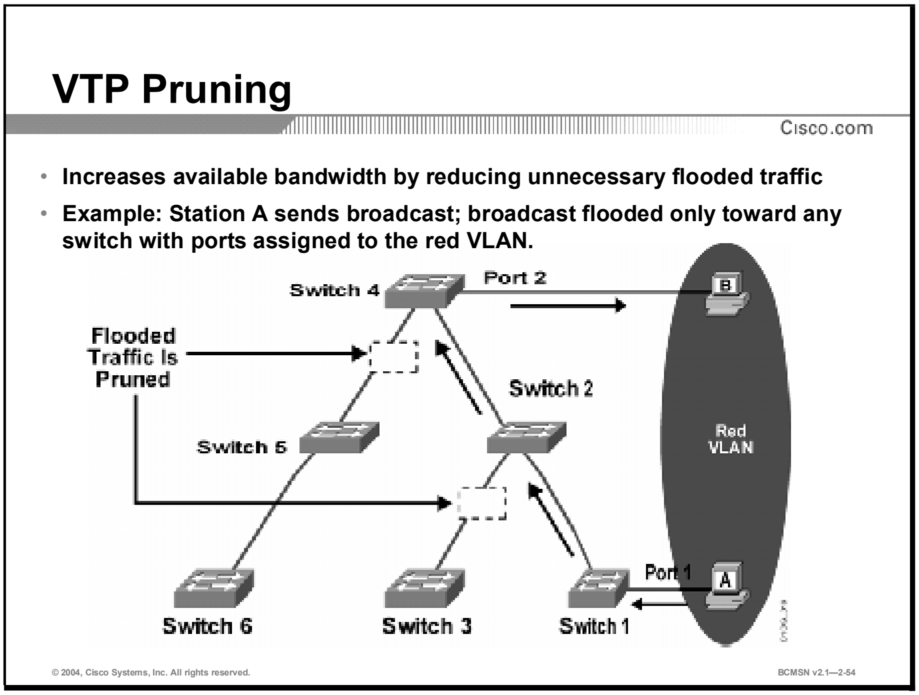

VTP Pruning

Cisco.com

- **Increases available bandwidth by reducing unnecessary flooded traffic**
- **Example: Station A sends broadcast; broadcast flooded only toward any switch with ports assigned to the red VLAN.**

BCMSN v2.1—2-54

VTP Pruning uses VLAN advertisements to determine when a trunk connection is flooding traffic needlessly.

By default, a trunk connection carries traffic for all VLANs in the VTP management domain. Commonly, some switches in an enterprise network do not have local ports configured in each VLAN. In the example, switches 1 and 4 support ports statically configured in the red VLAN.

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

The example shows a switched network with VTP Pruning enabled. The broadcast traffic from station A is not forwarded to switches 3, 5, and 6 because traffic for the red VLAN has been pruned on the links indicated on switches 2 and 4.

Note	You can implement VTP Pruning only on VTP servers, not on clients. Consider VTP Pruning support to minimize traffic on trunk links.

Be aware that VTP Pruning may have a negative impact on network update performance.

Note	A switch runs an instance of spanning tree for each VLAN it is aware of, even if no ports are active or if VTP Pruning is enabled. VTP Pruning prevents unnecessary flooded traffic but does not eliminate the switch knowledge of pruned VLANs. Spanning tree implementation is discussed in the module “Implementing Spanning Tree Protocol.”

Distinguishing Between VTP Versions 1 and 2

This topic compares the features of VTP versions 1 and 2.

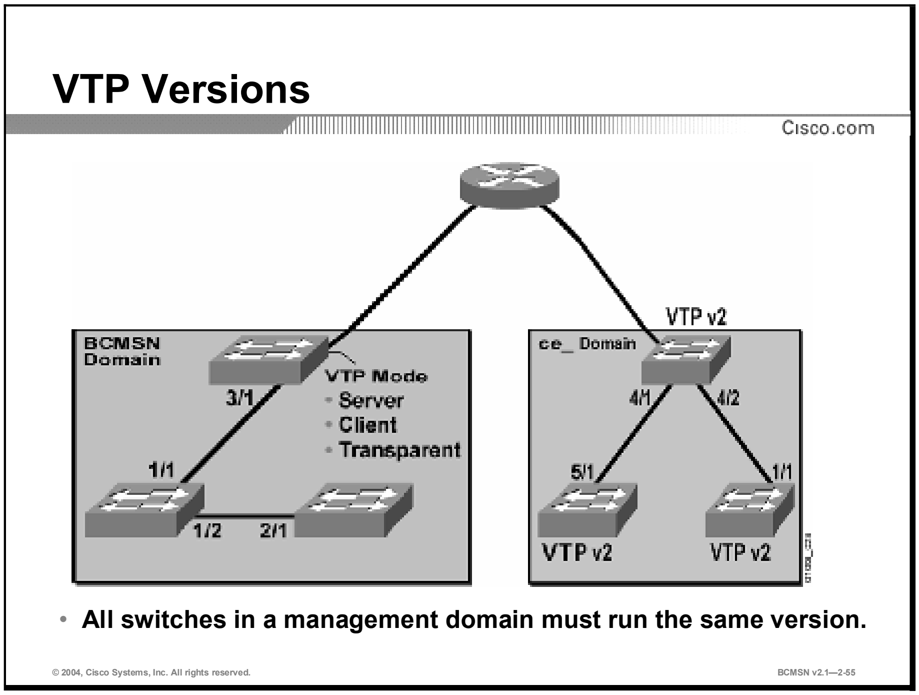

VTP Versions

Cisco.com

- **All switches in a management domain must run the same version.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-55

If you use VTP in your network, you must decide whether to use VTP version 1 or version 2. VTP version 1 is supported in Supervisor Engine software release 2.1 or later. VTP version 2 is supported in Supervisor Engine software release 3.1(1) and later.

Two different versions of VTP (version 1 and version 2) can run in your management domain. These two versions are not interoperable. If you choose to configure a switch in a domain for VTP version 2, you must configure all switches in the management domain to be in VTP version 2. VTP version 1 is the default. You may need to implement VTP version 2 if you need some of the specific features that VTP version 2 offers that are not offered in VTP version 1. A switch that is capable of running VTP version 2 can operate in the same domain as a switch running VTP version 1 if VTP version 2 remains disabled on the VTP version 2-capable switch.

VTP version 2 supports these features not supported in VTP version 1:

- **Token Ring support:** VTP version 2 supports Token Ring LAN switching and VLANs.
- **Unrecognized type, length, value (TLV) support:** A VTP server or client propagates configuration changes to its other trunks, even for TLVs that it is not able to parse. The unrecognized TLV is saved in NVRAM.
- **Version-independent transparent mode:** In VTP version 1, a VTP transparent switch inspects VTP messages for the domain name and version, forwarding a message only if the version and domain name match. Because only one domain is supported in the Supervisor Engine software, VTP version 2 forwards VTP messages in transparent mode, without checking the version.
- **Consistency checks:** In VTP version 2, VLAN consistency checks (such as VLAN names and values) are performed only when you enter new information through the command-line interface (CLI) or Simple Network Management Protocol (SNMP). Consistency checks are not performed when new information is obtained from a VTP message or when information is read from NVRAM. If the Message Digest 5 (MD5) on a received VTP message is correct, its information is accepted without consistency checks.

You must use VTP version 2 if you are running VTP in a Token Ring environment.

If all switches in a domain are capable of running VTP version 2, you need only enable VTP version 2 on one switch. The version number is propagated to the other VTP version 2-capable switches in the VTP domain.

Identifying the Procedure for Configuring VTP

This topic identifies the guidelines used to configure a VTP management domain.

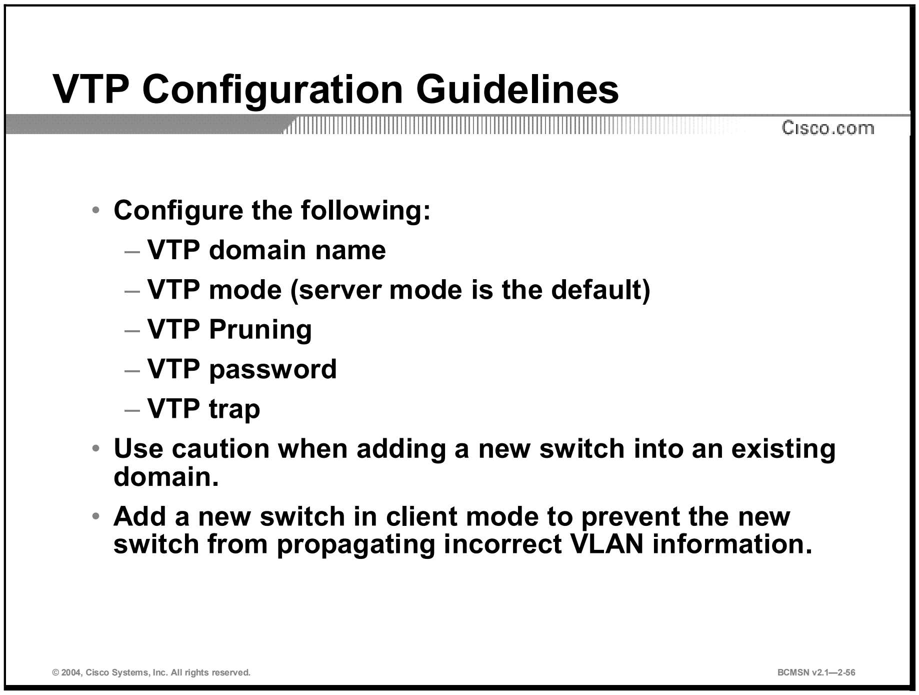

VTP Configuration Guidelines

- Configure the following:
  - VTP domain name
  - VTP mode (server mode is the default)
  - VTP Pruning
  - VTP password
  - VTP trap
- Use caution when adding a new switch into an existing domain.
- Add a new switch in client mode to prevent the new switch from propagating incorrect VLAN information.

BCMSN v2.1—2-56

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

Default VTP configuration values depend on the switch model and software version. For example, the default values for the Catalyst 4000 and 6000 series switches are as follows:

- **VTP domain name:** None
- **VTP mode:** Server
- **VTP pruning:** Disabled
- **VTP password:** None
- **VTP trap:** Disabled (SNMP traps communicating VTP status)

The VTP domain name can be specified or learned. By default, the domain name is not set.

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

VTP Pruning eligibility is one VLAN parameter that the VTP protocol advertises. Enabling or disabling VTP Pruning on a VTP server propagates the change throughout the management domain.

Creating a VTP Management Domain

This topic matches the configuration steps with the appropriate VTP mode.

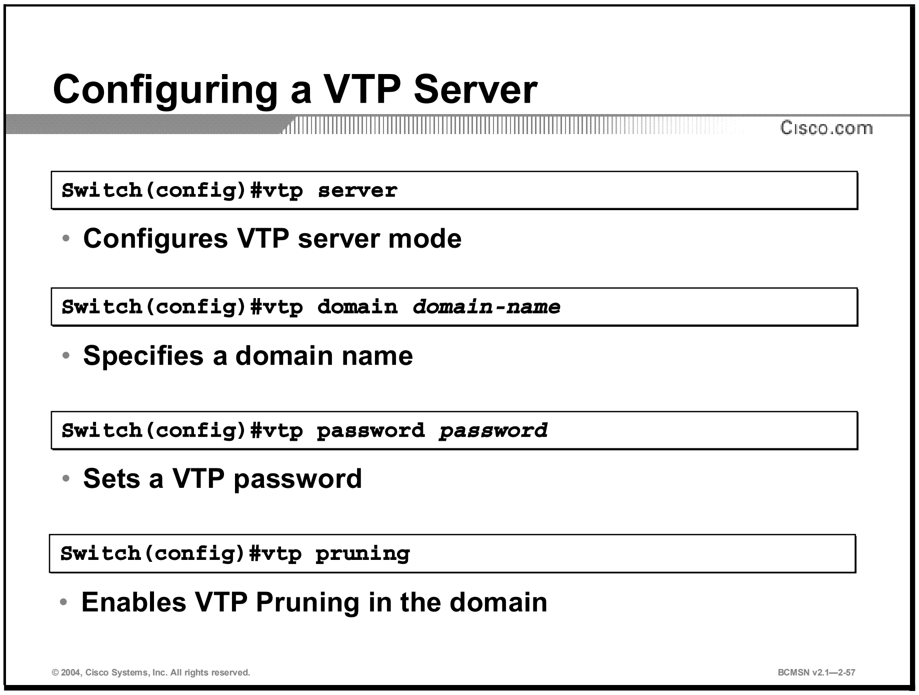

Configuring a VTP Server

Cisco.com

```
Switch(config)#vtp server
```

- Configures VTP server mode

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

- Specifies a domain name

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

- Sets a VTP password

```
Switch(config)#vtp pruning
```

- Enables VTP Pruning in the domain

BCMSN v2.1—2-57

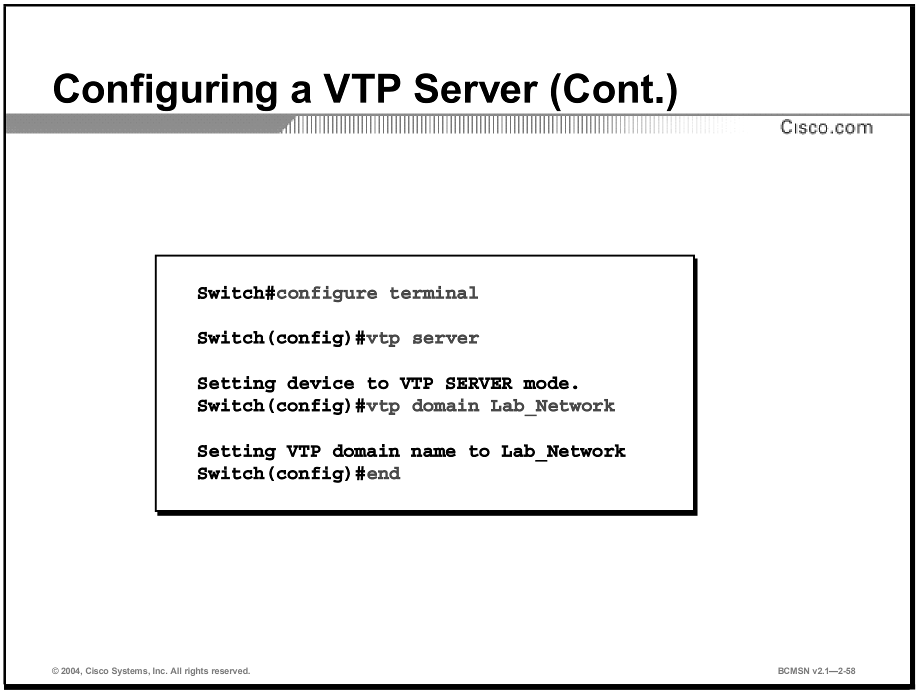

Configuring a VTP Server (Cont.)

Cisco.com

```
Switch#configure terminal

Switch(config)#vtp server

Setting device to VTP SERVER mode.
Switch(config)#vtp domain Lab_Network

Setting VTP domain name to Lab_Network
Switch(config)#end
```

BCMSN v2.1—2-58

To configure a VTP server, follow these steps from privileged EXEC mode:

Step	Action	Notes
1.	Enter global configuration mode. Switch#**configure terminal**	
2.	Configure the VTP mode as server. Switch(config)#**vtp server**	To revert to the default (server), enter **no vtp mode**.
3.	Configure the domain name. Switch(config)#**vtp domain** *domain_name*	Defines the VTP domain name, which can be up to 32 characters long. VTP server mode requires a domain name. If the switch has a trunk connection to a VTP domain, the switch learns the domain name from the VTP server in the domain. There is no option to clear the domain name.
4.	Enable VTP version 2. Switch(config)#**vtp v2-mode**	To revert to the default (VTP version 1), enter **vtp v1-mode**.
5.	Specify a VTP password. Switch(config)#**vtp password** *password_string*	Sets a password, which can be from 8 to 64 characters long, for the VTP domain. Use **no vtp password** to clear the password.
6.	Enable VTP Pruning in the management domain. Switch(config)#**vtp pruning**	Use the **no** version of this command to turn off VTP pruning.
7.	Exit global configuration mode. Switch(config)#**exit**	

To configure a VTP client, follow these steps from configuration mode:

Step	Action	Notes
1.	Enter global configuration mode. Switch#**configure terminal**	
2.	Configure the VTP mode as client. Switch(config)#**vtp client**	To revert to the default (server), enter **no vtp mode**.
3.	Exit global configuration mode. Switch(config)#**exit**	

Example: Configuring a Switch as a VTP Server

This example shows how to configure the switch as a VTP server:

```
Switch#configure terminal

Switch(config)#vtp mode server

Setting device to VTP SERVER mode.
Switch(config)#vtp domain Lab_Network

Setting VTP domain name to Lab_Network
Switch(config)#end
```

Verifying the VTP Configuration

This topic identifies the state of the VTP configuration using the output to the **show vtp status** command.

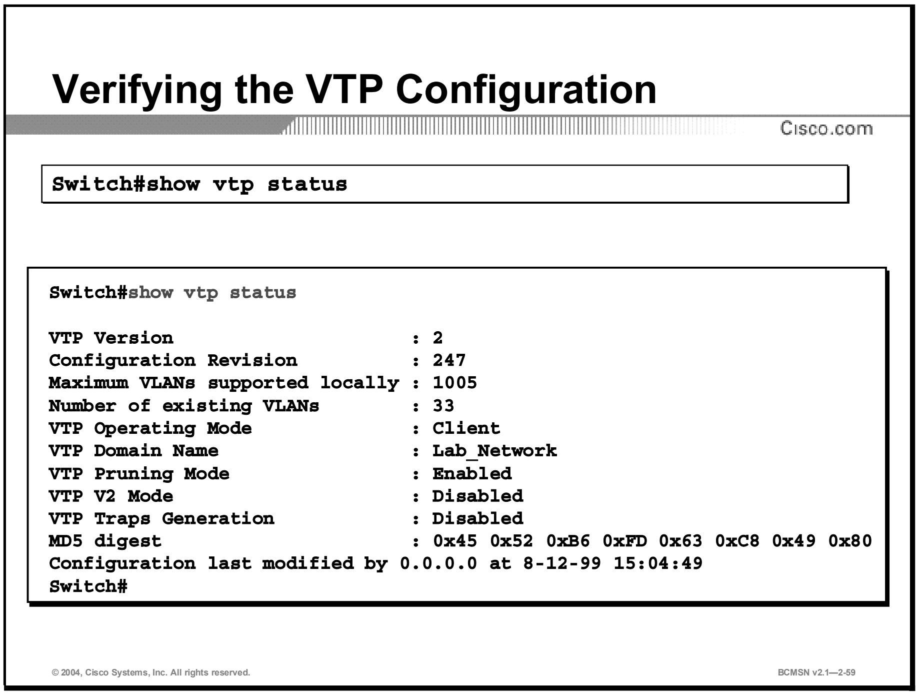

Verifying the VTP Configuration

```
Switch#show vtp status
```

```
Switch#show vtp status

VTP Version                     : 2
Configuration Revision          : 247
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 33
VTP Operating Mode              : Client
VTP Domain Name                 : Lab_Network
VTP Pruning Mode                : Enabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x45 0x52 0xB6 0xFD 0x63 0xC8 0x49 0x80
Configuration last modified by 0.0.0.0 at 8-12-99 15:04:49
Switch#
```

BCMSN v2.1—2-59

Use **show** commands to verify the VTP configuration. Use the **show vtp status** command to display information about the VTP configuration.

Note In this example, VTP version 2 is available (as shown by the "VTP Version" line of the output), but not enabled (as shown by the "VTP V2 Mode" line of the output).

Example: Displaying VTP Status

This example shows how to verify the VTP configuration using the **show vtp status** command. The output describes the VTP version, number of VLANs supported locally, VTP operating mode, VTP domain name, and VTP Pruning mode.

```
Switch#show vtp status


VTP Version                     : 2
Configuration Revision          : 247
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 33
VTP Operating Mode              : Server
VTP Domain Name                 : Lab_Network
VTP Pruning Mode                : Enabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x45 0x52 0xB6 0xFD 0x63
0xC8 0x49 0x80
Configuration last modified by 0.0.0.0 at 8-12-99 15:04:49
Local updater ID is 172.20.52.34 on interface Gi1/1 (first
interface found)
Switch#
```

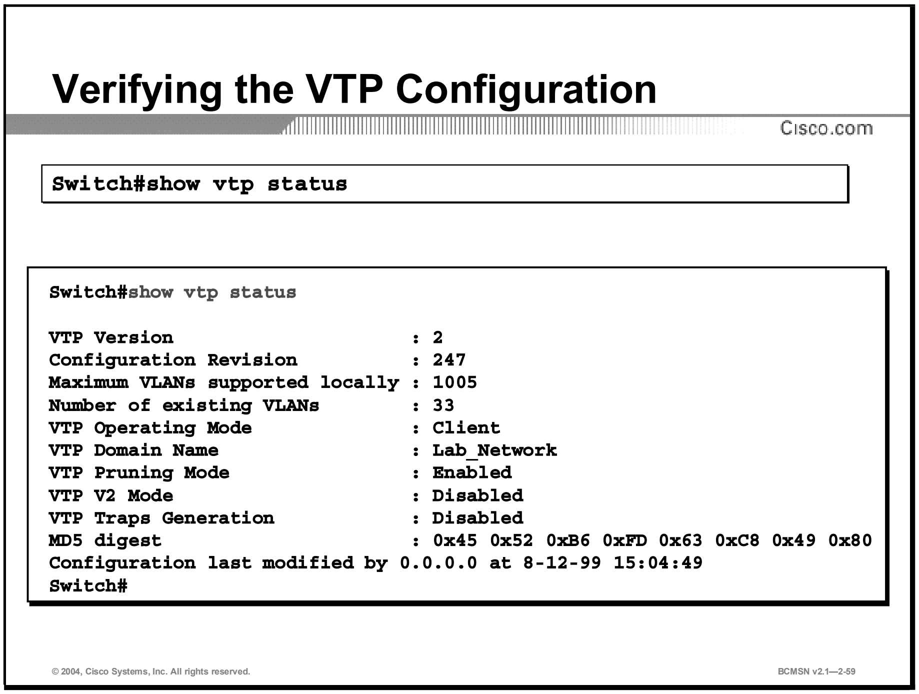

Verifying the VTP Configuration

Cisco.com

```
Switch#show vtp status
```

```
Switch#show vtp status

VTP Version                     : 2
Configuration Revision          : 247
Maximum VLANs supported locally : 1005
Number of existing VLANs        : 33
VTP Operating Mode              : Client
VTP Domain Name                 : Lab_Network
VTP Pruning Mode                : Enabled
VTP V2 Mode                     : Disabled
VTP Traps Generation            : Disabled
MD5 digest                      : 0x45 0x52 0xB6 0xFD 0x63 0xC8 0x49 0x80
Configuration last modified by 0.0.0.0 at 8-12-99 15:04:49
Switch#
```

BCMSN v2.1—2-59

Use the **show vtp counters** command to display statistics about VTP operation.

Example: Displaying VTP Statistics

This example shows how to display VTP statistics:

```
Switch# show vtp counters

VTP statistics:
Summary advertisements received    : 7
Subset advertisements received     : 5
Request advertisements received    : 0
Summary advertisements transmitted : 997
Subset advertisements transmitted  : 13
Request advertisements transmitted : 3
Number of config revision errors   : 0
Number of config digest errors     : 0
Number of V1 summary errors        : 0

VTP pruning statistics:

Trunk           Join Transmitted Join Received    Summary
advts received from
                                           non-pruning-capable
device
---------------- ---------------- ---------------- -----------
-------
Fa5/8               43071            42766            5
```

Selecting a Troubleshooting Approach for VTP

This topic identifies the guidelines used to troubleshoot specific VTP problems.

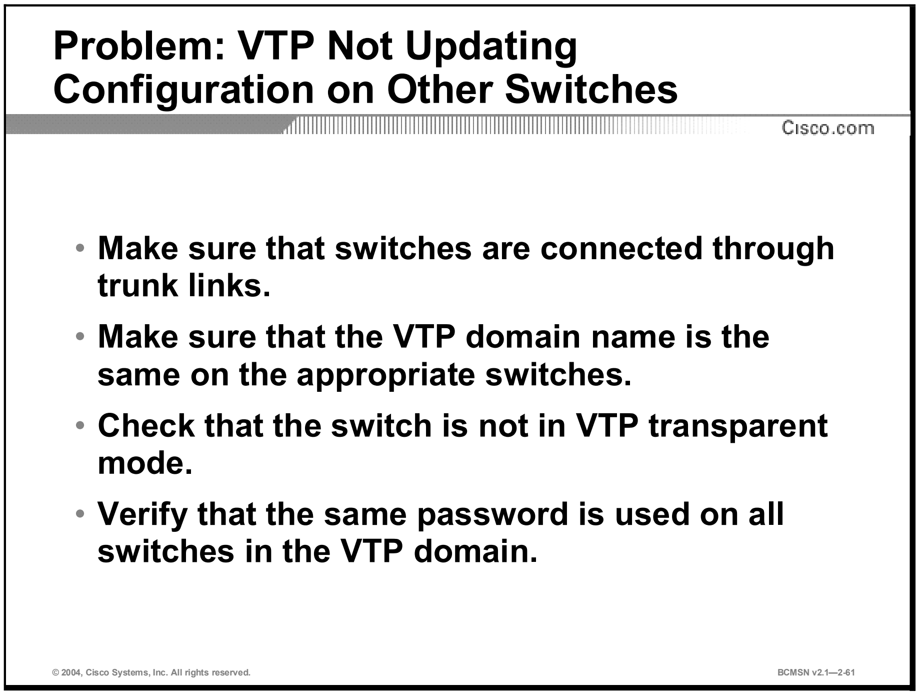

Problem: VTP Not Updating Configuration on Other Switches

Cisco.com

- Make sure that switches are connected through trunk links.
- Make sure that the VTP domain name is the same on the appropriate switches.
- Check that the switch is not in VTP transparent mode.
- Verify that the same password is used on all switches in the VTP domain.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—2-61

Problems with VTP configuration can frequently be traced to improperly configured trunk links, domain names, VTP modes, or passwords.

Problem: VTP is not updating the configuration on other switches when the VLAN configuration changes.

Suggested solutions to the problem are as follows:

- Make sure that the switches are connected through trunk links. VTP updates are exchanged only over trunk links. Check to make sure that the switches at each end of the link are using the same trunking protocol.
- Make sure that the VTP domain name (case sensitive) is the same on the appropriate switches. VTP updates are only exchanged between switches in the same VTP domain. Use the **show vtp status** command.
- Check if the switch is in VTP transparent mode. Only switches in VTP server or VTP client mode update their VLAN configuration based on VTP updates from other switches. Use the **show vtp status** command.
- If you are using VTP passwords, you must use the same password on all switches in the VTP domain. Make sure that the passwords on each switch match.

Note Make a backup copy of VLAN.dat or config.txt before troubleshooting.

Identifying How VTP Is Used in the Enterprise Composite Network Model

This topic identifies the benefits of applying VTP within an Enterprise Composite Network model.

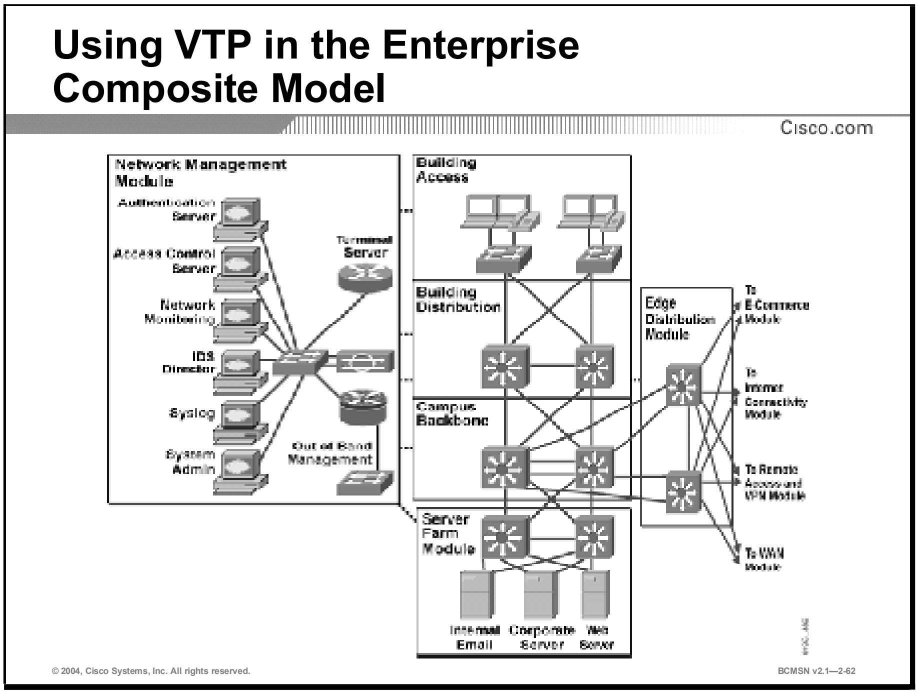

Using VTP in the Enterprise Composite Model

Cisco.com

BCMSN v2.1—2-62

Some of the benefits of applying VTP within an Enterprise Composite Network model are as follows:

- VTP domain is restricted to Building Access and Building Distribution switch blocks.
- VTP keeps VLAN information consistent between Building Distribution and Building Access switches.
- The failure domain is contained to the switch block.

Summary

This topic summarizes the key points discussed in this lesson.

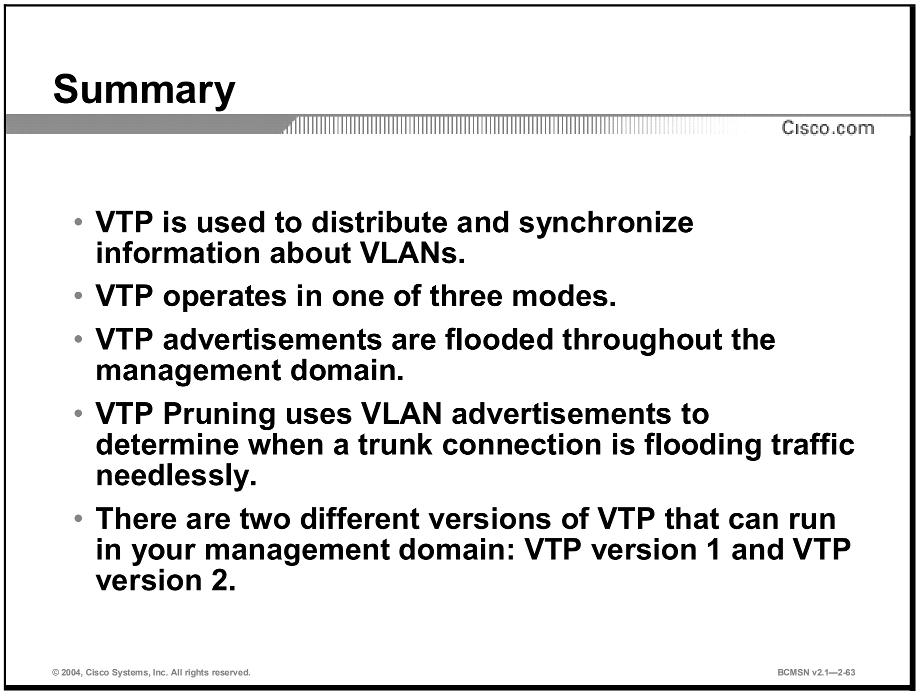

Summary

- VTP is used to distribute and synchronize information about VLANs.
- VTP operates in one of three modes.
- VTP advertisements are flooded throughout the management domain.
- VTP Pruning uses VLAN advertisements to determine when a trunk connection is flooding traffic needlessly.
- There are two different versions of VTP that can run in your management domain: VTP version 1 and VTP version 2.

BCMSN v2.1—2-63

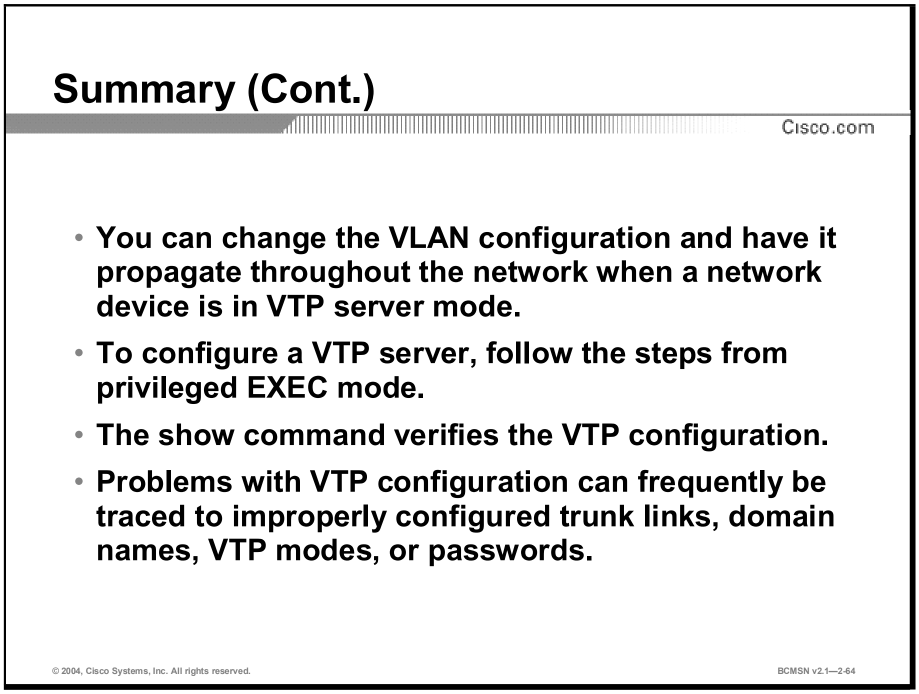

Summary (Cont.)

- You can change the VLAN configuration and have it propagate throughout the network when a network device is in VTP server mode.
- To configure a VTP server, follow the steps from privileged EXEC mode.
- The show command verifies the VTP configuration.
- Problems with VTP configuration can frequently be traced to improperly configured trunk links, domain names, VTP modes, or passwords.

BCMSN v2.1—2-64

References

For additional information, refer to this resource:

- The documentation that accompanied your Cisco Catalyst switch

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 2-1: Configuring VLANs and VTP

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Select the features that apply to VTP. (Choose two.)

A) distributes and synchronizes information about VLANs

B) allows for name changes in any type of domain

C) is a Layer 3 messaging protocol

D) maintains VLAN configuration consistency

Q2) Match the functions with the correct VTP mode.

_____ 1. client

_____ 2. server

_____ 3. transparent

A) cannot create, change, or delete VLANs

B) does not synchronize its VLAN configuration

C) creates, modifies, and deletes all VLANs

D) does not save VLAN configuration in NVRAM

Q3) List, in correct order, the steps on how VTP shares VLAN information in a management domain.

_____ 1. Step 1

_____ 2. Step 2

_____ 3. Step 3

_____ 4. Step 4

_____ 5. Step 5

A) VTP propagates Revision 4.

B) User adds new VLAN.

C) Revision 3 upgrades to Revision 4 on client.

D) Revision 3 upgrades to Revision 4 on server.

E) VTP synchronizes the new VLAN information on client.

Q4) Select the function used by VTP to control traffic within a management domain.

A) VTP ID

B) VTP advertisements

C) VTP Pruning

D) VTP version 1

Q5) Match the correct features with the appropriate VTP version.

_____ 1. VTP version 1

_____ 2. VTP version 2

A) domain name and version must match

B) consistency checks performed when you enter new information

C) default

D) version number propagated to other switches in the VTP domain

Q6) Select the guidelines used to configure a VTP management domain. (Choose two.)

A) add a new switch in client mode

B) set a different password for each switch in the domain

C) configure VTP pruning

D) change VLAN configuration in client mode and propagate it throughout the network

Q7) Match the correct configuration step with the appropriate VTP mode.

_____ 1. client

_____ 2. server

_____ 3. both

A) configure the domain name

B) enable VTP pruning in the management domain

C) exit global configuration mode

Q8) Select the different VTP configuration states displayed by the **show vtp status** command. (Choose two.)

A) VTP version

B) minimum VLANs supported

C) VTP password

D) number of existing VLANs

Q9) If the VTP is not updating the configuration on other switches, what steps would you take to solve the problem? (Choose two.)

A) Make sure that the VTP domain name is different on all switches.

B) Verify the switch is in VTP transparent mode.

C) Verify the switches are connected through trunk links.

D) Make sure that the same password is used on all switches in the VTP domain.

Q10) Select the benefit you can receive by applying VTP in an Enterprise Composite Network model.

A) VTP keeps VLAN information consistent between blocking switches.

B) VTP domain is restricted to the Building Access and Building Distribution switch blocks.

C) VTP assigns VID between Building Distribution and Building Access switches.

Quiz Answer Key

Q1) A, D

Relates to: Defining VTP

Q2) 1= A, D; 2=C; 3=B

Relates to: Identifying VTP Modes

Q3) 1=B, 2=D, 3=A, 4=C, 5=E

Relates to: Sharing VLAN Attributes in a Management Domain

Q4) C

Relates to: Controlling VLAN Traffic on a Trunk

Q5) 1=A, C; 2=B, D

Relates to: Distinguishing Between VTP Versions 1 and 2

Q6) A, C

Relates to: Identifying the Procedure for Configuring VTP

Q7) 2=A, B; 3=C

Relates to: Creating a VTP Management Domain

Q8) A, D

Relates to: Verifying the VTP Configuration

Q9) C, D

Relates to: Selecting a Troubleshooting Approach for VTP

Q10) B

Relates to: Identifying How VTP Is Used in the Enterprise Composite Network Model

Lesson Assessments

Overview

Use the lesson assessments here to test what you learned in this module. The correct answers and solutions are found in the Lesson Assessment Answer Key.

Outline

This section includes these assessments:

- Quiz 2-1: Implementing VLANs
- Quiz 2-2: Supporting Multiple VLANS Between Two Switches
- Quiz 2-3: Defining Trunking Protocols
- Quiz 2-4: Configuring Trunking Protocols
- Case Study 2-5: Maintaining VLAN Consistency Across the Network

Quiz 2-1: Implementing VLANs

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Select the characteristics that apply to VLANs
- Select the benefits that apply to VLANs
- Correctly identify the characteristics that apply to end-to-end and local VLANS
- Correctly select the conditions that no longer exist as a result of applying the Enterprise Composite Network model to a network design
- Identify the steps in the correct order necessary to create a generic Ethernet VLAN
- Correctly identify the steps necessary to create an Ethernet VLAN in global mode and the steps necessary to create an Ethernet VLAN in database mode
- Identify which command is used to verify a VLAN configuration
- Identify which commands are used to verify a VLAN port configuration
- Select the correct set of commands that delete a VLAN in both global and database mode
- Select the recommended approach to troubleshoot VLANs to resolve a specific problem

Quiz

Answer these questions:

Q1) Which two statements are characteristics of a VLAN? (Choose two.)

A) Membership is always associated with a switch port.

B) End users are members of the same broadcast domain.

C) It sets the switch port operational mode to be an access port.

D) Every data frame sent across a VLAN is either encapsulated or tagged.

Q2) A VLAN solves the scalability problems found in large flat networks by dividing the network into smaller _____.

A) switches

B) failure domains

C) broadcast storms

D) administrative domains

Q3) Which three statements are characteristics of an end-to-end VLAN? (Choose three.)

A) End-to-end VLANs tag frames sent between end users.

B) End-to-end VLANs are defined throughout the entire network.

C) End-to-end VLANs represent a bundle of multiple local VLANs.

D) An end-to-end VLAN may span several wiring closets or even several buildings.

E) End-to-end VLANs are usually associated with a workgroup, such as a department or project team.

Q4) When you apply the Enterprise Composite Network model to a network, what is one of the benefits you hope to see?

A) active redundant links at Layer 2

B) multiple customers with the same VID number

C) matching passwords on all switches

D) ports remain nonoperational until VLAN created

Q5) How is the VLAN database mode different than other configuration modes?

A) Catalyst switch determines the VID

B) changes not applied until you exit the session

C) must determine a name for the VLAN

D) can make VLAN configuration changes for upper range VLANs

Q6) What is the result of specifying a nonexistent VLAN number when assigning a port as a member of a VLAN?

A) An error message is displayed.

B) The VLAN is automatically created with default values.

C) The port remains nonoperational until the VLAN is created.

D) The port remains nonoperational until the VLAN is created and the port is reassigned.

Q7) Which command displays information about VLAN0003?

A) Switch#**show interface VLAN name VLAN0003**

B) Switch#**show no VLAN name VLAN0003**

C) Switch#**show VLAN name VLAN0003**

D) Switch#interface **show VLAN name VLAN0003**

Q8) When troubleshooting a VLAN, you should check the physical connections, VLAN configuration, and _____.

A) port configuration

B) VLAN membership

C) switch configuration

D) physical subnet

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Quiz 2-2: Supporting Multiple VLANs Between Two Switches

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Select the problems that occur when supporting multiple VLANs
- Identify the benefits that trunks provide in the Enterprise Composite Network model

Quiz

Q1) You are a new employee at a small company. You are one of two engineers who will be maintaining the company's network. You have learned that the network uses older switches and is not very complex. The network does not have native VLANs, or support for Voice over IP (VoIP) or multicasting. Because you do not have a lot of experience as a network engineer, you are pleased to see that their network is fairly easy to configure. Given this situation, what type of trunking protocol would you recommend for this network?

A) ISL

B) 802.1Q

Q2) You are a technical engineer at a large corporation. You are involved in a major expansion of the company's network. The company is adding support for VoIP and multilayer switches. The company also will incorporate FDDI backbones. Given this situation, what type of trunking protocol would you recommend for this network?

A) ISL

B) 802.1Q

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 100 percent.

Quiz 2-3: Defining Trunking Protocols

Complete this quiz to assess what you learned in the lesson.

Objectives

- Referencing Diagrams A and B, match the correct protocol with the appropriate diagram
- Referencing Diagram C, identify the place where a trunking protocol should be configured

Quiz

Q1) Using Diagram A and Diagram B, match the correct trunking protocol with the appropriate diagram.

_____ 1. ISL

_____ 2. 802.1Q

_____ 3. 802.3

_____ 4. VLANs

_____ 5. 802.5Q

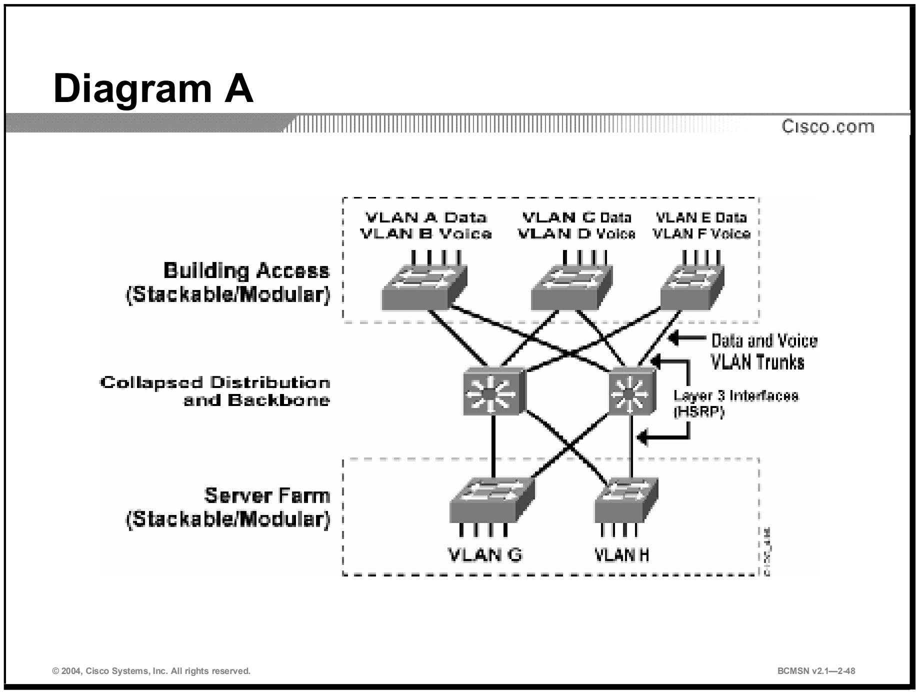

Diagram A

Cisco.com

VLAN A Data
VLAN B Voice
VLAN C Data
VLAN D Voice
VLAN E Data
VLAN F Voice
Building Access (Stackable/Modular)
Data and Voice VLAN Trunks
Collapsed Distribution and Backbone
Layer 3 Interfaces (HSRP)
Server Farm (Stackable/Modular)
VLAN G
VLAN H

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-48

Diagram A

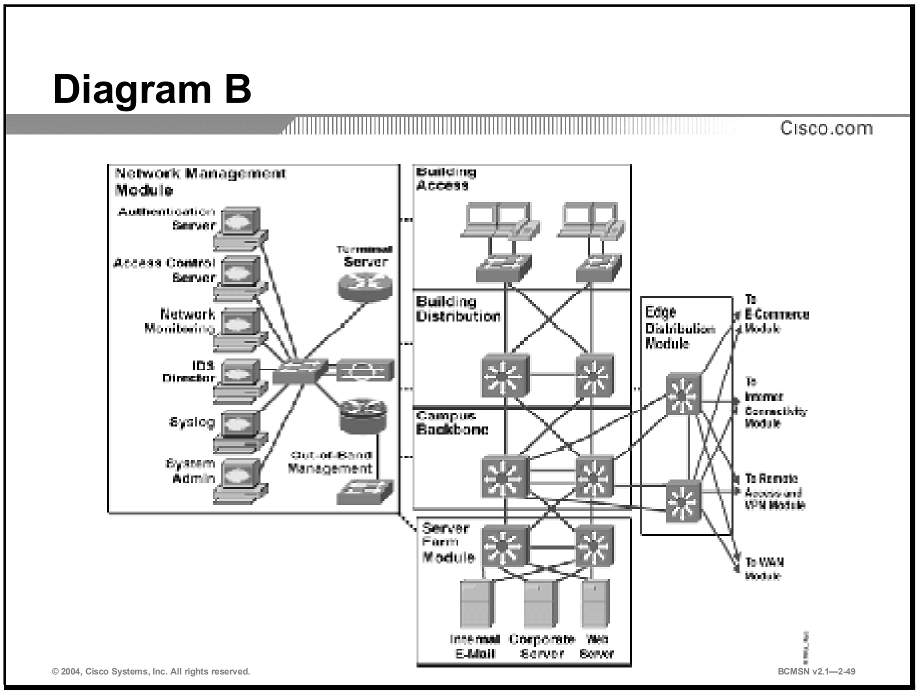

Diagram B

Cisco.com

BCMSN v2.1—2-49

Diagram B

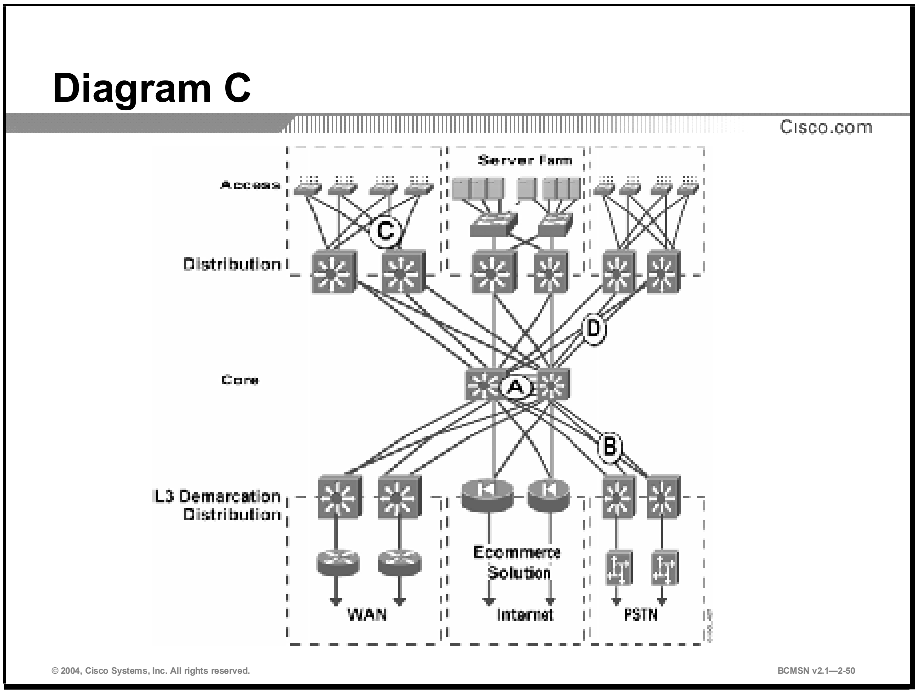

Diagram C

Cisco.com

BCMSN v2.1—2-50

Diagram C

Q2) Where on Diagram C would you apply your trunking protocol?

A) A

B) B

C) C

D) D

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 100 percent.

Quiz 2-4: Configuring Trunking Protocols

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Identify the steps and the correct order that apply to configuring both ISL or 801.2Q trunking protocols
- Correctly identify the commands used to configure an ISL trunk
- Correctly identify the commands used to verify the configuration of an ISL trunk
- Correctly identify the commands used to configure an 802.1Q trunk
- Correctly identify the commands used to verify 802.1Q trunk configuration
- Identify and resolve a specific trunk link problem in the network

Quiz

Answer these questions:

Q1) When configuring an 802.1Q trunking port, which command configures the allowable VLANs?

A) switch(config)#**switchport trunk encapsulation {isl | dot1q | negotiate}**

B) switch(config)#**interface {fastethernet | gigabitethernet}**

C) switch(config-if)#**switchport trunk allowed vlan {add | except | all | remove}**

D) switch(config-if)#**switchport mode {dynamic {auto | desirable} | trunk}**

Q2) What is the default native VLAN when configuring an 802.1Q trunk?

A) 10

B) 100

C) 1

D) 11

Q3) What is the command to display trunk information for 802.1Q trunking?

A) switch#**show running-config interface fastethernet 5/8**

B) switch#**show interfaces fastethernet 5/8 trunk**

C) switch#**configure terminal**

D) switch(config)#**interface {fastethernet | gigabitethernet}**

Q4) A device cannot establish a connection across a trunk link. What are two possible solutions?

A) The trunk ports should both be in trunk port operational mode.

B) The VLAN is automatically created with default values.

C) The native VLAN should be the same on both ends of the trunk.

D) The trunk encapsulation type configured on both ends of the link should be the same.

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Case Study 2-5: Maintaining VLAN Consistency Across the Network

Complete this case study to assess what you learned in the lesson.

Overview

VTP is used to distribute and synchronize information about VLANs configured throughout a switched network. VTP reduces the manual configuration needed in the network.

Relevance

In a large switched network, VTP allows you to manage the VLAN implementation.

Objectives

In this activity, learners will identify the process and commands used to propagate VLAN information throughout a VTP management domain. Upon completing this case study, learners will be able to meet these objectives:

- Referencing Diagram D and the job scenario, identify the command and operation mode used when configuring VLANs

Learner Skills and Knowledge

To benefit fully from this activity, you must have these prerequisite skills and knowledge:

- Basic knowledge of the components that make up the Enterprise Composite Network model
- Basic knowledge of switch operations
- Basic knowledge of router operations
- Basic knowledge of how data is routed between devices

Job Aid

The following diagram is used in this case study.

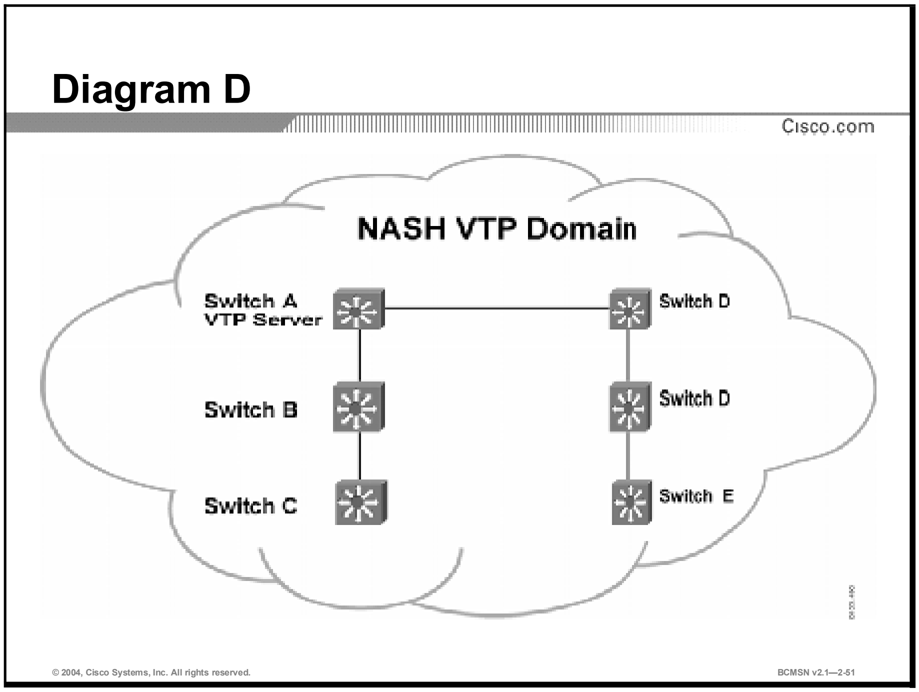

Diagram D

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—2-51

Diagram D

Your customer is adding new switches and a second VLAN on the NASH VTP Domain diagrammed below. Switch A is the VTP server for the domain. Switches A, B, and C are in VLAN4. The new Switches D, E and F are in VLAN5.

Outline

Answer the following questions regarding the scenario and diagram above.

Q1) What commands are needed to implement VLAN5 in the VTP domain and in which switches should the commands be interned?

____

Q2) What mode should the new switches be in before adding them to the VTP domain?

____

Q3) What commands are needed on switches B and C with VLAN5?

____

Case Study Verification

You have completed this activity when your case study solution has been presented to the class and you have justified any major deviations from the case study solution supplied by the instructor.

Presentation of Case Study

You have completed this activity when your case study solution has been presented to the class and you have justified any major deviations from the case study solution supplied by the instructor.

Lesson Assessment Answer Key

Quiz 2-1: Implementing VLANs

Q1) A, B

Q2) C

Q3) B, D, E,

Q4) A

Q5) B

Q6) C

Q7) C

Q8) A

Quiz 2-2: Supporting Multiple VLANs Between Two Switches

Q1) A

Q2) B

Quiz 2-3: Defining Trunking Protocols

Q1) 1=B, 2=A, 3=A, 4=A, 5=B

Q2) C

Quiz 2-4: Configuring Trunking Protocols

Q1) D

Q2) C

Q3) B

Q4) A, D

Case Study: 2-5: Maintaining VLAN Consistency Across the Network

Q1) Switch A will have to be configured for VLAN5, and the interface will have to be configured and activated. Switches D, E, and F should be configured with a VTP domain name of "NASH" and a VTP mode of "Client." All other commands are not needed to get VTP running.

Q2) All new switches should be added in client mode.

Q3) No commands are needed. VTP will update all switches in the NASH VTP Domain.

Module 3

Implementing Spanning Tree Protocol

Overview

Because organizations rely heavily on their multilayer switched network for conducting business, high availability is a primary concern. One method of ensuring high availability is to provide redundancy of devices, modules, and links throughout the network. Network redundancy introduces the potential for bridging loops—packets looping endlessly between devices, crippling the network. The Spanning Tree Protocol (STP) is designed to identify and prevent such loops.

Upon completing this module, you will be able to:

- Identify the bridging-loop issues and solutions
- Identify the features that describe the operations of STP
- Identify the steps STP takes to establish a loop free topology in a new network
- Configure and verify STP on a switch device

Outline

The module contains these components:

- Preventing Bridging Loops Using Spanning Tree Protocol
- Defining Spanning Tree Protocol Operations
- Establishing a Loop-Free Topology in a New Network
- Configuring Spanning Tree Protocol
- Lesson Assessments

Preventing Bridging Loops Using Spanning Tree Protocol

Overview

Spanning Tree Protocol (STP) is a Layer 2 link management protocol that provides path redundancy while preventing undesirable loops in switched or bridged networks. The STP operation is transparent to end stations. The STP runs on Layer 2 switches, bridges, and routers configured to operate as bridges.

Relevance

Redundant switched and bridged topologies cause problems in the network. It is important to know why these problems occur and how STP addresses them.

Objectives

Upon completing this lesson, you will be able to:

- Identify the features that apply to transparent bridging
- Identify the behavior of flooded unicast frames in a bridged loop
- Identify the behavior of broadcast frames in a bridged loop
- Select a definition that best describes how STP uses spanning tree algorithm to prevent bridging loops in a redundant network

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Identifying the Features of Transparent Bridging
- Identifying Flooded Unicast Traffic
- Identifying Broadcast Traffic
- Preventing Bridging Loops in a Redundant Network
- Summary
- Quiz

Identifying the Features of Transparent Bridging

This topic identifies the features that apply to transparent bridging.

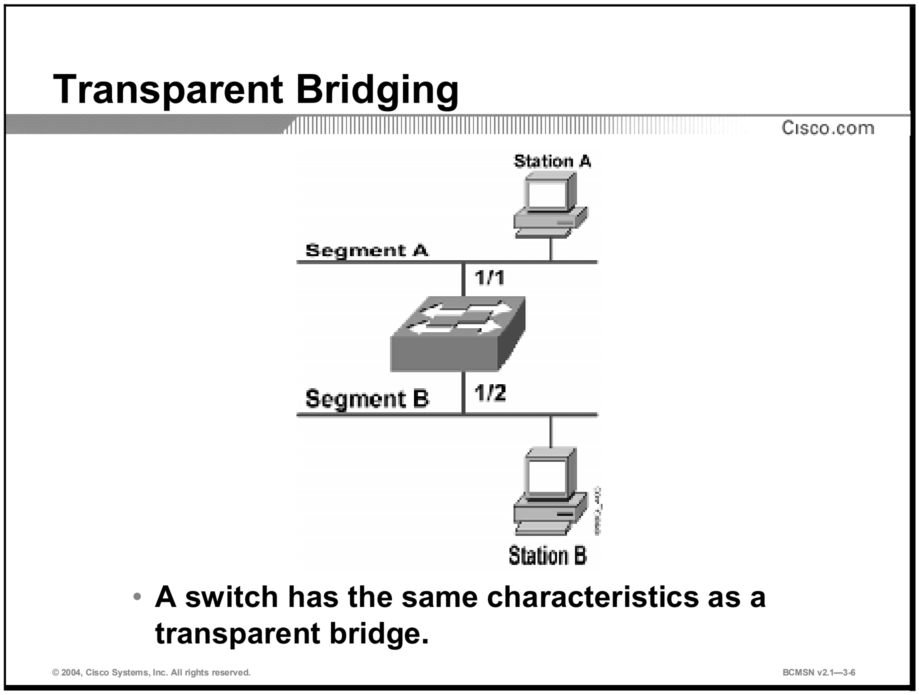

Transparent Bridging

Cisco.com

- **A switch has the same characteristics as a transparent bridge.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—3-6

The basic STP functionality of a switch is identical to that of a transparent bridge. To understand STP, it is important first to look at the behavior of a transparent bridge without spanning tree.

By definition, a transparent bridge has these characteristics:

- Must not modify the frames that are forwarded.
- Learns addresses by "listening" on a port for the source address of a device. If a source address comes into a specific port, the bridge assumes that the source address can be found by sending out from that port. The bridge then builds a table indicating that frames can reach the source by sending the frames out that same port. A bridge is always listening and learning.
- Must forward all broadcasts out all ports, except for the port that initially received the broadcast.
- Forwards the frame out all ports—except for the port that initially received the frame— if a destination address is unknown (sometimes called an "unknown unicast").

Transparent bridging by its very definition must be transparent to the devices on the network. End stations do not need to be modified to support the process of bridging.

In a simple bridge environment without any redundant links, transparent bridging works. However, transparent bridging begins to have problems as soon as a redundant path is added to the bridged network.

Identifying Flooded Unicast Traffic

This topic identifies the behavior of flooded unicast frames in a bridged loop.

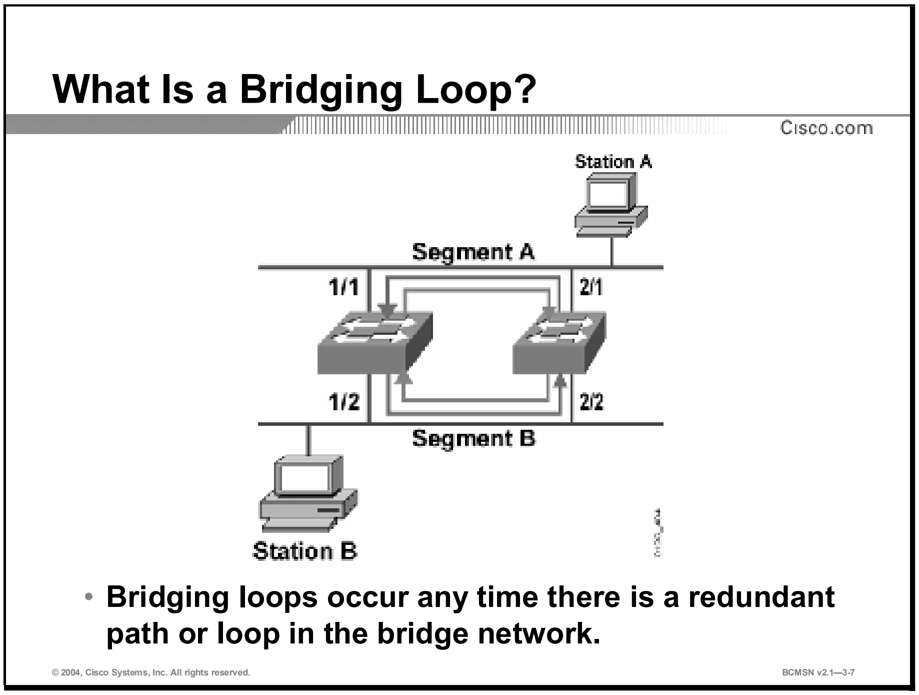

What Is a Bridging Loop?

Cisco.com

- **Bridging loops occur any time there is a redundant path or loop in the bridge network.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—3-7

Example: Flooded Unicast Frames and Bridging Loops

Station A has two potential paths to station B by way of the switches. What happens if station A sends to station B, but neither of the switches has station B in its address table?

- Station A transmits the frame to segment A. Both bridges on segment A pick up the frame on their ports 1/1 and 2/1, respectively. Both switches populate their respective address tables indicating that station A resides on segment A on ports 1/1 and 2/1.
- Both switches forward the frame to segment B. Notice that not only will station B receive the frame, but that both switches also see the frame coming from the *other* switch. Because one of the basic characteristics of transparent bridging is to "listen" to the source addresses to "learn" the correct port to use for that address, each switch relearns station A as residing on ports 1/2 and 2/2. The switches now incorrectly assume that all frames for station A should be sent to segment B.
- The packet is then forwarded again to segment A, where the frame originated. Because neither switch is aware of the other, each switch will continually forward the frame on the other port. This loop will go on forever. Sometimes, the information that the switch has learned is correct and the frame will make it to the destination. Other times, the switch will believe that the destination is on the same segment as the receiving port and will not forward the frame.
- If station A had originally sent a broadcast, the problem would actually be much worse than a bridging loop. Because bridges always retransmit a broadcast and never mark the frame, bridges actually can create broadcasts in an exponential fashion when a bridge loop occurs. This process of creating new broadcasts does not stop until the loop is shut down. Eventually, the bridging loop brings down the network through a broadcast storm.

Identifying Broadcast Traffic

This topic identifies the behavior of broadcast frames in a bridged loop.

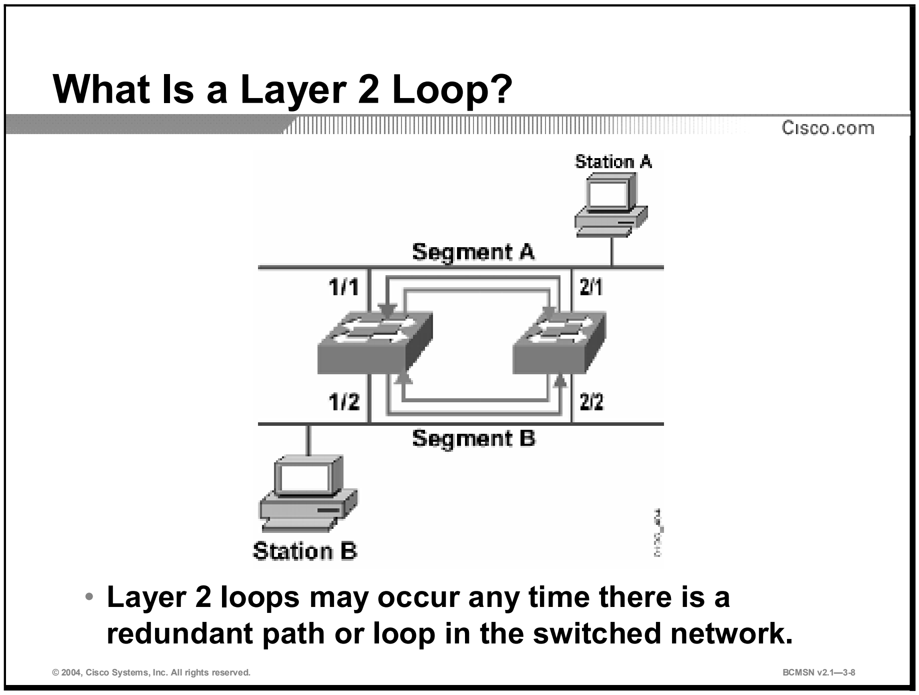

What Is a Layer 2 Loop?

- **Layer 2 loops may occur any time there is a redundant path or loop in the switched network.**

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—3-8

Example: Broadcast Frames and Bridging Loops

Station A has two potential paths to station B by way of the two intermediate switches. What happens if station A sends to station B, while a Layer 2 loop exists without STP?

- Station A transmits the frame destined for station B to segment A. Both bridges on segment A pick up the frame on their switch ports 1/1 and 2/1, respectively. Both switches populate their respective MAC tables indicating that station A resides on segment A on switch ports 1/1 and 2/1.
- Both switches forward the frame to segment B. Notice that not only will station B receive the frame, but that both switches also see the same frame, with station A's MAC address in the source address (SA) field, coming from the *other* switch. The switches will now incorrectly forward all frames for station A to segment B. When station B responds to station A, the frame will be dropped by both switches because it will be received on the same switch ports that are also the destination switch ports, according to the MAC table of each switch.
- If station A transmits, both switches will "relearn" that station A resides on segment A. Any frames destined for station A are then forwarded again to segment A. Station B will be able to communicate with station A until station A responds, at which point, both switches will again relearn that station A is on segment B. This transaction will effectively cause the loss of connectivity to station A. The same scenario would happen to all stations on segments A and B. The network experiences the effects of a Layer 2 loop. The loop manifests itself as the ability to get to station A and station B some of the time.

- If station A, or any station, sends a broadcast, the effects of the Layer 2 loop would be much worse. The destination MAC address would be FF-FF-FF-FF-FF-FF. This would cause each switch to forward the frame out all switch ports except the switch port upon which the frame was received. The broadcast frame would also be forwarded to the originating switch, which would again forward the same broadcast out all switch ports. This broadcast would continue until the loop is shut down or the switch could no longer handle the load.

Preventing Bridging Loops in a Redundant Network

This topic identifies how STP uses the spanning tree algorithm (STA) to prevent bridging loops in a redundant network.

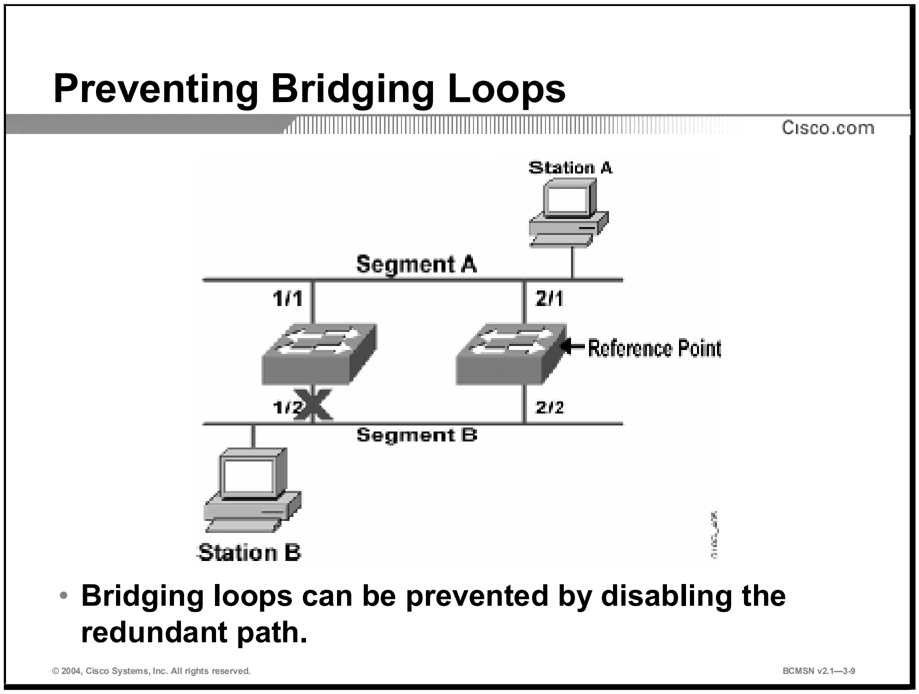

Preventing Bridging Loops

Cisco.com

- **Bridging loops can be prevented by disabling the redundant path.**

BCMSN v2.1—3-9

STP was created to overcome the problems of transparent bridging in redundant networks. The purpose of STP is to avoid and eliminate loops in the network by imposing a loop-free path. STP does this by determining where there are loops in the network and shutting down links that are redundant. In this way, STP ensures that there will be only one path to every destination and that a bridging loop can never occur. In the case of a link failure, the bridge would know that a redundant link exists and would bring up the link that was previously shut down.

This technology means that some ports will need to be disabled or put into nonforwarding (blocking) mode. The ports remain aware of the topology of the network and can be enabled if a failure occurs on the link-forwarding data.

The STP executes an STA. To find the redundant links, the STA chooses a reference point in the network and calculates the redundant paths to that reference point. If the STA finds a redundant path, STA chooses which path will forward frames and which redundant paths are blocked. This technique effectively severs the redundant links within the network.

Summary

This topic summarizes the key points discussed in this lesson.

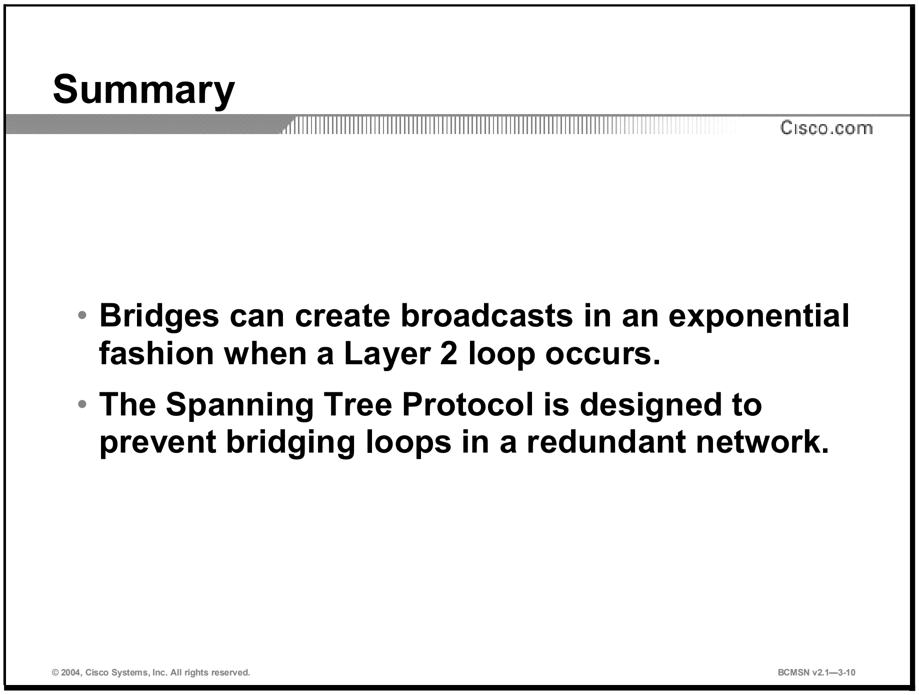

Summary

Cisco.com

- Bridges can create broadcasts in an exponential fashion when a Layer 2 loop occurs.
- The Spanning Tree Protocol is designed to prevent bridging loops in a redundant network.

BCMSN v2.1—3-10

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Select the features that apply to transparent bridging. (Choose two.)

A) End stations must be modified to support the bridging process.

B) Frames are often modified.

C) Transparent bridging forwards all broadcasts out all ports except for the port that initially received the broadcast.

D) Transparent bridging functions the same as the STP functionality of a switch.

Q2) Given a Layer 2 network topology in which switches have active redundant paths, select the behavior of the flooded unicast frames in a bridged loop.

A) The frame is forwarded to both its intended destination and its origination.

B) The frame is forwarded to its intended destination.

C) The frame is returned to its origination.

D) The frame goes back and forth between two switches, never reaching its destination.

Q3) Given a Layer 2 network topology in which switches have active redundant paths, select the behavior of the broadcast frames in a bridged loop.

A) The broadcast frames are all returned to their origination.

B) The broadcast frames get to their destination some of the time.

C) The broadcast frames all flow to their destination.

Q4) Select the best description of how spanning tree uses the spanning tree algorithm to prevent bridging loops in a redundant network. (Choose two.)

A) Some ports are disabled.

B) Some ports send out a broadcast message.

C) Alternate switches are introduced.

D) Some redundant paths are blocked.

Quiz Answer Key

Q1) C, D

Relates to: Identifying the Features of Transparent Bridging

Q2) A

Relates to: Identifying Flooded Unicast Traffic

Q3) B

Relates to: Identifying Broadcast Traffic

Q4) A, D

Relates to: Preventing Bridging Loops in a Redundant Network

Defining Spanning Tree Protocol Operations

Overview

STP is configured on a per-VLAN basis, designating root and secondary root bridges, and influencing the likelihood of ports being selected for the forwarding state.

Relevance

Spanning tree is a critical feature of a multilayer switched network. The configuration of spanning tree can have a huge impact on network performance. Proper configuration methods help ensure that the network operates at full efficiency.

Objectives

Upon completing this lesson, you will be able to:

- Select the bridge ID that is most likely to be chosen as the root bridge by STP
- Match each bridge ID with the correct description
- Select the correct path based on path cost to determine the forwarding path between a device and the root switch
- Match the features the correctly apply to the appropriate port role
- Select the features that best apply to BPDU messaging
- Match the correct BPDU field name with the appropriate definition
- Identify the correct timer that affects the transition of the switch port from one STP state to another
- Match the correct features and function to the appropriate STP port state
- List the steps involved in topology changes in the correct order

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Selecting the Root Switch
- Distinguishing Between Bridge IDs
- Determining the Forwarding Path
- Defining Spanning Tree Port Roles
- Using BPDU Messaging
- Identifying Selected BPDU Field Definitions
- Transitioning from One STP State to Another
- Identifying Spanning Tree Port States
- Identifying Topology Changes
- Summary
- Quiz

Selecting the Root Switch

This topic identifies the BID that is most likely to be chosen as the root switch by STP.

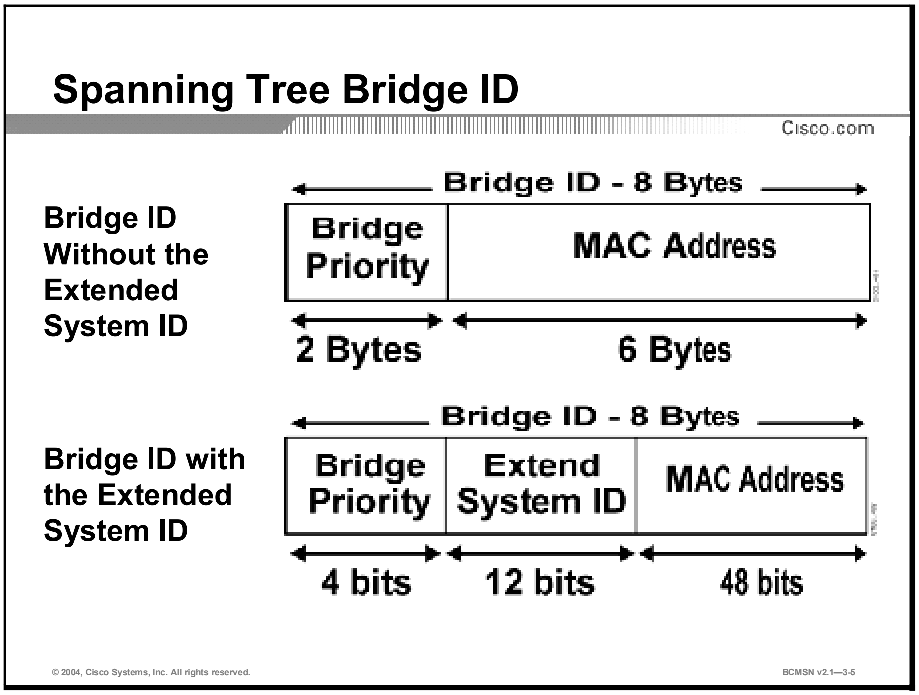

Spanning Tree Bridge ID

Cisco.com

Bridge ID Without the Extended System ID

Bridge ID - 8 Bytes

Bridge Priority	MAC Address
2 Bytes	6 Bytes

Bridge ID with the Extended System ID

Bridge ID - 8 Bytes

Bridge Priority	Extend System ID	MAC Address
4 bits	12 bits	48 bits

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—3-5

STP uses the concepts of root switches, root ports, and designated ports to establish a loop-free path through the network. An Ethernet Layer 2 loop-free environment is accomplished by blocking various links to provide a single Layer 2 path through which to forward frames. There are various factors and processes that STP leverages to determine what path should be blocked and what path should forward.

Spanning tree calls for each bridge entity to be assigned a unique identifier called the "bridge ID" (BID). The BID, which is a 64-bit (8-byte) value, is used in determining the Layer 2 network reference point. This reference point is called a "root bridge" or "switch." STP calculates which redundant path will be blocked and which will be used to forward data. The switch with the lowest numerical BID is considered to be the most favorable reference point or "root switch."

STP can calculate which paths are blocked and which paths may forward on a per-STP process basis, called an "instance." STP can also calculate blocked paths on a per-VLAN basis, where each instance or VLAN requires a unique BID. The BID can be local at the switch itself, or on other interconnected switches and their STP instances. It is more common than not to have multiple STP processes running; therefore, a unique root switch and BID per instance or VLAN is vital. The BID comprises a bridge priority value (4 bits), an extended system ID value (12 bits), and a bridge MAC address (48 bits or 6 bytes).

Distinguishing Between Bridge IDs

This topic matches each BID with the correct description.

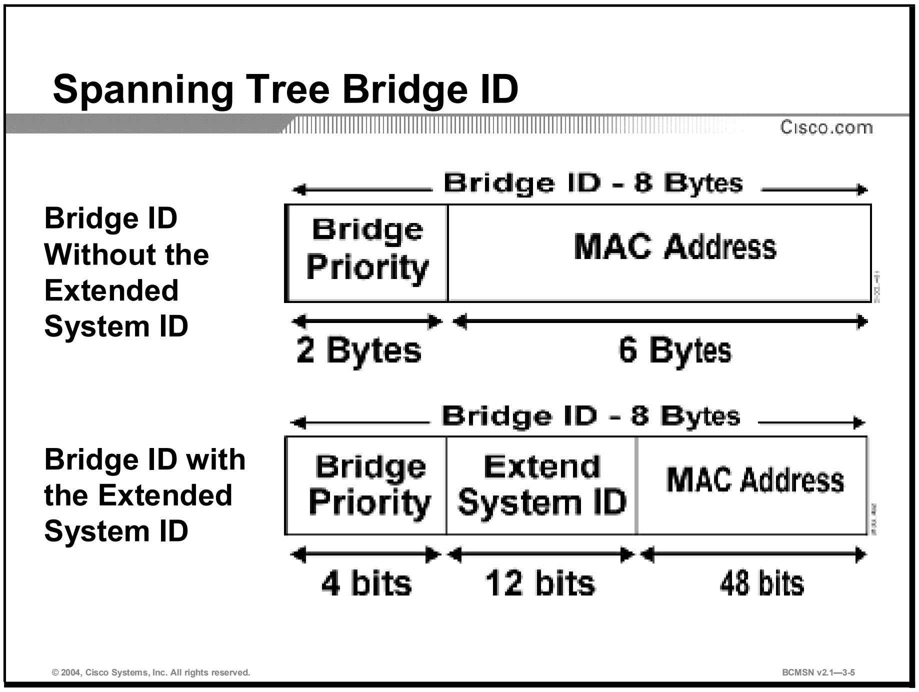

Spanning Tree Bridge ID

Cisco.com

Bridge ID Without the Extended System ID

Bridge ID - 8 Bytes

Bridge Priority	MAC Address
2 Bytes	6 Bytes

Bridge ID with the Extended System ID

Bridge ID - 8 Bytes

Bridge Priority	Extend System ID	MAC Address
4 bits	12 bits	48 bits

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—3-5

The BID was made up of only a bridge priority value (two bytes) and a bridge MAC address (six bytes). The BID was always unique by virtue of using a unique MAC address for each STP instance or VLAN. The MAC addresses were allocated from a pool of MAC addresses that are factory assigned to the switch or module. Depending on the platform, there was usually either a pool of 64 or 1024 MAC addresses available for use in making the BID unique for STP operations. The default priority was 32,768 (1000 0000 0000 0000 in binary, or 0x8000 in hex). This is a midrange value of the maximum total value of 65,535 (1111 1111 1111 1111 in binary, or 0xFFFF in hex) available in the 2-byte priority field. Many Catalyst switch platforms still support this method of comprising the BID. The drawback would be that for each STP instance or VLAN that is created, a single unique MAC address must be allocated for use in the BID for STP. For switch platforms that only use a pool of 64 MAC addresses, there would only be the ability to support 64 VLANs in the range of 1 to 1024. Many Catalyst switches that only support 64 addresses will default to, the extended system ID feature to be enabled.

The extended system ID feature redefines the bridge priority bits to include only the most significant 4 bits of a BID. The remaining 12 bits have been reallocated as a system ID and allow for 4096 BIDs to be uniquely identified. The purpose of this change was to no longer rely on MAC addresses to make the BID locally unique and to allow only one base MAC address to be spent in the BID for a single switch. The extended system ID value is the VLAN ID (VID). For example, an STP instance running on VLAN 1000 would have an extended system ID value of 1000. The 802.1Q trunking protocol also uses 12 bits in its VID field in the 802.1Q tag to uniquely identify 4096 VLANs. Therefore, there will always be a unique BID for any STP instance or VLAN created without using more that one base MAC address. This is why the extended system ID feature is the default.

Determining the Forwarding Path

This topic identifies the forwarding path between a device and the root bridge.

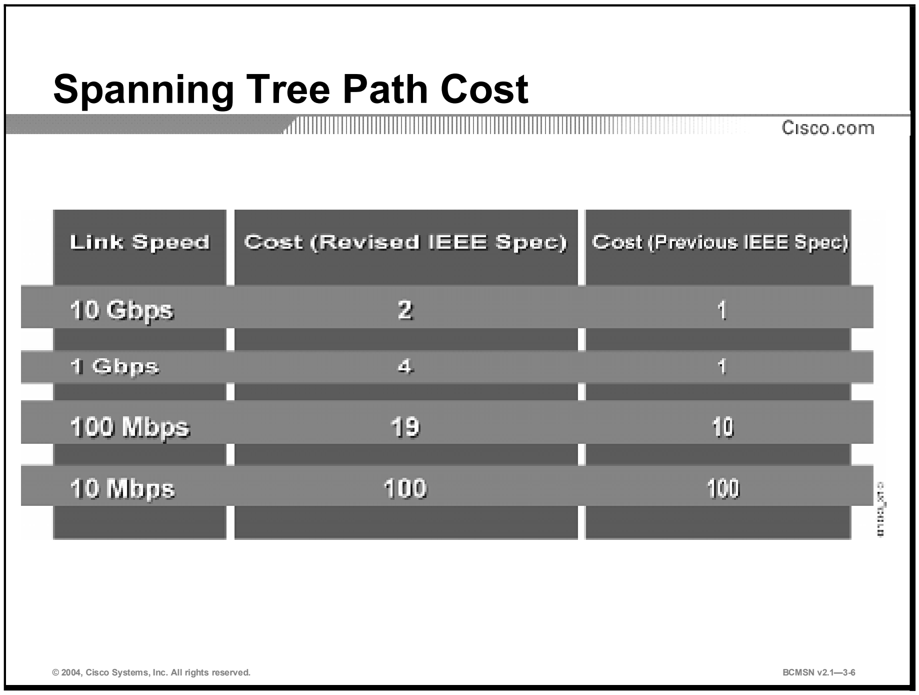

Spanning Tree Path Cost

Cisco.com

Link Speed	Cost (Revised IEEE Spec)	Cost (Previous IEEE Spec)
10 Gbps	2	1
1 Gbps	4	1
100 Mbps	19	10
10 Mbps	100	100

BCMSN v2.1—3-6

One factor in determining which path to forward frames is the path cost. The spanning tree path cost is an accumulated total cost of the path from an intermediate switch to the reference point or root switch of the Layer 2 network. The path cost is based on the bandwidth of all intermediate links in the path to the root switch. A path cost is used to determine the best path to the root switch. The lowest cost is considered to be the best path.

Example: Determining the Forward Path

There are two redundant paths that exist to the root switch: one with a path cost of 38 and the other with a path cost of 23. The path cost of 38 indicates that there are two Fast Ethernet links to cross to get to the root switch. The path cost of 23 indicates that there is one Fast Ethernet link and one Gigabit Ethernet link to cross to reach the root switch. The path with the cost of 23 would be used because it has the lower numerical path cost value. A root switch will have a path cost of “zero” because it is the root switch and there is no path.

The IEEE 802.1D specification has been revised. In the older specification, the cost value was calculated based on 1000 Mbps, or Gigabit Ethernet being the maximum Ethernet bandwidth available. The new specification adjusts the calculation using a nonlinear scale to accommodate higher-speed interfaces. Once a bridge signals a topology change, it starts sending topology change notifications (TCNs) on its root port. The designated bridge receives the TCN, acknowledges it, and generates another one for its own root port—and so on—until the TCN hits the root bridge.

Defining Spanning Tree Port Roles

This topic identifies the features that apply to the appropriate port role.

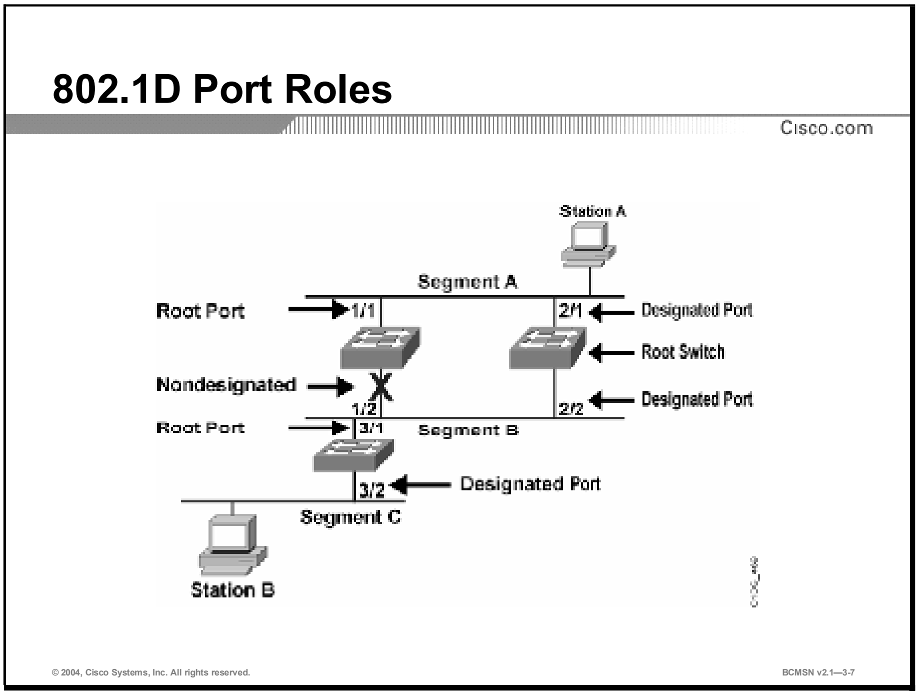

802.1D Port Roles

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—3-7

When STP has determined a forwarding path, the switch ports will have assumed various roles that define their specific function and operation. There are four 802.1D port roles.

- **Root port:** This port exists on nonroot or designated switches only and is the switch port with the least path cost to the root switch. Root ports are responsible for forwarding frames to and from an intermediate segment facing toward the root switch from the perspective of the designated switch. Root ports are also capable of populating the MAC table. Only one root port is allowed per switch.
- **Designated port:** This port exists on root and nonroot switches. For root switches, all switch ports are designated ports. For nonroot or designated switches, the designated port is the switch port with the least path cost to the root switch that is not a root port. The designated port forwards frames toward the root switch from the perspective of that switch. Designated ports are also capable of populating the MAC table. Only one designated port per segment is allowed. If multiple switches exist on the same segment, an election process determines the designated switch, and the corresponding switch port begins forwarding frames for the segment.
- **Nondesignated port:** This is the switch port that is blocking data frames and is not populating the MAC table on any given segment.
- **Disabled port:** This is a port that is shut down.

By examining the switch port roles on a switch, the forwarding path can be determined to ensure that data frames are taking the desired path.

Using BPDU Messaging

This topic identifies the features that apply to bridge protocol data unit (BPDU) messaging.

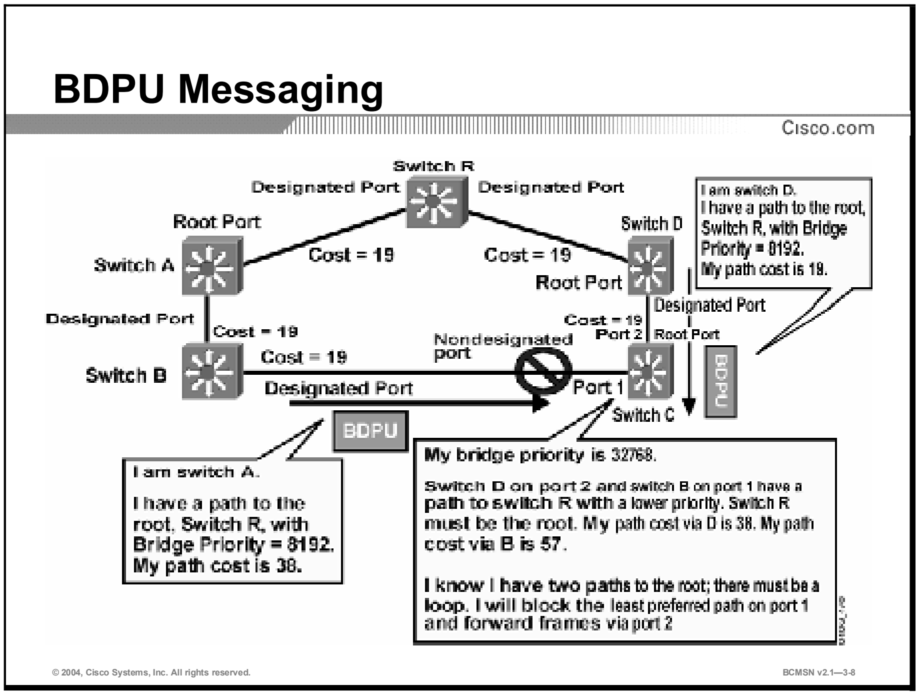

BDPU Messaging

Cisco.com

I am switch D.
I have a path to the root, Switch R, with Bridge Priority = 8192.
My path cost is 19.

I am switch A.

I have a path to the root, Switch R, with Bridge Priority = 8192.
My path cost is 38.

My bridge priority is 32768.

Switch D on port 2 and switch B on port 1 have a path to switch R with a lower priority. Switch R must be the root. My path cost via D is 38. My path cost via B is 57.

I know I have two paths to the root; there must be a loop. I will block the least preferred path on port 1 and forward frames via port 2

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—3-8

Switches running STP exchange configuration messages with other switches at regular 2-second intervals. The 2-second intervals are the default to help ensure a stable, loop-free, Layer 2 topology.

Overall, the exchange of BPDUs yields the following final results:

- The election of a root switch as a Layer 2 topology point of reference
- The determination of the best lowest cost path to the root switch
- The election of a designated switch and corresponding designated port for every switched segment
- The removal of loops in the switched network by transitioning redundant switch ports to a disabled port role
- Determination of the “active topology” for each instance or VLAN running STP

The active topology is the final set of communication paths that are created by switch ports forwarding frames. Once the active topology has been established, the switched network must reconfigure the activity topology if a link failure occurs. The switch uses a special BPDU called a topology change notification (TCN). When a topology change occurs, a bridge sends TCNs out on the root port. The designated bridge receives the TCN, acknowledges it, and generates another TCN for its own root port. This process continues until the root switch receives the TCN.

A TCN is sent when these events occur:

- There is a link failure.
- A switch port begins forwarding data frames, and the bridge already has a designated port.
- A nonroot bridge receives a TCN on its designated port, then propagates the TCN or retransmits the TCN out its root port toward the root switch.

Once the root bridge registers the topology change, the root switch generates configuration BPDUs with the topology change (TC) bit set. Every switch in the network relays these BPDUs to propagate the change throughout the network.

Example: Learning Topology Changes Using BPDU Messaging

In this example, two switches have a path to the root switch and have identified the BID of the root that contains the bridge priority value of 8192. (Priorities are covered in the module “Enhancing Spanning Tree Protocol.) Switch C is the receiving switch and only receives the two BPDUs. Switch C can deduce several characteristics regarding the STP topology. The bridge priority is a lower numerical value and therefore, is preferred to be the root. The path cost to the root is a lower numerical value and therefore better. There is also a Layer 2 loop in the STP topology; thus, switch C will disable port blocking and data frames from being transmitted out port 1 and dropping any data frames received. However, switch B will still send BPDUs, and switch C will continue to evaluate them.

Identifying Selected BPDU Field Definitions

This topic identifies the correct BPDU field name with its appropriate definition.

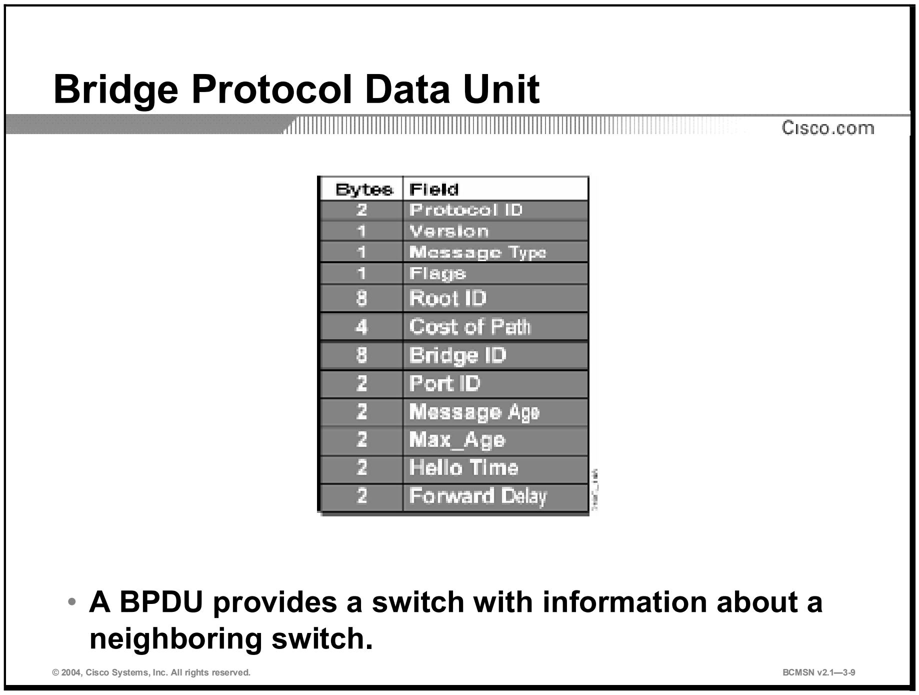

Bridge Protocol Data Unit

Bytes	Field
2	Protocol ID
1	Version
1	Message Type
1	Flags
8	Root ID
4	Cost of Path
8	Bridge ID
2	Port ID
2	Message Age
2	Max_Age
2	Hello Time
2	Forward Delay

- **A BPDU provides a switch with information about a neighboring switch.**

BCMSN v2.1—3-9

The information provided in a BPDU includes the following:

- The BID of the transmitting switch
  - The switch priority (located in the BID)
  - The extended system ID
  - The MAC address of the transmitting switch
- The transmitting switch port ID
- The root ID
- The path cost from the transmitting switch to the calculated root switch
- The STP timer parameter values

The switch compares the BPDUs and evaluates the roles that each port should play in the STP topology.

Transitioning from One STP State to Another

This topic identifies the timer that affects the transition of a switch port from one STP state to another.

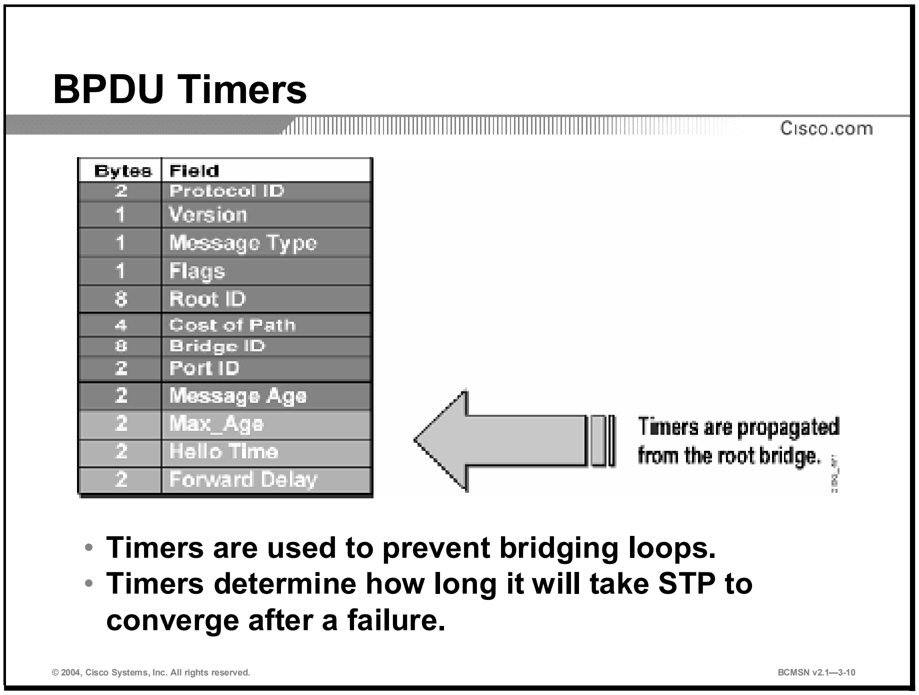

BPDU Timers

Cisco.com

Bytes	Field
2	Protocol ID
1	Version
1	Message Type
1	Flags
8	Root ID
4	Cost of Path
8	Bridge ID
2	Port ID
2	Message Age
2	Max_Age
2	Hello Time
2	Forward Delay

Timers are propagated from the root bridge.

- Timers are used to prevent bridging loops.
- Timers determine how long it will take STP to converge after a failure.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—3-10

BPDU timers specify a set period of time by which ports wait for topology information during propagation delays. The following three timers affect STP performance and state changes:

- **hello time:** The hello time is the time between each BPDU that is sent on a port. This is equal to 2 seconds by default, but can be tuned to be between 1 and 10 seconds.
- **forward delay:** The forward delay is the time spent in the listening and learning state. This is by default equal to 15 seconds for each state, but can be tuned to be between 4 and 30 seconds.
- **max_age:** The max age timer controls the maximum length of time a switch port saves its configuration BPDU information. This is 20 seconds by default, but can be tuned to be between 6 and 40 seconds.

When STP is enabled, every switch in the network goes through the blocking state and the transitory states of listening and learning at power up. The ports then stabilize to the forwarding or blocking state. During a topology change, a port temporarily implements the listening and learning states for a specified period called the “forward delay interval.”

Initially, all switch ports start in the blocking state when they listen for BPDUs. When a switch first boots up, it thinks it is the root switch and transitions to the listening state for the forward delay period of 15 seconds by default. In the listening state, a port is able to send and receive BPDUs to determine the active topology. During the listening state, the bridge performs the these functions:

- It elects the root bridge.
- It elects the root ports on the nonroot bridges.
- It elects the designated ports on each segment.

Switch ports that are not the designated or root ports will transition back to the blocking state. A switch port in the learning state populates its MAC address table with MAC addresses heard on the ports, but it does not forward user data frames. If a port is still a designated or root port at the end of the learning state, the port will transition to the forwarding state. Ports that are not the designated or root port will transition back to the blocking state. In the forwarding state, a port is capable of sending and receiving user data.

Identifying Spanning Tree Port States

This topic identifies the features and functions that apply to each port state.

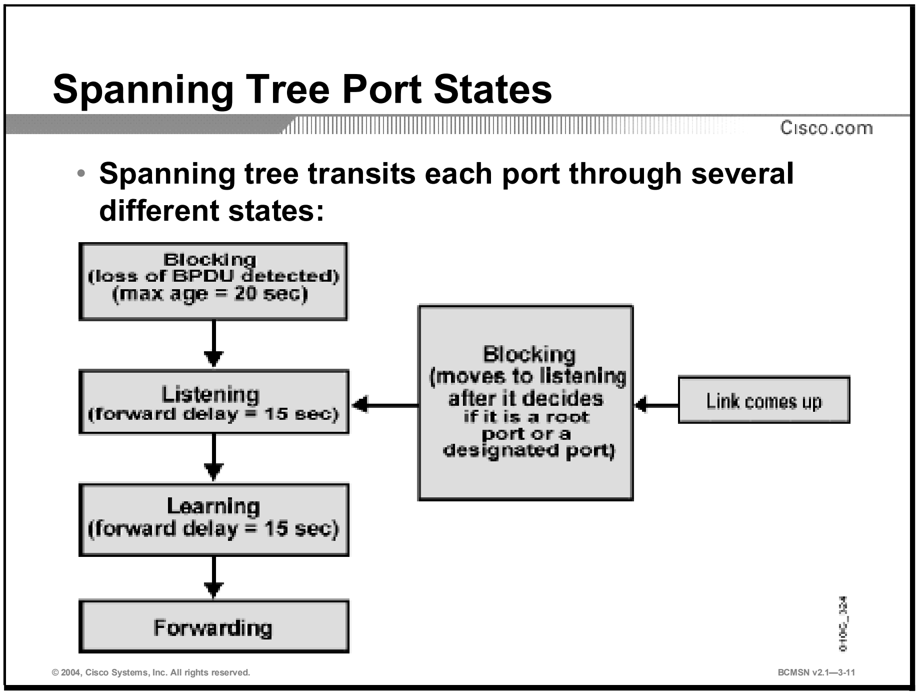

Spanning Tree Port States

- **Spanning tree transits each port through several different states:**

Blocking (loss of BPDU detected) (max age = 20 sec)

Listening (forward delay = 15 sec)

Learning (forward delay = 15 sec)

Forwarding

Blocking (moves to listening after it decides if it is a root port or a designated port)

Link comes up

Propagation delays can occur when protocol information passes through a switched LAN. As a result, topology changes can take place at different times and at different places in a switched network. When a Layer 2 interface transitions directly from nonparticipation in the spanning tree topology to the forwarding state, the interface can create temporary data loops. Ports must wait for new topology information to propagate through the switched LAN before starting to forward frames.

Each Layer 2 interface on a switch that uses spanning tree exists in one of these five states:

- **Blocking:** In this state, the Layer 2 interface is a nondesignated port and does not participate in frame forwarding but does receive BPDUs. The port begins to evaluate the BPDUs to determine where the root switch is and what port role (root, designated, or nondesignated) the switch port should assume in the final active STP topology.
- **Listening:** This state is the first transitional state after the blocking state when spanning tree determines that the Layer 2 interface could participate in frame forwarding according to the BPDUs that the switch has received thus far. At this point, the switch port is not only receiving BPDUs, it is also transmitting its own BPDUs and informing any adjacent switches that the switch port is preparing to participate in the active topology.
- **Learning:** In this state, the Layer 2 interface prepares to participate in the active topology and frame forwarding and begins to populate the MAC table.
- **Forwarding:** In this state, the Layer 2 interface is considered part of the active topology and forwards frames and also sends and receives BPDUs.
- **Disabled:** In this state, the Layer 2 interface does not participate in spanning tree and does not forward frames.

Identifying Topology Changes

This topic identifies the steps involved in a topology change.

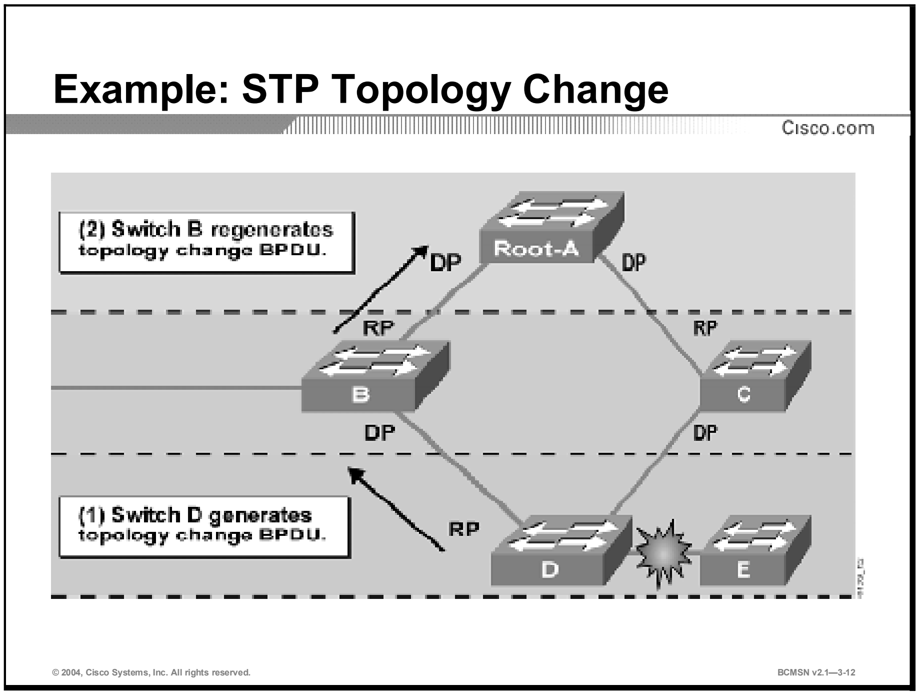

Example: STP Topology Change

Cisco.com

(2) Switch B regenerates topology change BPDU.
DP
Root-A
DP
RP
RP
B
C
DP
DP
(1) Switch D generates topology change BPDU.
RP
D
E

BCMSN v2.1—3-12

Example: Identifying Topology Changes

The steps that occur in a topology change are as follows:

Step 1 Switch D notices that a change to a link has occurred.

Step 2 Switch D sends a TCN BPDU out the root port destined ultimately for the root switch. The switch will send out the TCN BPDU until the designated switch responds with a topology change acknowledgement.

Step 3 Switch B, the designated switch, sends out a topology change acknowledgement to the originating switch D. The designated switch also sends a TCN BPDU out the root port destined for either the designated switch or the root switch. (This is a propagation TCN.)

Step 4 When the root switch receives the topology change message, the root switch changes its configuration BPDU to indicate that a topology change is occurring. The root switch sets the topology change in the configuration for a period of time equal to the sum of the forward delay and max_age parameters, which is approximately 50 seconds.

Step 5 A switch receiving the topology change configuration message from the root bridge uses the forward delay timer to age out entries in the MAC address table. This time specification allows the switch to age out MAC address, switch port, and VLAN mapping entries faster than the normal five-minute default. The bridge continues this process until it no longer receives topology change configuration messages from the root bridge.

Step 6 The backup link, if there is one, is enabled and the address table is repopulated.

Summary

This topic summarizes the key points discussed in this lesson.

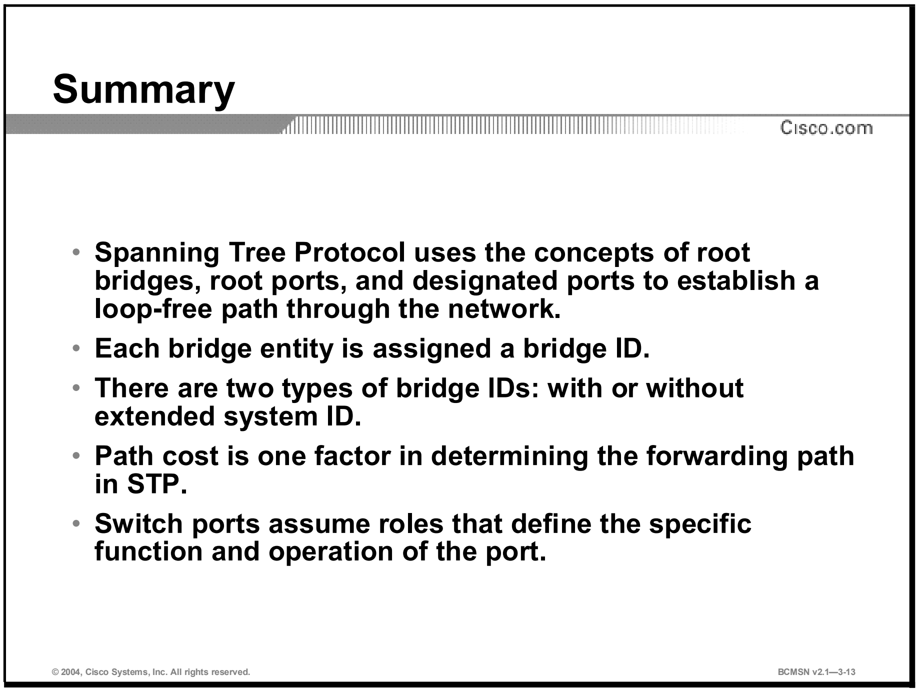

Summary

- Spanning Tree Protocol uses the concepts of root bridges, root ports, and designated ports to establish a loop-free path through the network.
- Each bridge entity is assigned a bridge ID.
- There are two types of bridge IDs: with or without extended system ID.
- Path cost is one factor in determining the forwarding path in STP.
- Switch ports assume roles that define the specific function and operation of the port.

BCMSN v2.1—3-13

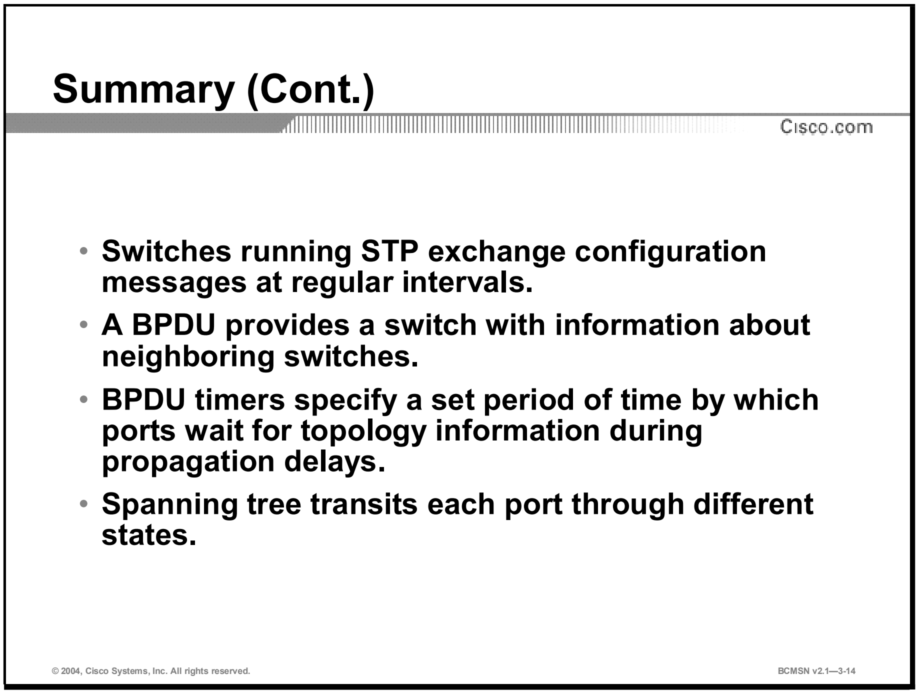

Summary (Cont.)

- Switches running STP exchange configuration messages at regular intervals.
- A BPDU provides a switch with information about neighboring switches.
- BPDU timers specify a set period of time by which ports wait for topology information during propagation delays.
- Spanning tree transits each port through different states.

BCMSN v2.1—3-14

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Select the most likely root bridge based on the bridge IDs.

A) 54

B) 107

C) 34

D) 79

Q2) Match each bridge ID with the correct description.

_____ 1. bridge ID with extended system ID

_____ 2. bridge ID without extended system ID

A) This bridge ID is made up of a bridge priority value and a bridge MAC address.

B) Bridge priority bits include the four most significant bits of a bridge ID.

C) A single MAC address must be allocated for each STP instance.

D) This bridge ID does not rely on MAC addresses to make it locally unique.

Q3) Based on the diagram, select the path cost to determine the forwarding path between a device and the root switch.

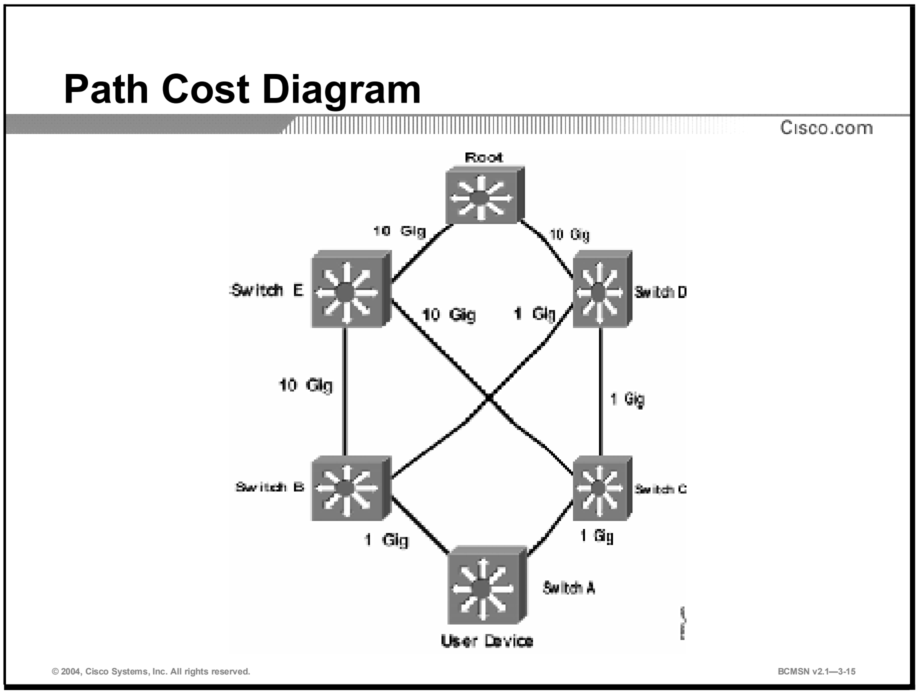

Path Cost Diagram

Cisco.com

BCMSN v2.1—3-15

A) Switch A, Switch C, Switch D, Root

B) Switch A, Switch B, Switch E, Root

C) Switch A, Switch C, Switch E, Root

D) Switch A, Switch B, Switch D, Root

Q4) Match the correct features with the appropriate port role.

_____ 1. designated port

_____ 2. root port

_____ 3. nondesignated port

_____ 4. disabled port

A) switch port with the least cost path to the root switch

B) port that is shut down

C) capable of populating the MAC table

D) exists on nonroot or designated switches only

E) blocks data frames

Q5) Select the features that correctly apply to BPDU messaging. (Choose two.)

A) determines the lowest cost path to the root switch

B) the election of a root switch as a Layer 3 topology point of reference

C) the election of a designated switch and corresponding designated port for every switched segment

D) removal of loops in the switched network by disabling ports

Q6) Match the BPDU field name with the appropriate field definition.

_____ 1. port ID

_____ 2. root ID

_____ 3. bridge ID

_____ 4. bost of path

_____ 5. flags

A) contains switch priority, extended system ID, and MAC address of the transmitting switch

B) combination of the switch port number and port priority value

C) value from the transmitting switch to the calculated root switch

D) BID of the root switch according to the transmitting switch

Q7) Match the timer with the appropriate time period

_____ 1. 15 seconds

_____ 2. 2 seconds

_____ 3. 20 seconds

A) hello time

B) forward delay

C) max_age

Q8) Match the STP port states with their appropriate features.

_____ 1. forwarding

_____ 2. learning

_____ 3. blocking

_____ 4. disabled

_____ 5. listening

A) The Layer 2 interface does not participate in spanning tree.

B) The Layer 2 interface receives BPDUs and evaluates them to determine the location of the root switch.

C) This is the first transitional state after the blocking state.

D) The Layer 2 interface forwards frames and also sends and receives BPDUs.

E) The Layer 2 interface populates the MAC table.

Q9) Match each topology change step with the correct step number.

_____ 1. Step 1

_____ 2. Step 2

_____ 3. Step 3

_____ 4. Step 4

_____ 5. Step 5

_____ 6. Step 6

A) Switch A sends a TCN BPDU out the root port destined for the root switch.

B) Any switch receiving the topology change message from the root bridge uses the forward delay timer to age out entries in the MAC address table.

C) Switch A notices that a change to a link has occurred.

D) The backup link is enabled, and the address table is repopulated.

E) The designated switch sends out a topology change acknowledgement to Switch A.

F) The root switch changes its configuration BPDU to indicate that a topology change is occurring.

Quiz Answer Key

Q1) C

Relates to: Selecting the Root Switch

Q2) 1=B, D; 2=A, C

Relates to: Distinguishing Between Bridge IDs

Q3) C

Relates to: Determining the Forwarding Path

Q4) 1=C, 2=A, D, 3=E, 4=B

Relates to: Defining Spanning Tree Port Roles

Q5) A, C

Relates to: Using BPDU Messaging

Q6) 1=B, 2=D, 3=A, 4=C

Relates to: Identifying Selected BPDU Field Definitions

Q7) 1=B, 2=A, 3=C

Relates to: Transitioning from One STP State to Another

Q8) 1=D, 2=E, 3=B, 4=A, 5=C

Relates to: Identifying Spanning Tree Port States

Q9) 1=C, 2=A, 3=E, 4=F, 5=B, 6=D

Relates to: Identifying Topology Changes

Establishing a Loop-Free Topology in a New Network

Overview

Spanning tree establishes a loop-free network topology by determining the root bridge, forming an association with the root switch, selecting the root, and selecting the designated port. Spanning tree will behave in certain ways depending on the makeup of the network topology.

Relevance

Spanning tree is a critical feature of a multilayer switched network. The configuration of spanning tree can have a huge impact on network performance.

Objectives

Upon completing this lesson, you will be able to:

- Select the correct criteria used to determine the root bridge
- Identify the process steps to determine the root switch
- Select the root
- Select the steps used by switches to determine the forwarding path to the root switch
- Identify the switch port that will become the root port based on given path costs and port IDs
- Select the features of the designated switch port
- Match the features to the correct STP implementation
- Predict the behavior of the STP with the Enterprise Composite Network model

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Identifying Root Bridge Criteria
- Identifying Root Switch Process Steps
- Selecting the Root
- Forming an Association with the Root Switch
- Selecting the Root Port
- Selecting the Designated Port
- Comparing STP and Per VLAN Spanning Tree+
- Predicting STP Behavior in the Enterprise Composite Network Model
- Summary
- Quiz

Identifying Root Bridge Criteria

This topic identifies the criteria used to determine the root bridge.

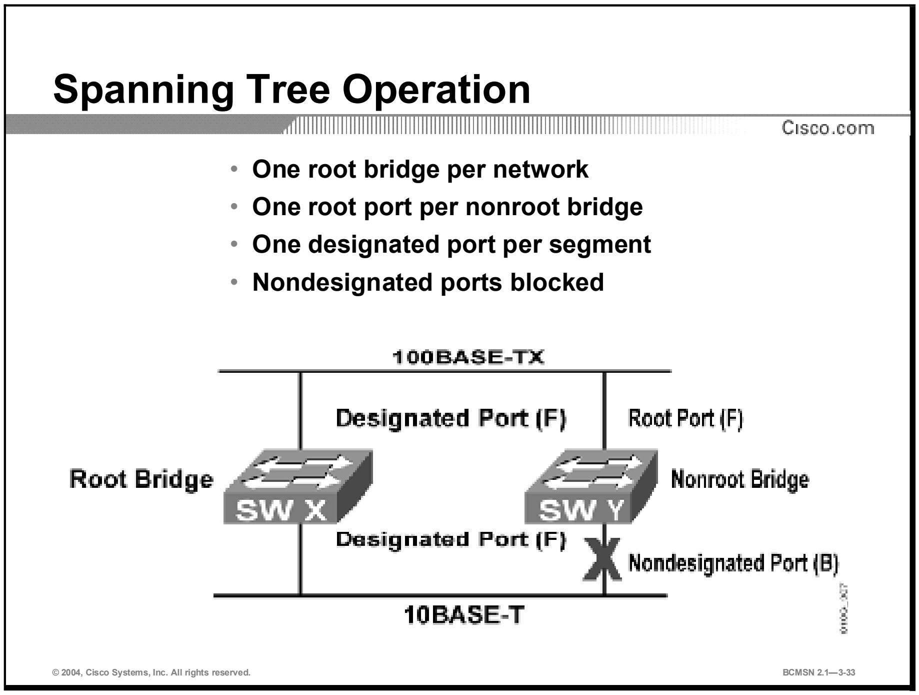

Spanning Tree Operation

Cisco.com

- One root bridge per network
- One root port per nonroot bridge
- One designated port per segment
- Nondesignated ports blocked

BCMSN 2.1—3-33

STP initially converges on a logically loop-free network topology by performing these steps:

Step 1 **Elects one root switch:** STP initially receives BPDUs from all attached segments on disabled ports in the blocking state. Only one switch can act as the root switch in a given STP topology. On the root switch, all ports are designated ports and are in the forwarding state. This is true except in the case that a switch port on the root switch is attached to itself. In that case, the switch port moves to the blocking state. In the example, switch X is elected as the root bridge.

Step 2 **Selects the root port on all nonroot switches:** STP establishes one root port on each nonroot switch. The root port is the lowest cost path from the nonroot switch to the root switch. In the example, the lowest cost path from switch Y to the root switch is through the 100BASE-TX Fast Ethernet link.

Step 3 **Selects the designated port on each segment:** STP establishes one designated port per nonroot switch per segment. The designated port is selected on the switch that has the lowest path cost of all the switches on that segment to the root switch. In the example, the 10BASE-T Ethernet port on switch Y is a nondesignated port because the switch port on the Fast Ethernet segments offers a lower cost path. Nondesignated ports receive BPDUs but do not forward traffic to logically break the loop topology.

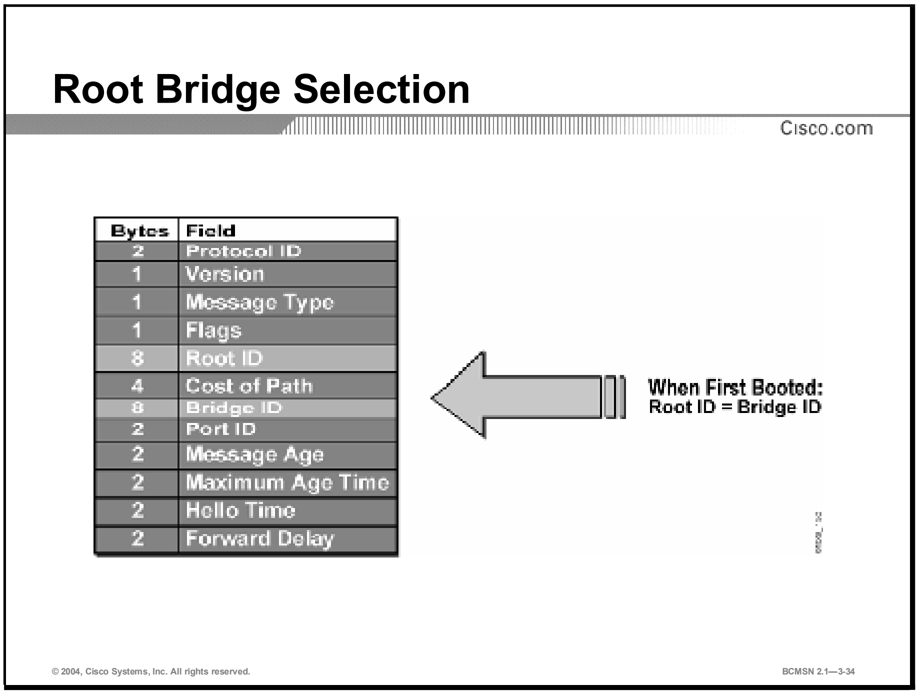

Root Bridge Selection

Cisco.com

Bytes	Field
2	Protocol ID
1	Version
1	Message Type
1	Flags
8	Root ID
4	Cost of Path
8	Bridge ID
2	Port ID
2	Message Age
2	Maximum Age Time
2	Hello Time
2	Forward Delay

When First Booted:
Root ID = Bridge ID

BCMSN 2.1—3-34

The first step in creating the loop-free spanning tree is to elect a root switch. The root switch is the reference point that all switches use to determine if there are loops in the network.

Identifying Root Switch Process Steps

This topic identifies the process steps to determine the root switch.

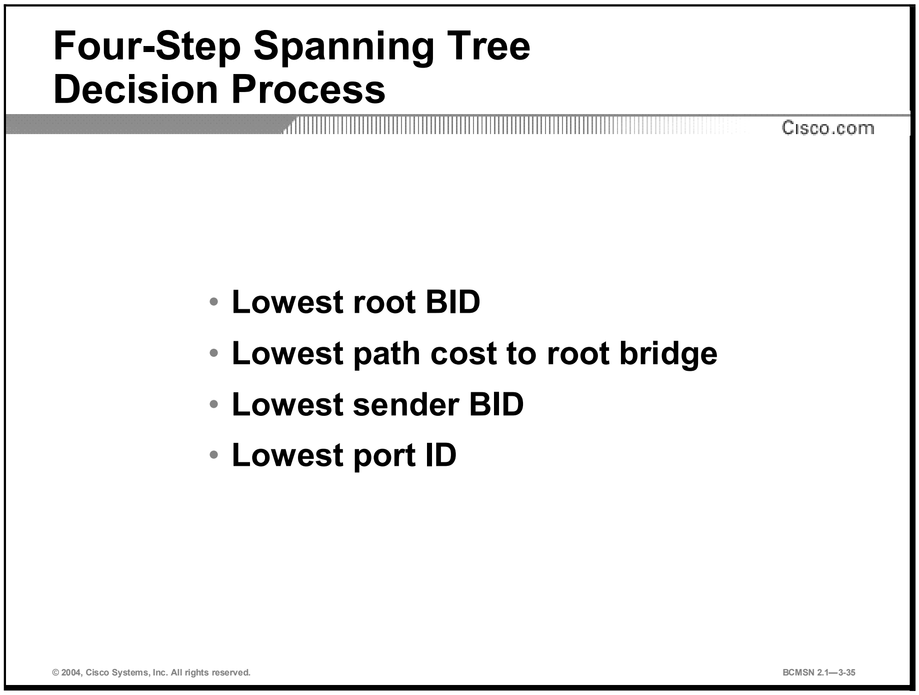

Four-Step Spanning Tree Decision Process

Cisco.com

- **Lowest root BID**
- **Lowest path cost to root bridge**
- **Lowest sender BID**
- **Lowest port ID**

BCMSN 2.1—3-35

When determining the loop-free topology, STP identifies the root switch by evaluating BPDUs. Four criteria are used in the decision-making process, in the following order:

- Lowest root BID
- Lowest path cost to the root switch
- Lowest sender BID
- Lowest local port ID

When STP determines the root switch of an STP topology, the switch first examines the received BPDUs for the lowest root BID. If there are equal BID values reported by two different switches claiming to be the root, the switch considers the path cost to each of the two switches. The switch may have two or more equal cost paths to the prospective root switches. If this situation occurs, STP examines the BID of the switches that sent the BPDUs for the two eligible root switches. If the BIDs are again equal, STP looks at the switch port ID and selects the lowest one as the root port. Generally, the root is determined on the first criterion, the lowest root BID.

Selecting the Root

This topic describes the process for selecting the root.

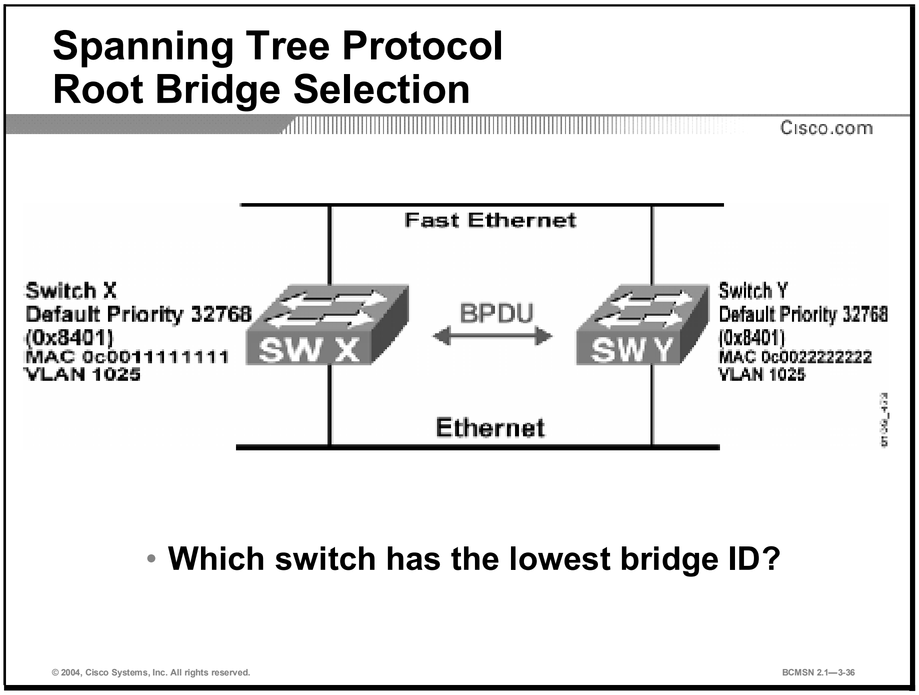

Spanning Tree Protocol Root Bridge Selection

Cisco.com

- **Which switch has the lowest bridge ID?**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-36

At startup, the switch assumes that it is the root bridge and sets the root ID equal to its BID. If all devices have the same priority for the same LAN, the bridge with the lowest MAC address becomes the root bridge.

Both switches are using the same default priority and are calculating for the same VLAN. The switch with the lowest MAC address will be the root switch. In the example, switch X is the root switch with a bridge ID of 0x8401:0c0011111111.

Forming an Association with the Root Switch

This topic identifies the features used by switches to determine the forwarding path to the root switch.

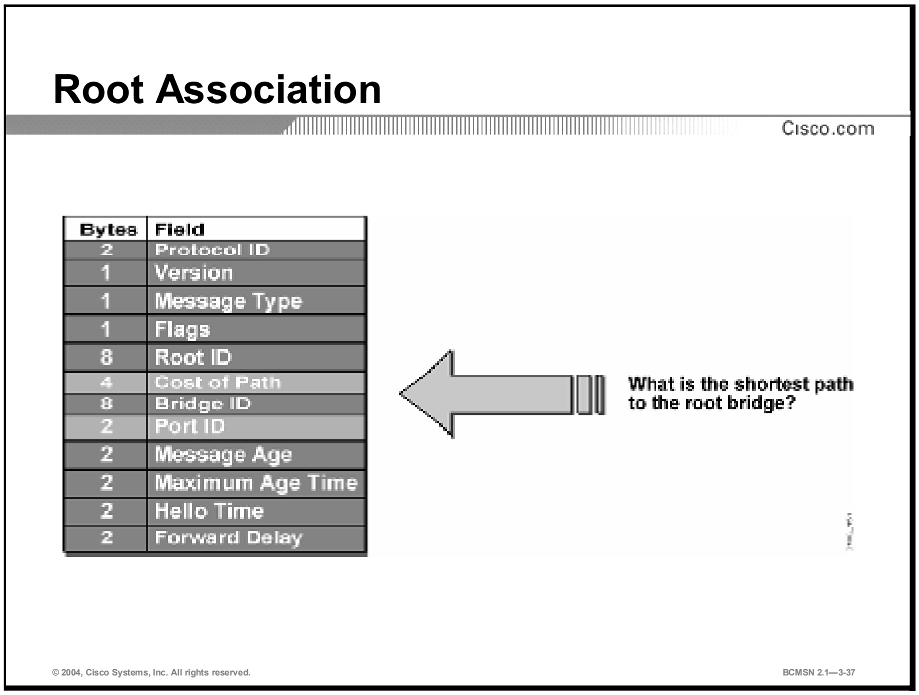

Root Association

Bytes	Field
2	Protocol ID
1	Version
1	Message Type
1	Flags
8	Root ID
4	Cost of Path
8	Bridge ID
2	Port ID
2	Message Age
2	Maximum Age Time
2	Hello Time
2	Forward Delay

What is the shortest path to the root bridge?

BCMSN 2.1—3-37

After the root switch has been elected, all nonroot switches must form an association with the root switch. Each switch does this by listening to BPDUs as they come in on all ports. Receiving BPDUs on multiple ports indicates a redundant path to the root switch.

The switch looks at these two components in the BPDU to determine which switch ports will forward data and which switch ports will block data:

- Path cost
- Port ID

The switch looks at the path cost first to determine which port is receiving the lowest cost path. The path is calculated based on the link speed and the number of links the BPDU traversed. If a port has the lowest cost, that port is eligible to be placed in forwarding mode. All other ports that are receiving BPDUs continue in blocking mode.

If the path cost is equal, as in the case of like parallel links, the bridge goes to the port ID as a "tiebreaker." The port with the lowest port ID forwards, and all other ports continue to block.

Selecting the Root Port

This topic identifies how a switch port is selected as the root port.

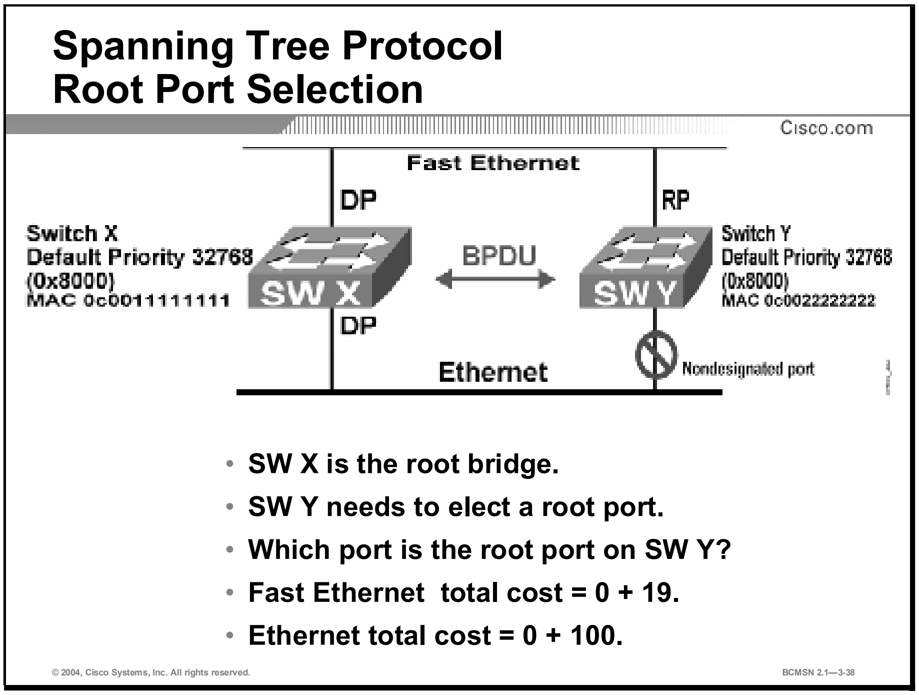

Spanning Tree Protocol Root Port Selection

- SW X is the root bridge.
- SW Y needs to elect a root port.
- Which port is the root port on SW Y?
- Fast Ethernet total cost = 0 + 19.
- Ethernet total cost = 0 + 100.

BCMSN 2.1—3-38

Each nonroot switch running STP must elect a root port. The root port is determined by the lowest path cost to the root switch. If two or more switch ports have the same path cost to the root switch, the switch port with the lowest port ID becomes the root port.

Example: Selecting the Root Port

Switch Y receives a BPDU from the root switch (switch X) on its switch port on the Fast Ethernet segment and another on its switch port on the Ethernet segment. The root path cost in both cases is zero. The local path cost on the Fast Ethernet switch port is 19, while the local path cost on the Ethernet switch port is 100. As a result, the switch port on the Fast Ethernet segment has the lowest path cost to the root switch and is elected the root port for switch Y.

Selecting the Designated Port

This topic identifies the features that apply to designated switch ports.

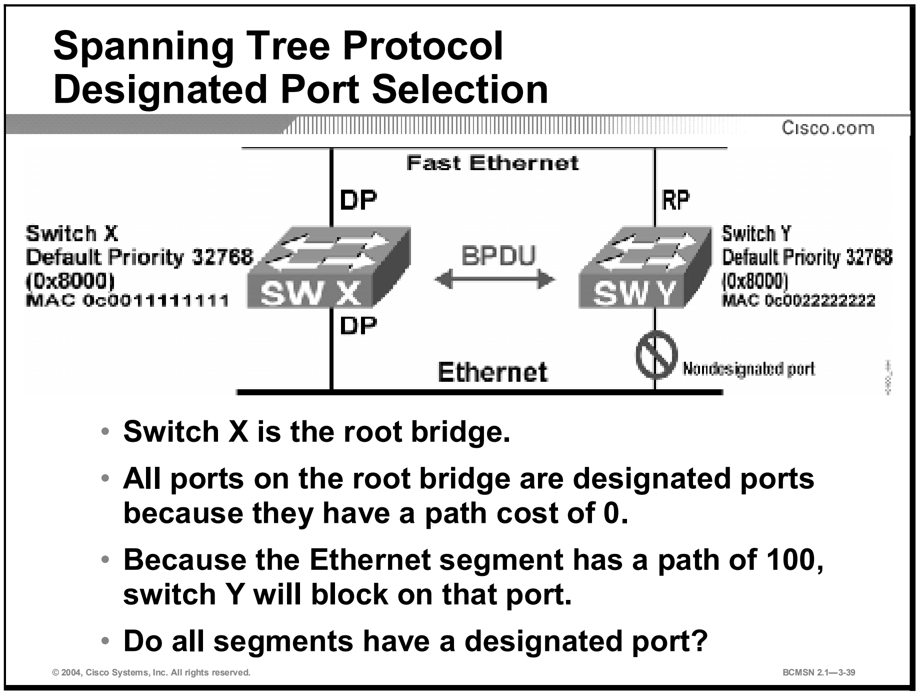

Spanning Tree Protocol Designated Port Selection

- Switch X is the root bridge.
- All ports on the root bridge are designated ports because they have a path cost of 0.
- Because the Ethernet segment has a path of 100, switch Y will block on that port.
- Do all segments have a designated port?

BCMSN 2.1—3-39

STP selects one designated port per segment to forward traffic. Other switch ports on the segment become nondesignated ports and continue blocking. The nondesignated ports receive BPDUs but do not forward data traffic to prevent loops. The switch port on the segment with the lowest path cost to the root switch is elected the designated port. If multiple switch ports on the same segment have the same cost, the switch port with the lowest port ID becomes the designated port.

Because ports on the root bridge all have a root path cost of zero, all ports on the root bridge are designated ports.

Comparing STP and Per VLAN Spanning Tree+

This topic identifies the features that apply to Per VLAN Spanning Tree+ (PVST+).

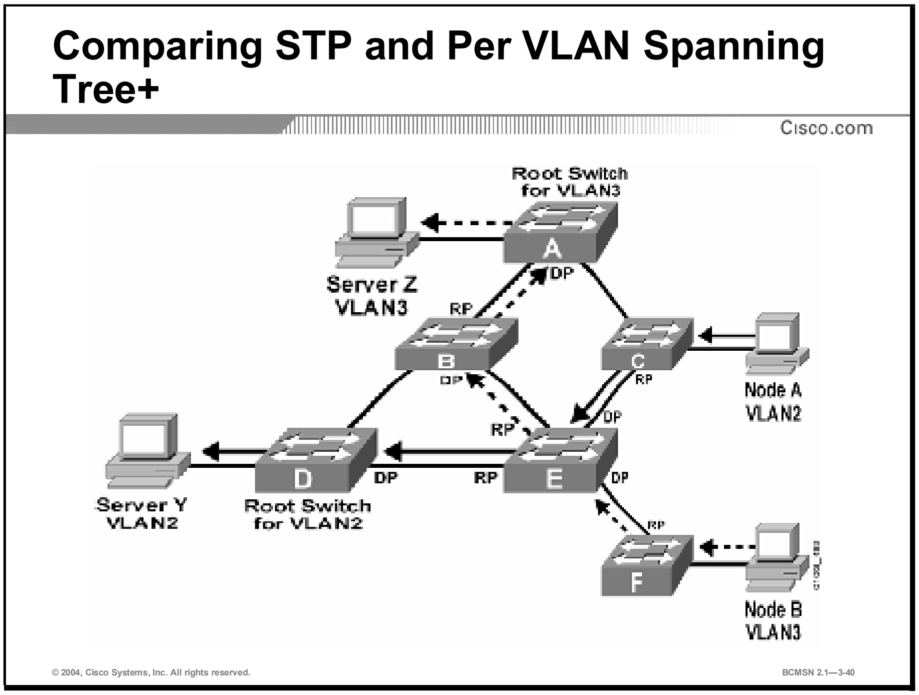

Comparing STP and Per VLAN Spanning Tree+

Cisco.com

BCMSN 2.1—3-40

PVST+ maintains a separate spanning tree instance for each VLAN.

PVST+ is fully compatible with the 802.1Q trunking protocol and with ISL. PVST+ runs the same STA that 802.1D does and provides the same functionality, to prevent Layer 2 loops. The difference is that PVST+ is still a Cisco proprietary protocol and runs a separate instance of the STA for each VLAN. This means that for every VLAN created, a separate root switch, a separate set of designated switches, and associated port roles and states are calculated.

Example: Comparing STP and PVST+

It is possible to create different logical topologies using the VLANs and PVST+ on your network. In the example, there is a switched network that has implemented two different logical topologies by defining two different root switches on two different VLANs. The network administrator has adjusted the PVST+ parameters to manually set the root switch to be the switch closest to the destination resource. The only difference is that the STP parameters, such as bridge priority, have been configured on a per-VLAN basis.

With 802.1D, only one root switch can be calculated for the entire switched network. One or both of the VLANs have to take an indirect or suboptimal path to reach the resource servers. Switch B is equidistant between the two servers and might be the logical root switch selection. However, with the Enterprise Composite Network model and the use of local VLANs and resource servers being centralized in a server farm, this particular type of network configuration is not practical, efficient, or optimal. The ability to define a separate root switch on a per-VLAN basis still has merit.

One of the biggest advantages of PVST+ is the ability to load balance end-user traffic across the switched domain uplinks. The functionality of STP is now extended from providing only redundancy where there is an idle backup link to using the backup link as a primary link for another VLAN.

When trunking with ISL, the PVST+ BPDU information is encapsulated, and the VID is identified in the ISL header. When trunking with 802.1Q, the BPDU is tagged with the VID or is received on the native VLAN. The Port VLAN ID (PVID) is used to identify the source VLAN.

Note The PVST+ configuration is not necessarily automatic. You need to plan and configure it manually to ensure this traffic flow.

Predicting STP Behavior in the Enterprise Composite Network Model

This topic identifies the behavior of STP when applied to the Enterprise Composite Network model.

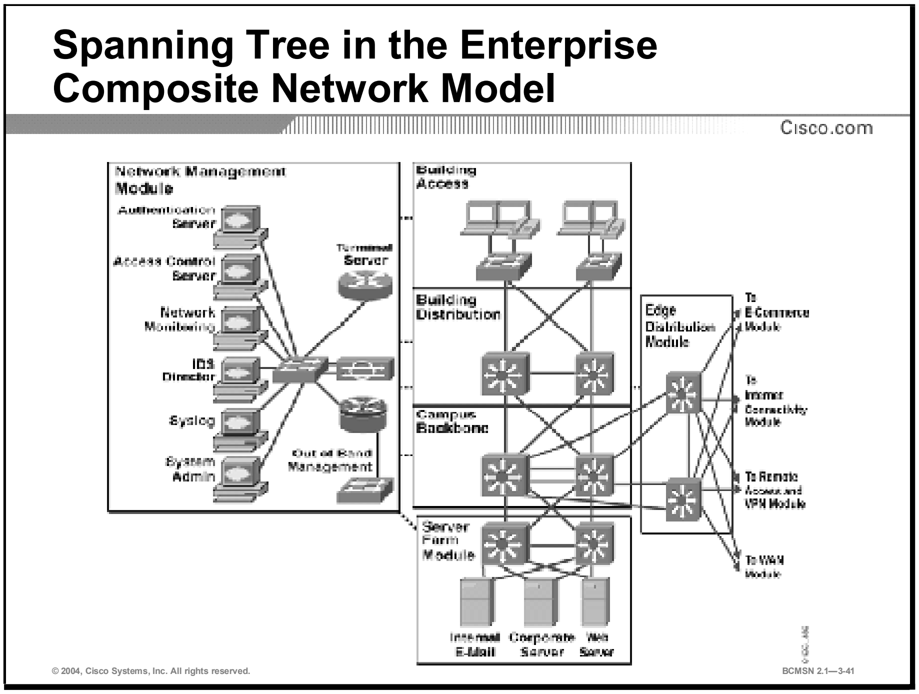

Spanning Tree in the Enterprise Composite Network Model

BCMSN 2.1—3-41

Using the Enterprise Composite Network model reduces the need for STP in present-day networks. In a typical Layer 2 LAN network, the end user has one link to the access-layer switch. The access-layer switch has two uplinks to the Building Distribution layer on different Layer 2 switches. One switch at the Building Distribution layer will be the root; the other will be the backup root, thus specifying STP operation and configuration.

Summary

This topic summarizes the key points discussed in this lesson.

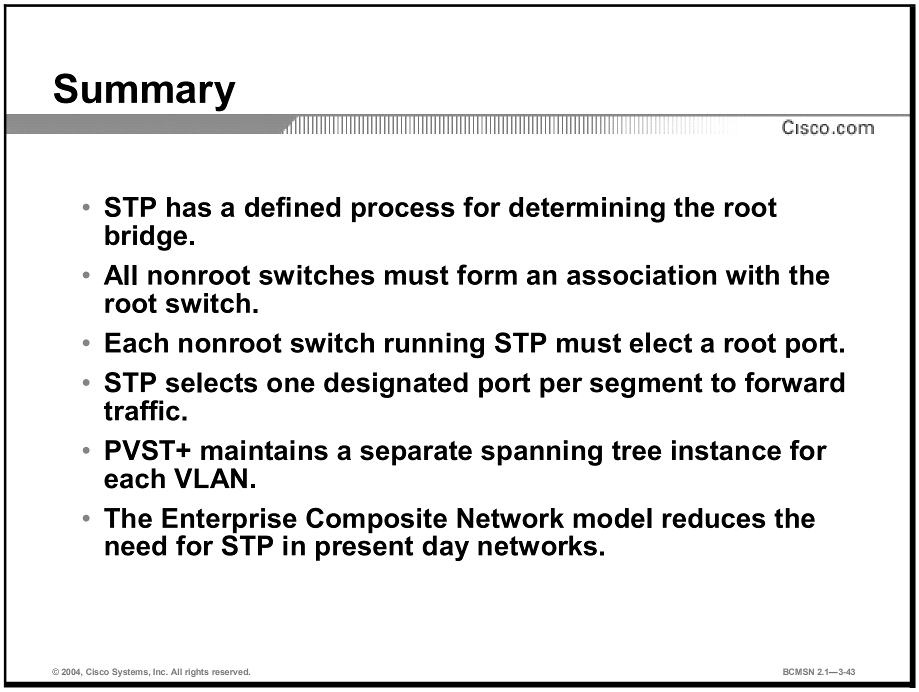

Summary

- STP has a defined process for determining the root bridge.
- All nonroot switches must form an association with the root switch.
- Each nonroot switch running STP must elect a root port.
- STP selects one designated port per segment to forward traffic.
- PVST+ maintains a separate spanning tree instance for each VLAN.
- The Enterprise Composite Network model reduces the need for STP in present day networks.

BCMSN 2.1—3-43

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Select the criteria used to determine the root bridge. (Choose two.)

A) one designated port per segment

B) two root bridges per network

C) nondesignated ports forwarded

D) one root port per nonroot bridge

Q2) Match each step with the correct order to determine the root switch.

_____ 1. Step 1

_____ 2. Step 2

_____ 3. Step 3

_____ 4. Step 4

A) lowest sender BID

B) lowest root BID

C) lowest local port ID

D) lowest path cost to root switch

Q3) Select the root switch based on the following bridge IDs.

A) 0x8402:0c0003233333333

B) 0x8402:0c0002222222222

C) 0x8402:0c0002444444444

D) 0x8402:0c0003332333344

Q4) Select the steps used by switches to determine the forwarding path to the root switch. (Choose two.)

A) switch looks at bridge ID

B) switches form an association with the root switch

C) switch looks at path cost

D) switches watch packet speed

Q5) Switch A receives a BPDU message from the root switch. The message is received on its Ethernet segment and its Fast Ethernet segment. The root path cost on the Ethernet segment is 5, while the root path cost on the Fast Ethernet segment is 6. The local path cost on both switch ports is 15. Which switch port is elected as the root port for switch A?

A) Fast Ethernet segment

B) Ethernet segment

C) both segments

D) neither segment

Q6) Select the features that apply to designated switch ports. (Choose two.)

A) STP selects two designated ports per segment.

B) All nondesignated ports block traffic.

C) Switch ports with the lowest path cost to the root switch are elected as the designated port.

D) If multiple switch ports on the same segment have the same cost, the switch port with the lowest bridge ID becomes the designated port.

Q7) Match the appropriate features with the correct STP implementation.

_____ 1. STP implementation

_____ 2. PVST+ implementation

A) It runs a separate instance of the STA for each VLAN.

B) It provides greater flexibility and granularity for Layer 2 switched networks.

C) BPDU is tagged with the VLAN ID or is received on the native VLAN.

D) For every VLAN created, a separate root switch, a separate set of designated switches, and associated port roles and states are calculated.

E) It provides only redundancy.

Q8) Based on the diagram, why is spanning tree no longer an issue? (Choose two.)

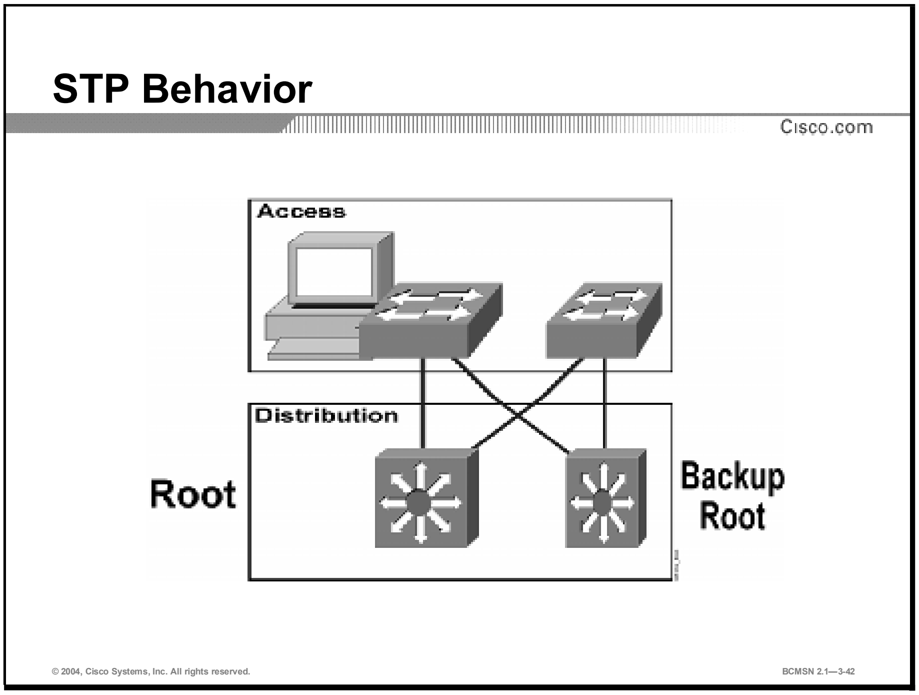

STP Behavior

Cisco.com

Access

Distribution

Root

Backup Root

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-42

A) The end user has one link to the Building Access layer.

B) There is no backup root to cause confusion.

C) The access-layer switch has two uplinks to the Distribution layer switch.

D) The root exists at the access-layer switch.

Quiz Answer Key

Q1) A, D

Relates to: Identifying Root Bridge Criteria

Q2) 1=B, 2=D, 3=A, 4=C

Relates to: Identifying Root Switch Process Steps

Q3) B

Relates to: Selecting the Root

Q4) B, C

Relates to: Forming an Association with the Root Switch

Q5) B

Relates to: Selecting the Root Port

Q6) B, C

Relates to: Selecting the Designated Port

Q7) 1=A, B, D; 2=C, E

Relates to: Comparing STP and Per VLAN Spanning Tree+

Q8) A, C

Relates to: Predicting STP Behavior in the Enterprise Composite Network Model

Configuring Spanning Tree Protocol

Overview

Cisco Catalyst switches support spanning tree in PVST+ mode by default. You configure STP on a per-VLAN basis, designating root and secondary root bridges, and influencing the likelihood of ports being selected for the forwarding state.

Relevance

Spanning tree is a critical feature of a multilayer switched network, and the configuration of spanning tree can have a huge impact on network performance. Understanding the proper configuration methods will help ensure that your network operates at full efficiency.

Objectives

Upon completing this lesson, you will be able to:

- Identify the commands that enable spanning tree on a per-VLAN basis
- Identify the commands that configure the switch as the root bridge
- Identify the commands that set the port cost and VLAN port cost on a specific switch
- Identify the command and variables that verify a spanning tree configuration for a VLAN and interface

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Enabling Spanning Tree
- Configuring the Root Bridge
- Setting the Port Cost
- Verifying the STP Configuration
- Summary
- Quiz

Enabling Spanning Tree

This topic identifies the commands that enable spanning tree on a per-VLAN basis.

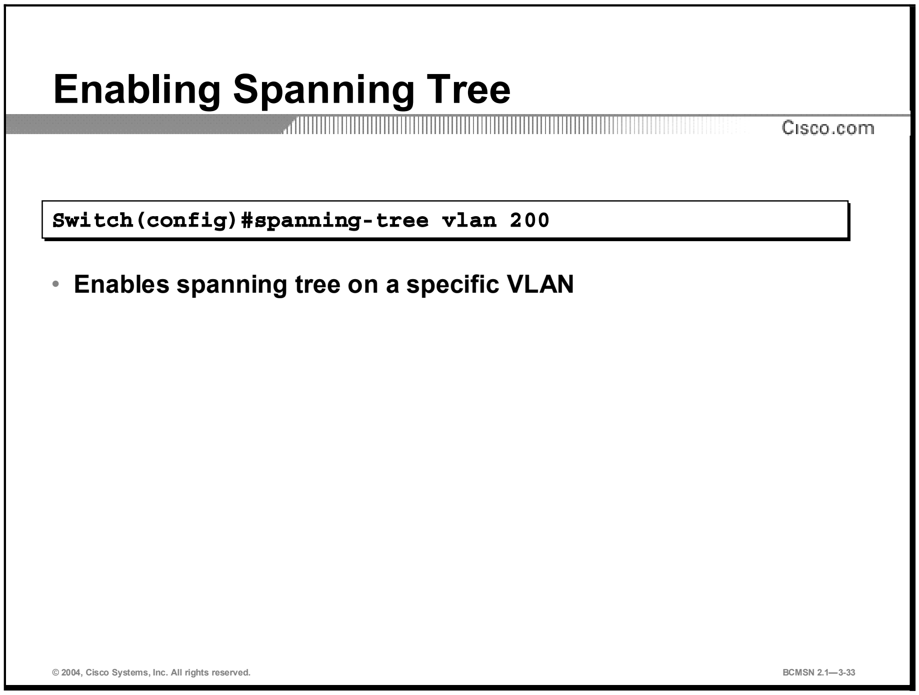

Enabling Spanning Tree

Cisco.com

```
Switch(config)#spanning-tree vlan 200
```

- **Enables spanning tree on a specific VLAN**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-33

You enable spanning tree on a per-VLAN basis. The switch maintains a separate instance of spanning tree for each VLAN (except on VLANs on which you have disabled a spanning tree). By default, spanning tree is enabled on all VLANs; therefore, there is no need to take action to enable STP. If you have disabled spanning tree and need to re-enable it for a particular VLAN, you can use the following command from global configuration mode:

```
Switch(config)#spanning-tree vlan vlan_ID
```

This same command is used with additional arguments to configure various features of STP.

Example: Enabling Spanning Tree

This example shows how to enable spanning tree on VLAN200:

```
Switch#configure terminal
Switch(config)#spanning-tree vlan 200
Switch(config)#end
```

Configuring the Root Bridge

This topic identifies the commands to configure a switch as the root bridge.

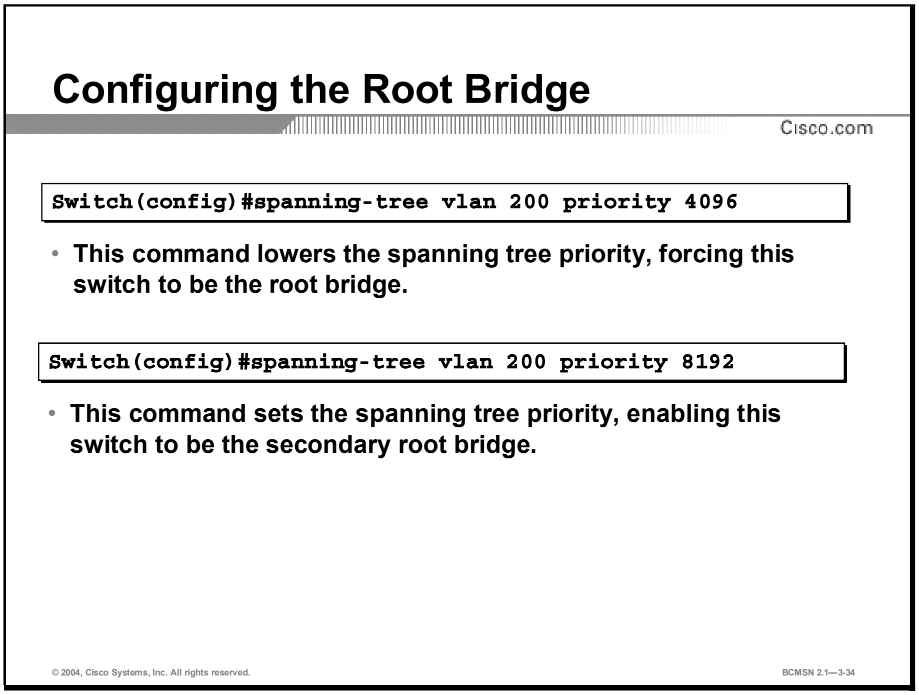

Configuring the Root Bridge

Cisco.com

```
Switch(config)#spanning-tree vlan 200 priority 4096
```

- **This command lowers the spanning tree priority, forcing this switch to be the root bridge.**

```
Switch(config)#spanning-tree vlan 200 priority 8192
```

- **This command sets the spanning tree priority, enabling this switch to be the secondary root bridge.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-34

The switch with the lowest BID will become the root bridge for a VLAN. You can use a configuration command to specify that this switch is the root bridge for the VLAN.

A Catalyst switch running PVST+ maintains an instance of spanning tree for each active VLAN configured on the switch. A BID, consisting of the bridge priority and the bridge MAC address, is associated with each instance. For each VLAN, the switch with the lowest BID will become the root bridge for that VLAN. Whenever the bridge priority changes, the BID also changes. This result is the recomputation of the root bridge for the VLAN.

To configure a switch to become the root bridge for a specified VLAN, lower its priority from the default value with the command **spanning-tree vlan** *vlan-ID* **priority** *value*. Assuming the other bridges in the VLAN retain their default priority, this switch will become the root bridge. The value 4096 is used by convention.

Note The root switch for each instance of spanning tree should be a Building Distribution switch. On a network with a collapsed backbone and Building Distribution layer, one of these backbone switches and or distribution switches should be the root switch. Do not configure a Building Access switch as the spanning tree primary root.

A secondary root is a switch that may become the root bridge for the specified VLAN if the primary root bridge fails. You specify a switch as the secondary root bridge by setting the priority to a value between the low value of the root bridge (4096) and the default value (32768). The priority value 8192 is used.

You can run this command on more than one switch to configure multiple backup root switches.

Setting the Port Cost

This topic identifies the commands that set the port cost and VLAN port cost on a specific switch.

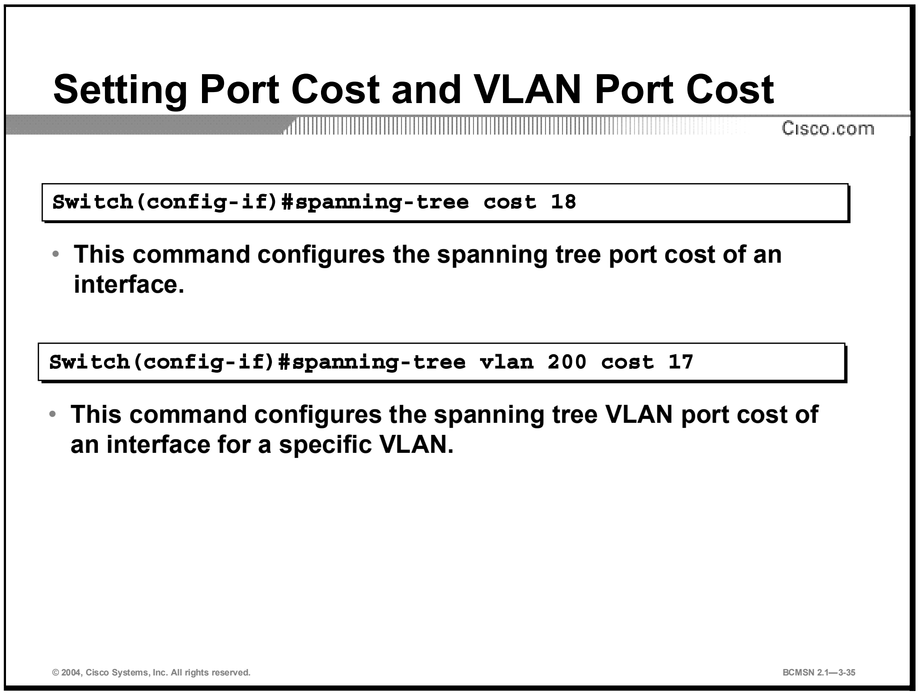

Setting Port Cost and VLAN Port Cost

Cisco.com

```
Switch(config-if)#spanning-tree cost 18
```

- **This command configures the spanning tree port cost of an interface.**

```
Switch(config-if)#spanning-tree vlan 200 cost 17
```

- **This command configures the spanning tree VLAN port cost of an interface for a specific VLAN.**

BCMSN 2.1—3-35

Spanning tree considers port cost when selecting an interface to put into the forwarding state. You can assign lower cost values to interfaces that you want spanning tree to select first.

The default value for spanning tree port path cost is derived from the interface media speed. In the event of a loop, spanning tree considers port cost when selecting an interface to put into the forwarding state. You can assign lower cost values to interfaces that you want spanning tree to select first, and higher cost values to interfaces that you want spanning tree to select last. If all interfaces have the same cost value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks other interfaces. The possible cost range is 1 through 200,000,000 (the default is media-specific).

Spanning tree uses the port cost value when the interface is configured as an access port and uses VLAN port cost values when the interface is configured as a trunk port.

To configure the spanning tree port cost or VLAN port cost of an interface, use one of these commands:

```
Switch(config-if)#spanning-tree cost port_cost
Switch(config-if)#spanning-tree vlan vlan_ID cost port_cost
```

Example: Changing the Spanning Tree Port Cost

This example shows how to change the spanning tree port cost of a Fast Ethernet interface to 17, making it more likely to be chosen as a forwarding port than another Fast Ethernet interface configured with the default port cost (19 for Fast Ethernet):

```
Switch#configure terminal
Switch(config)#interface fastethernet 5/8
Switch(config-if)#spanning-tree cost 17
Switch(config-if)#end
```

This example shows how to configure the spanning tree VLAN port cost of a Fast Ethernet interface to 20, making it less likely to be chosen as a forwarding port than another Fast Ethernet interface configured with the default port cost (19 for Fast Ethernet):

```
Switch# configure terminal
Switch(config)#interface fastethernet 5/8
Switch(config-if)#spanning-tree vlan 200 cost 20
Switch(config-if)#end
```

Verifying the STP Configuration

This topic identifies the commands and variables that verify a spanning tree configuration for a VLAN and an interface.

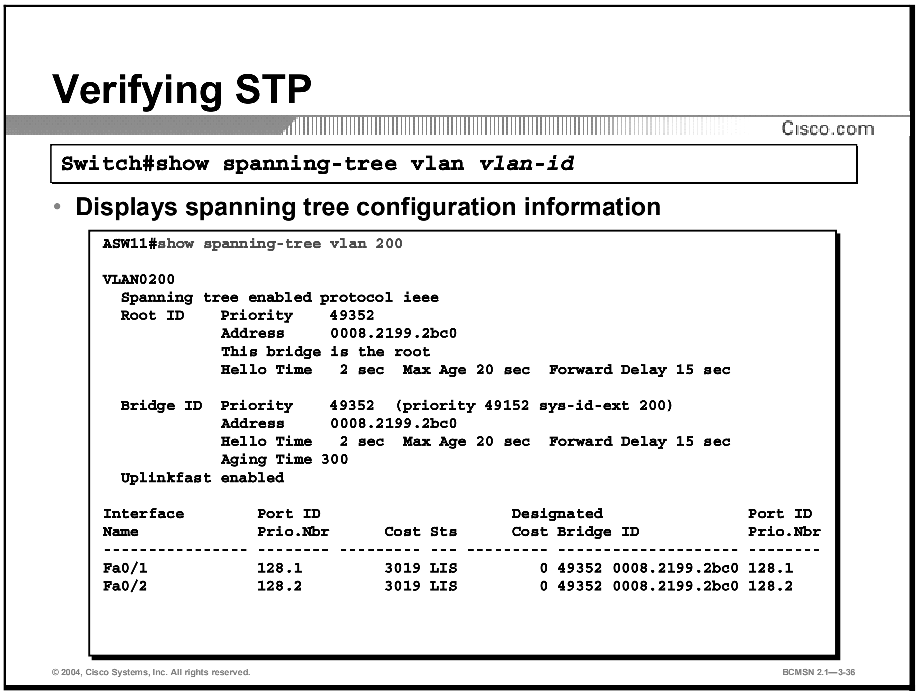

Verifying STP

```
Switch#show spanning-tree vlan vlan-id
```

- **Displays spanning tree configuration information**

```
ASW11#show spanning-tree vlan 200

VLAN0200
  Spanning tree enabled protocol ieee
  Root ID    Priority    49352
             Address     0008.2199.2bc0
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    49352  (priority 49152 sys-id-ext 200)
             Address     0008.2199.2bc0
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300
  Uplinkfast enabled

Interface        Port ID                       Designated               Port ID
Name             Prio.Nbr      Cost Sts      Cost Bridge ID            Prio.Nbr
---------------- -------- --------- --- --------- -------------------- --------
Fa0/1            128.1         3019 LIS         0 49352  0008.2199.2bc0 128.1
Fa0/2            128.2         3019 LIS         0 49352  0008.2199.2bc0 128.2
```

BCMSN 2.1—3-36

You can use a variety of **show** commands to display configuration and operation information about spanning tree. The **show spanning-tree** command takes several arguments to display a variety of information about the STP configuration. Without any arguments, it will display general information about all STP configurations. The complete syntax is as follows:

```
Switch#show spanning-tree [bridge-group | active |
backbonefast | {bridge [id]}| detail | inconsistentports |
{interface interface interface-number} | root | summary
[total] | uplinkfast | {vlan vlan-id} | {port-channel number}
| pathcost-method]
```

Refer to your software documentation for a complete explanation of each parameter.

Example: Verifying Spanning Tree

This example shows how to verify that spanning tree is enabled on VLAN200. This example is on a nonroot bridge:

```
Switch#show spanning-tree vlan 200

 VLAN200 is executing the ieee compatible Spanning Tree
protocol
  Bridge Identifier has priority 32768, address 0050.3e8d.6401
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 16384, address 0060.704c.7000
  Root port is 264 (FastEthernet5/8), cost of root path is 38
  Topology change flag not set, detected flag not set
  Number of topology changes 0 last change occurred 01:53:48
ago
  Times:  hold 1, topology change 24, notification 2
          hello 2, max age 14, forward delay 10
  Timers: hello 0, topology change 0, notification 0

 Port 264 (FastEthernet5/8) of VLAN200 is forwarding
   Port path cost 19, Port priority 128, Port Identifier
129.9.
   Designated root has priority 16384, address 0060.704c.7000
   Designated bridge has priority 32768, address
00e0.4fac.b000
   Designated port id is 128.2, designated path cost 19
   Timers: message age 3, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   BPDU: sent 3, received 3417
   Timers:message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state:1
   Link type is point-to-point by default
   BPDU:sent 187, received 1

Port 129 (FastEthernet3/1) of VLAN1002 is forwarding
   Port path cost 19, Port priority 128, Port Identifier
128.129.
   Designated root has priority 32768, address 0003.6b10.ebe9
   Designated bridge has priority 32768, address
0003.6b10.ebe9
   Designated port id is 128.129, designated path cost 0
   Timers:message age 0, forward delay 0, hold 0
```

Example: Verifying a Spanning Tree Interface Configuration

This example shows how to display the details of the interface configuration for an interface that is configured as an access port in several VLANs:

```
Switch#show spanning-tree interface fastethernet 3/1 detail

Port 129 (FastEthernet3/1) of VLAN0001 is forwarding
   Port path cost 19, Port priority 128, Port Identifier
128.129.
   Designated root has priority 32768, address 0003.6b10.e800
   Designated bridge has priority 32768, address
0003.6b10.e800
   Designated port id is 128.129, designated path cost 0
   Timers:message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state:1
   Link type is point-to-point by default
   BPDU:sent 187, received 1

Port 129 (FastEthernet3/1) of VLAN1002 is forwarding
   Port path cost 19, Port priority 128, Port Identifier
128.129.
   Designated root has priority 32768, address 0003.6b10.ebe9
   Designated bridge has priority 32768, address
0003.6b10.ebe9
   Designated port id is 128.129, designated path cost 0
   Timers:message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state:1
   Link type is point-to-point by default
   BPDU:sent 94, received 2

 Port 129 (FastEthernet3/1) of VLAN1003 is forwarding
   Port path cost 19, Port priority 128, Port Identifier
128.129.
   Designated root has priority 32768, address 0003.6b10.ebea
   Designated bridge has priority 32768, address
0003.6b10.ebea
   Designated port id is 128.129, designated path cost 0
   Timers:message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state:1
   Link type is point-to-point by default
   BPDU:sent 94, received 2
 Port 129 (FastEthernet3/1) of VLAN1004 is forwarding
   Port path cost 19, Port priority 128, Port Identifier 
128.129.
   Designated root has priority 32768, address 0003.6b10.ebeb
   Designated bridge has priority 32768, address 
0003.6b10.ebeb
   Designated port id is 128.129, designated path cost 0
   Timers:message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state:1
   Link type is point-to-point by default
   BPDU:sent 95, received 2

 Port 129 (FastEthernet3/1) of VLAN1005 is forwarding
   Port path cost 19, Port priority 128, Port Identifier 
128.129.
   Designated root has priority 32768, address 0003.6b10.ebec
   Designated bridge has priority 32768, address 
0003.6b10.ebec
   Designated port id is 128.129, designated path cost 0
   Timers:message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state:1
   Link type is point-to-point by default
   BPDU:sent 95, received 2
Switch#
```

Summary

This topic summarizes the key points discussed in this lesson.

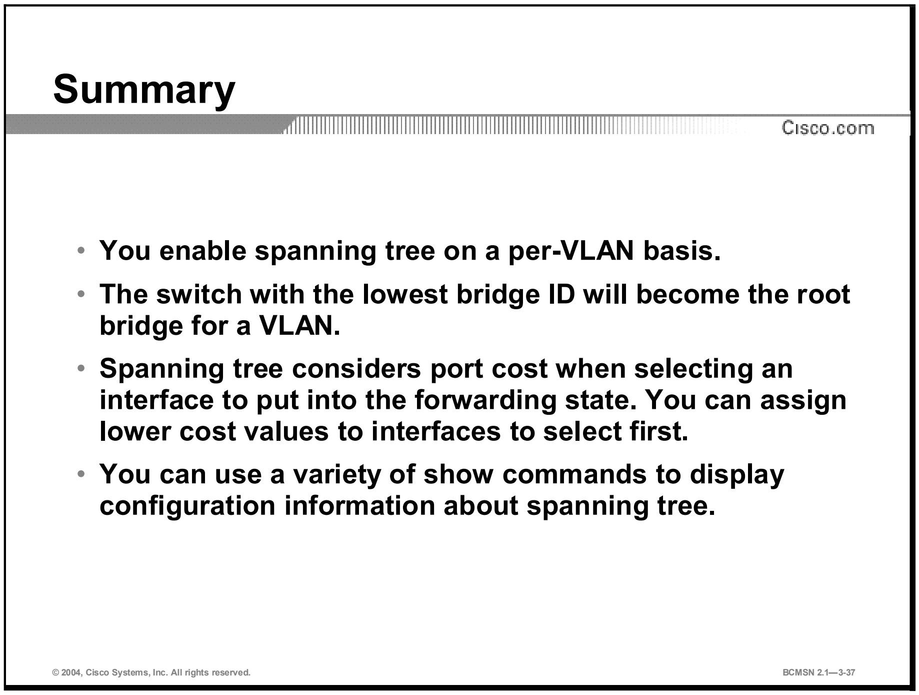

Summary

- You enable spanning tree on a per-VLAN basis.
- The switch with the lowest bridge ID will become the root bridge for a VLAN.
- Spanning tree considers port cost when selecting an interface to put into the forwarding state. You can assign lower cost values to interfaces to select first.
- You can use a variety of show commands to display configuration information about spanning tree.

BCMSN 2.1—3-37

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 3-1: Implementing Spanning Tree Protocol

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which command correctly enables spanning tree for VLAN200 on a Catalyst switch?

A) switch(config)#**spanning-tree vlan 200**

B) switch(config-if)#**spanning-tree vlan 200**

C) switch(config)#**spanning-tree mode pvst vlan 200**

D) switch(config-if)#**spanning-tree mode pvst vlan 200**

Q2) Which command correctly configures the Catalyst switch as the root bridge for VLAN10, assuming that the other switches in the VLAN maintain their default priority?

A) switch(config)#**spanning-tree 10 4096**

B) switch(config-if)#**spanning-tree vlan 10 4096**

C) switch(config)#**spanning-tree vlan 10 priority 4096**

D) switch(config-if)#**spanning-tree vlan 10 priority 4096**

Q3) Which command correctly sets the port cost of an interface in VLAN200 to 15?

A) switch(config)#**spanning-tree cost 15 vlan 200**

B) switch(config)#**spanning-tree vlan 200 cost 15**

C) switch(config-if)#**spanning-tree vlan 200 cost 15**

D) switch(config-if)#**spanning-tree cost 15 vlan 200**

Q4) Which command correctly displays STP configuration information specific to interface FastEthernet 5/6?

A) switch#**show spt interface fastethernet 5/6**

B) switch#**show spanning-tree fastethernet 5/6 detail**

C) switch#**show spanning-tree interface fastethernet 5/6**

D) switch#**show interface fastethernet 5/6 spanning-tree detail**

Quiz Answer Key

Q1) A

Relates to: Enabling Spanning Tree

Q2) C

Relates to: Configuring the Root Bridge

Q3) C

Relates to: Setting the Port Cost

Q4) C

Relates to: Verifying the STP Configuration

Lesson Assessments

Overview

Use the lesson assessments here to test what you learned in this module. The correct answers and solutions are found in the Lesson Assessment Answer Key.

Outline

This section includes these assessments:

- Quiz 3-1: Preventing Bridging Loops Using Spanning Tree Protocol
- Quiz 3-2: Defining Spanning Tree Protocol Operations
- Quiz 3-3: Establishing a Loop-Free Topology in a New Network
- Quiz 3-4: Configuring Spanning Tree Protocol

Quiz 3-1: Preventing Bridging Loops Using Spanning Tree Protocol

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Identify the features that apply to transparent bridging
- Identify the behavior of flooded unicast frames in a bridged loop
- Identify the behavior of broadcast frames in a bridged loop
- Select a definition that best describes how STP uses STA to prevent bridging loops in a redundant network

Quiz

Answer these questions:

Q1) How does STP enforce a loop-free path in the network?

A) by shutting down redundant links

B) by shutting down redundant bridges

C) by dropping packets form redundant links

D) by dropping packets form redundant bridges

Q2) What is the first step in creating a loop-free spanning tree?

A) elect a root bridge

B) determine which timers to use

C) allow the topology to converge

D) determine the spanning tree cost

Q3) A simple bridge environment can function smoothly. Problems arise when _____ are added.

A) transparent bridges

B) redundant links

C) blocking bridges

D) broadcast storms

Q4) What two behaviors will bring about a broadcast storm?

A) always retransmitting a broadcast

B) blocked bridges

C) “listening” to the source address

D) never marking the frame

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 67 percent or better.

Quiz 3-2: Defining Spanning Tree Protocol Operations

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Select the bridge ID that is most likely to be chosen as the root bridge by STP
- Match each bridge ID with the correct description
- Calculate path cost to determine the forwarding path between a device and the root switch
- Match the features the correctly apply to the appropriate port role
- Select the features that best apply to BPDU messaging
- Match the correct BPDU field name with the appropriate definition
- Identify the correct timer that affects the switch port transitioning from one STP state to another
- Match the correct features and function to the appropriate STP port state
- List the steps involved in topology changes in the correct order

Quiz

Answer these questions:

Q1) Why is a unique bridge ID per VLAN vital?

A) to discover where redundant links exist

B) because multiple STP processes are often running at the same time

C) because it enables switches to determine cost path

D) because it reduces the amount of flooding

Q2) The STP path cost is the accumulated total cost of the path from _____ to _____.

A) designated switch to the intermediate switch

B) root switch to the designated switch

C) designated switch to the blocked switch

D) intermediate switch to the root switch

Q3) STP exchanges configuration messages with other switches at _____ intervals.

A) 1-second

B) 2-second

C) 50-second

D) 30-second

Q4) Under which conditions is a TCN sent out? (Choose two.)

A) when there is a link failure

B) when switch ports begin blocking data

C) when switch ports begin forwarding data

D) when BPDUs are sent from designated ports to root ports

Q5) When a port is in the transitional state of listening and learning, what tasks can it perform?

A) can only send BPDUs

B) can only receive BPDUs

C) can receive and send BPDUs

D) cannot receive or send BPDUs

Q6) What is the time interval for ports to transition from the blocking state to the forwarding state?

A) 2 to 10 seconds

B) 15 to 19 seconds

C) 15 to 30 seconds

D) 30 to 50 seconds

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 75 percent or better.

Quiz 3-3: Establishing a Loop-Free Topology in a New Network

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Identify the steps to determine the root switch
- Identify the root ID of a switch port at start up
- Select the criteria to determine the forwarding path of a switch
- Predict the behavior of the switch ports after the designated port is determined
- Identify the criteria for determining the port cost

Quiz

Answer these questions:

Q1) What steps are taken when electing the root switch?

A) determine which switch ports will be forwarding data

B) evaluate incoming BPDUs to see which ones are reporting the lowest root bridge ID

C) select the designated port

D) determine the nondesignated ports

Q2) When the system first boots up, the root ID equals the _____.

A) designated port

B) nondesignated port

C) bridge ID

D) root port

Q3) What two BPDU components does the switch look at to determine which will be forwarding data? (Choose two.)

A) path cost

B) bridge ID

C) port ID

D) designated port

Q4) When one switch port has been selected as a designated port, what do the other switch ports continue to do?

A) receive BPDUs and continue to forward data

B) receive BPDUs but do not forward data

C) do not receive BPDUs but continue to forward data

D) do not receive BPDUs and do not forward data

Q5) How is the path cost determined when choosing which switch ports will be forwarding data?

A) link speed and number of links that BPDUs have to cross

B) number of links BPDUs have to cross

C) link speed

D) root bridge ID

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 75 percent or better.

Quiz 3-4: Configuring Spanning Tree Protocol

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Identify the commands that enable spanning tree on a per-VLAN basis
- Identify those commands that configure the switch as the root bridge
- Identify those commands that set the port cost and VLAN cost on a specific switch
- Identify the command and variables that verify a spanning tree configuration for a VLAN and interface

Quiz

Answer these questions:

Q1) Which command correctly enables spanning tree for VLAN101 on a Catalyst switch?

A) switch(config)#**spanning-tree 101**

B) switch(config-if)#**spanning-tree 101**

C) switch(config)#**spanning-tree vlan 101**

D) switch(config-if)#**spanning-tree vlan 101**

Q2) Which command correctly configures the Catalyst switch as a likely secondary root bridge for VLAN10?

A) switch(config)#**spanning-tree vlan 10 priority 4096**

B) switch(config)#**spanning-tree vlan 10 priority 8192**

C) switch(config-if)#**spanning-tree vlan 10 priority 4096**

D) switch(config-if)#**spanning-tree vlan 10 priority 8192**

Q3) Which command changes the spanning tree port cost of a Fast Ethernet interface to 25?

A) switch(config-if)#**spanning-tree cost 25**

B) switch(config-if)#**spanning-tree vlan 200 cost 25**

C) switch(config-if)#**spanning-tree cost 17**

D) switch(config-if)#**spanning-tree vlan 200 cost 20**

Q4) Which command correctly displays STP configuration information for all VLANs?

A) switch#**show spanning-tree all**

B) switch#**show spanning-tree bridge**

C) switch#**show spanning-tree detail**

D) switch#**show spanning-tree all detail**

Q5) Spanning tree uses VLAN port cost values when the interface is configured as _____.

A) an access port

B) a trunk port

C) designated port

D) root port

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 75 percent or better.

Lesson Assessment Answer Key

Quiz 3-1: Preventing Bridging Loops Using Spanning Tree Protocol

Q1) A

Q2) A

Q3) B

Q4) A, D

Quiz 3-2: Defining Spanning Tree Protocol Operations

Q1) B

Q2) D

Q3) B

Q4) A, C

Q5) C

Q6) D

Quiz 3-3: Establishing a Loop Free Topology in a New Network

Q1) B

Q2) C

Q3) A, C

Q4) B

Q5) A

Quiz 3-4: Configuring Spanning Tree Protocol

Q1) C

Q2) B

Q3) A

Q4) B

Q5) B

Module 4

Enhancing Spanning Tree Protocol

Overview

Cisco has developed several proprietary features to enhance Spanning Tree Protocol (STP). In addition, new standards have evolved to define Rapid Spanning Tree (RSTP) and Multiple Spanning Tree (MST). The Cisco EtherChannel feature provides a mechanism for treating multiple Ethernet interfaces as a single port, providing load balancing and effectively higher bandwidth capabilities.

Upon completing this module, you will be able to:

- Apply the appropriate protocol to enhance STP operation
- Select the benefits that RSTP provides
- Correctly match the appropriate STP commands with the appropriate protocol
- Match the appropriate features and commands to the correct protocol
- Select an approach for troubleshooting STP

Outline

The module contains these components:

- Optimizing Spanning Tree Protocol
- Accelerating Spanning Tree Convergence
- Configuring Spanning Tree Enhancements
- Tuning Spanning Tree Protocol
- Selecting a Troubleshooting Approach
- Lesson Assessments

Optimizing Spanning Tree Protocol

Overview

Cisco has developed several features to enhance STP, including PortFast, UplinkFast, and BackboneFast. STP has evolved to include standards for RSTP and MST.

Relevance

Proper configuration of STP versions and enhancements can have a significant impact on overall network performance.

Objectives

Upon completing this lesson, you will be able to:

- Identify the correct features that apply to PortFast
- Identify the correct features that apply to UplinkFast
- Identify the correct features that apply to BackboneFast
- Select the appropriate characteristic that applies to EtherChannel
- Match the correct features with the appropriate PAgP and LACP protocols
- Identify how STP enhances the Enterprise Composite Network model

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Applying PortFast
- Applying UplinkFast
- Applying BackboneFast
- Applying EtherChannel
- Comparing Port and Link Aggregation Protocols
- Using STP Enhancements in the Enterprise Composite Network Model
- Summary
- Quiz

Applying PortFast

This topic identifies the features of PortFast.

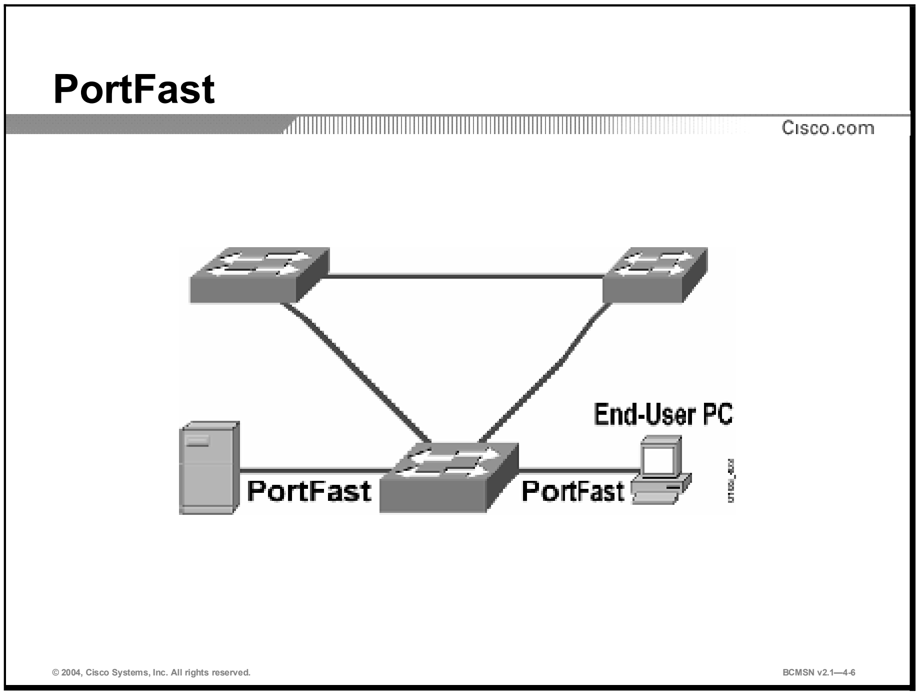

PortFast

Cisco.com

BCMSN v2.1—4-6

Spanning Tree PortFast causes an interface configured as a Layer 2 access port to enter the forwarding state immediately, bypassing the listening and learning states. You can use PortFast on Layer 2 access ports connected to a single workstation or server to allow those devices to connect to the network immediately, rather than waiting for spanning tree to converge. If the interface receives a bridge protocol data unit (BPDU), spanning tree can put the port into the blocking state.

Example: Applying PortFast

In the example, a server and workstation are attached to an access switch through ports that have been configured with PortFast. PortFast interfaces do not transition through all STP states, but they transition directly to forwarding.

If a link that is not attached to this port fails, the port does not transition directly to the forwarding state because it is already forwarding. This scenario does not affect the state of the PortFast port. The transition directly to the forwarding state occurs when a device is connected to the switch port.

Caution Because the purpose of PortFast is to minimize the time that access ports must wait for spanning tree to converge, it should be used only on access ports. If you enable PortFast on a port connecting to another switch, you risk creating a spanning tree loop.

Applying UplinkFast

This topic identifies the features of UplinkFast.

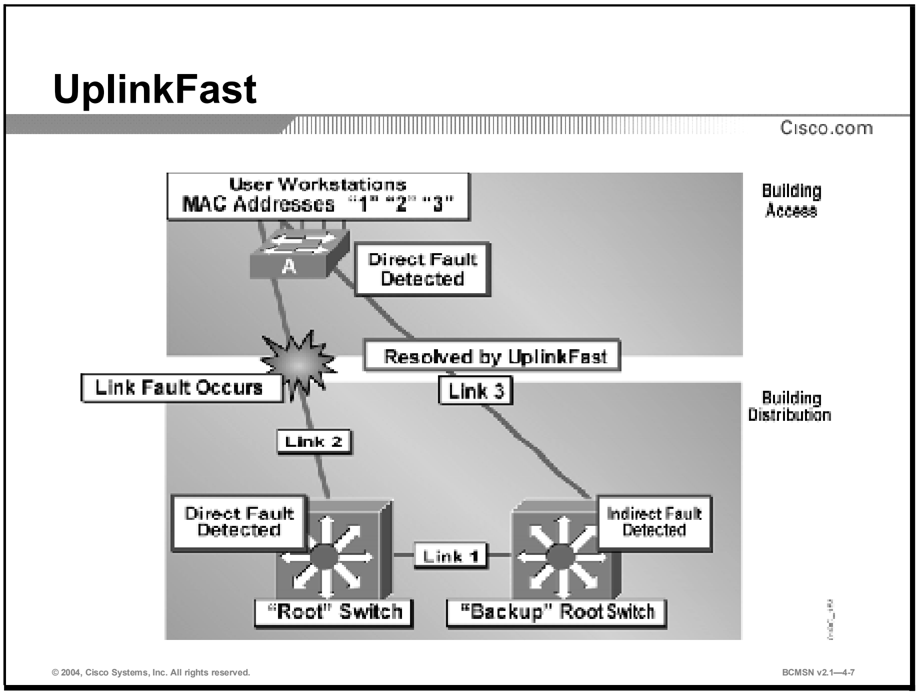

UplinkFast

Cisco.com

BCMSN v2.1—4-7

Spanning Tree UplinkFast provides fast convergence after a direct link failure. Layer 2 convergence is the speed and ability of a group of Layer 2 devices running STP to agree on a loop-free topology after a change in that topology. An uplink group is a set of Layer 2 interfaces (per VLAN), only one of which is forwarding at any given time. Specifically, an uplink group consists of the root port (which is forwarding) and a set of blocked ports. The uplink group provides an alternate path in case the currently forwarding link fails.

Example: Applying UplinkFast

The figure shows an example of a topology in which switch A is deployed in the Building Access submodule with uplink connections to the root switch over link 2 and the backup root switch over link 3. (Both switches are in the Building Distribution submodule.) Initially, the port on switch A connected to link 2 is in the forwarding state, and the port connected to link 3 is in the blocking state.

When switch A detects a link failure on the currently active link 2 on the root port (a direct link failure), UplinkFast unblocks the blocked port on switch A and transitions it to the forwarding state without going through the listening and learning states. This switchover occurs within 5 seconds.

Note UplinkFast is most useful in wiring-closet switches with at least one blocked port. This feature might not be useful for other types of applications.

Applying BackboneFast

This topic identifies the features of BackboneFast.

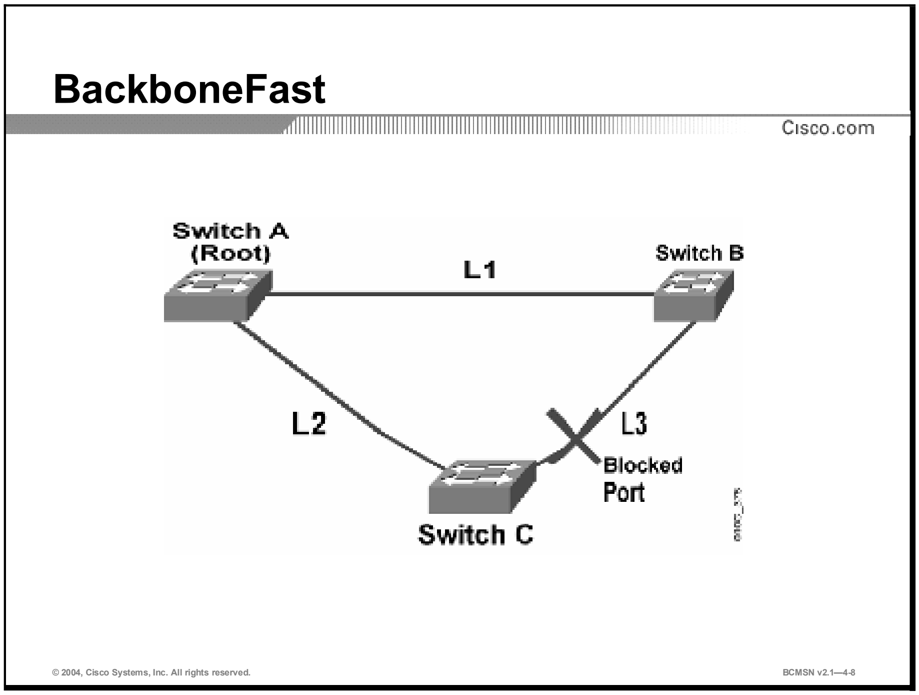

BackboneFast

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-8

BackboneFast addresses the situation where an indirect failure causes a topology change and a switch must find an alternative path through an intermediate switch.

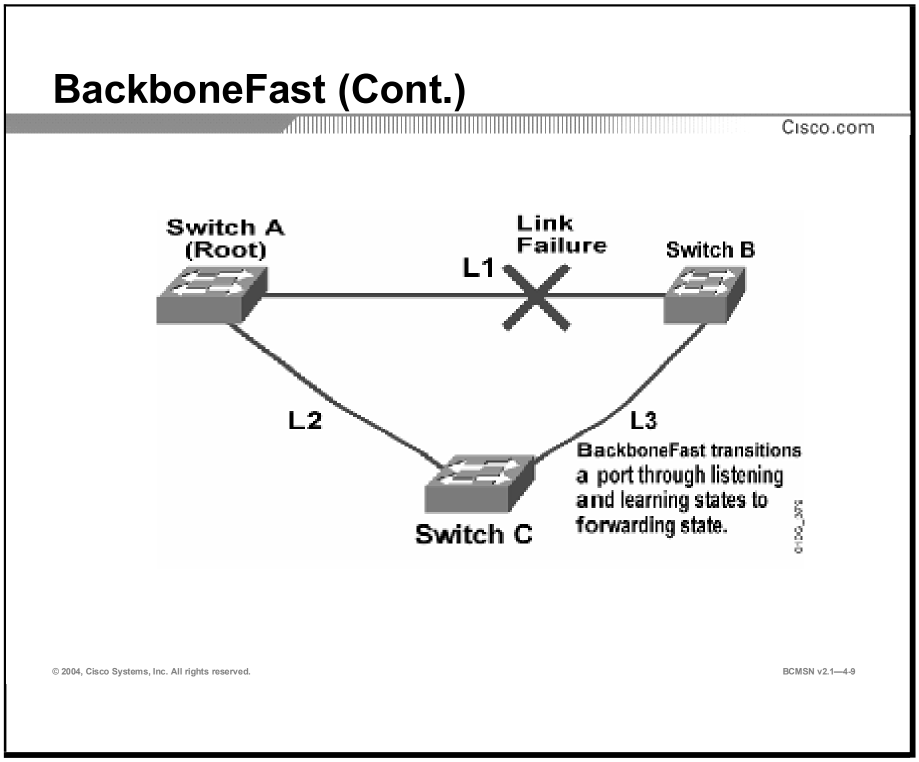

BackboneFast (Cont.)

BackboneFast transitions a port through listening and learning states to forwarding state.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-9

Example: BackboneFast Operation

BackboneFast is best illustrated by the failure of the link between the root and the backup root switch. The backup root switch does not assume that the root switch is still available. The backup switch will immediately block all previously forwarding ports and transmit configuration BPDUs claiming root responsibility. Since the root switch failure from the perspective of the backup switch is recent, the backup switch will set the designated and root port parameters to 1.

When the access switch receives the BPDU of the backup root switch, the access switch sees the BPDU as inferior; this is because its own root port is still active, and the backup root switch is claiming designated and root switch status. The access switch then transmits a special root query message to explicitly determine if the root switch is still active. Upon receipt of a response to the root query message, the access switch sends a BPDU using its known root switch parameters to the backup root switch and cycles the port adjacent to the backup root switch through the listening and learning states.

This differs from standard 802.1d spanning tree operation in that normally the blocked port does not process the received BPDUs until the max_age interval has expired. By using the BackboneFast feature, the network recovers from an indirect failure in two times the forward delay time, which is 30 seconds by default, rather than max_age plus two times forward delay time, which is 50 seconds by default.

Applying EtherChannel

This topic identifies the characteristics that apply to EtherChannel.

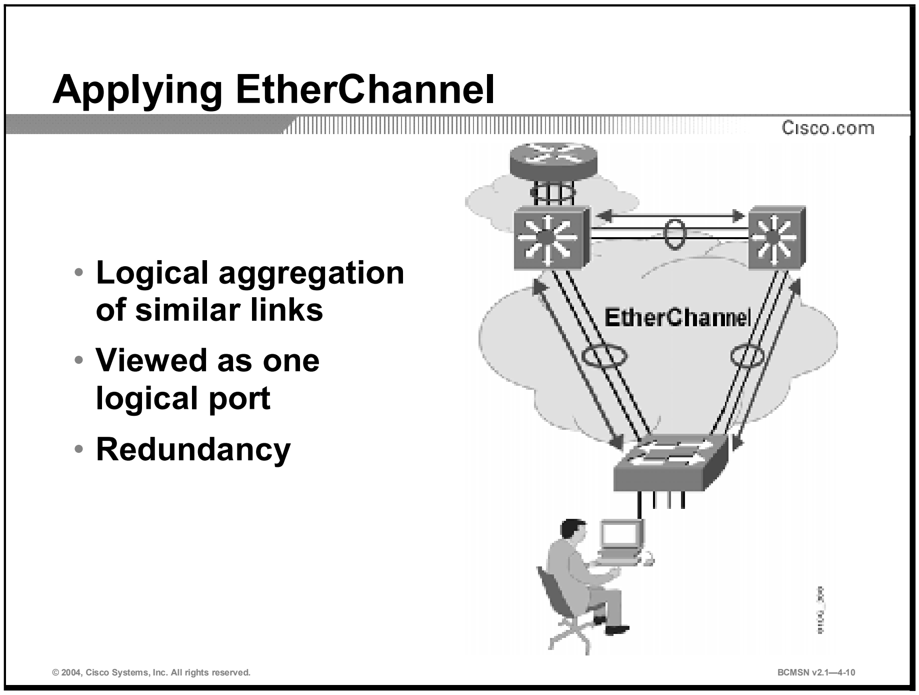

Applying EtherChannel

Cisco.com

- Logical aggregation of similar links
- Viewed as one logical port
- Redundancy

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-10

EtherChannel bundles individual Ethernet links into a single logical link that provides bandwidth up to 1600 Mbps (Fast EtherChannel full duplex) or 16 Gbps (Gigabit EtherChannel) between two Catalyst switches. All interfaces in each EtherChannel must be the same speed and must all be configured as either Layer 2 or Layer 3 interfaces.

If a segment within an EtherChannel fails, traffic previously carried over the failed link switches to the remaining segments within the EtherChannel. When the segment fails, a Simple Network Management Protocol (SNMP) trap is sent, identifying the switch, the EtherChannel, and the failed link. Inbound broadcast and multicast packets on one segment in an EtherChannel are blocked from returning on any other segment of the EtherChannel.

Each EtherChannel has a numbered port channel interface. A configuration applied to the port channel interface affects all physical interfaces assigned to that interface.

After you configure an EtherChannel, the configuration you apply to the port channel interface affects the EtherChannel. The configuration you apply to the physical interface affects only the interface where you apply the configuration. To change the parameters of all ports in an EtherChannel, apply configuration commands to the port channel interface. (Such commands can be STP commands or commands to configure a Layer 2 EtherChannel as a trunk.)

Use the option that provides the greatest variety in your configuration. If the traffic on a channel is going to a MAC address, the switch always uses the same link in the channel. Using source addresses or IP addresses might result in better load balancing.

Note	Load balancing operates at the switch level rather than at the per-channel level, and it is applied globally for all channels on the switch.

Comparing Port and Link Aggregation Protocols

This topic discusses the features the Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP).

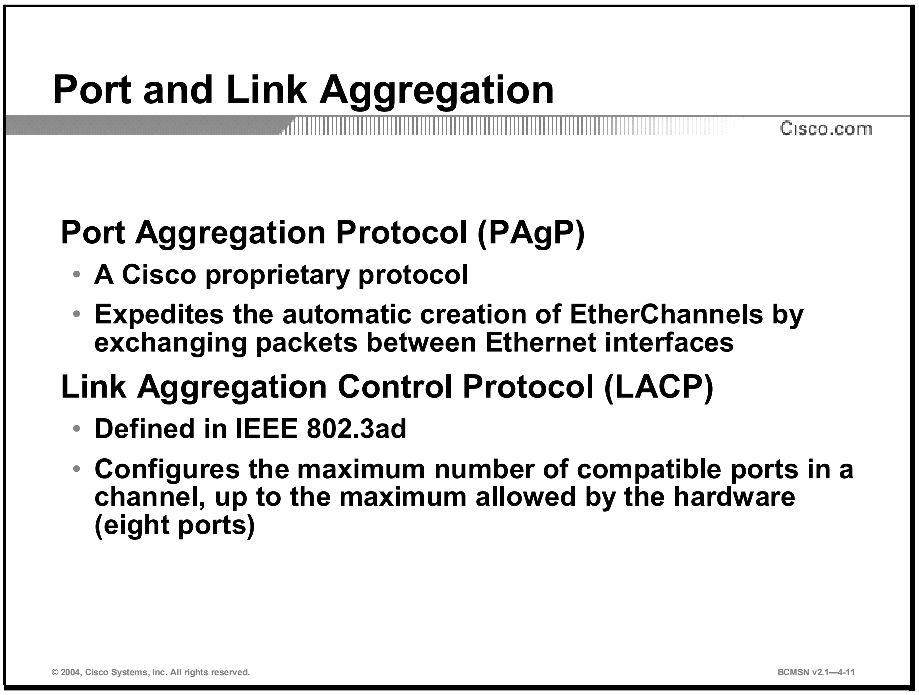

Port and Link Aggregation

Cisco.com

Port Aggregation Protocol (PAgP)

- **A Cisco proprietary protocol**
- **Expedites the automatic creation of EtherChannels by exchanging packets between Ethernet interfaces**

Link Aggregation Control Protocol (LACP)

- **Defined in IEEE 802.3ad**
- **Configures the maximum number of compatible ports in a channel, up to the maximum allowed by the hardware (eight ports)**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-11

PAgP and LACP are two different protocols that allow ports with similar characteristics to form a channel through dynamic negotiation with adjoining switches.

- PAgP is a Cisco proprietary protocol that can be run only on Cisco switches that are released by licensed vendors.
- LACP, which is defined in IEEE 802.3ad, allows Cisco switches to manage Ethernet channeling with devices that conform to the 802.3ad specification.

Note PAgP expedites the automatic creation of EtherChannels by exchanging packets between Ethernet interfaces.

PAgP learns the capabilities of interface groups dynamically and informs the other interfaces. When PAgP identifies correctly matched Ethernet links, PAgP groups the links into an EtherChannel. The EtherChannel is then added to the spanning tree as a single bridge port.

EtherChannel includes the following three user-configurable modes: on, auto, and desirable. Only auto and desirable are PAgP modes.

- **On:** This is the mode that forces the interface to channel without PAgP.
- **Auto:** This PAgP mode places an interface in a passive negotiating state in which the interface responds to the PAgP packets it receives, but it does not initiate PAgP negotiation.
- **Desirable:** This PAgP mode places an interface in an active negotiating state in which the interface initiates negotiations with other interfaces by sending PAgP packets.

Interfaces configured in the on mode do not exchange PAgP packets. The default mode for PAgP is auto mode.

To form an EtherChannel, interfaces use both the auto and desirable modes when negotiating with partner interfaces.

The EtherChannel modes that use LACP are as follows:

- **On:** This mode forces the port to channel without LACP.
- **Off:** This mode prevents the port from channeling.
- **Passive:** This LACP mode places a port into a passive negotiating state. In this state, the port responds to the LACP packets that it receives, but it does not initiate LACP packet negotiation (default).
- **Active:** This LACP mode places a port into an active negotiating state. In this state, the port initiates negotiations with other ports by sending LACP packets.

The following parameters are used in configuring LACP:

- **System priority:** Each switch running LACP must have a system priority. The system priority can be specified automatically or through the command-line interface (CLI). The switch uses the MAC address and the system priority to form the system ID.
- **Port priority:** Each port in the switch must have a port priority. The port priority can be specified automatically or through the CLI. The port priority and the port number form the port identifier. The switch uses the port priority to decide which ports to put in standby mode when a hardware limitation prevents all compatible ports from aggregating.
- **Administrative key:** Each port in the switch must have an administrative key value, which can be specified automatically or through the CLI. The administrative key defines the ability of a port to aggregate with other ports, determined by the following:
  - Port physical characteristics, such as data rate, duplex capability, and point-to-point or shared medium
  - Configuration constraints that you establish

When enabled, LACP attempts to configure the maximum number of compatible ports in a channel. If LACP is not able to aggregate all the ports that are compatible (for example, the remote system might have more restrictive hardware limitations), all the ports that cannot be actively included in the channel are put in hot standby state and used only if one of the channeled ports fails.

Using STP Enhancements in the Enterprise Composite Network Model

This topic identifies how STP enhances the Enterprise Composite Network model.

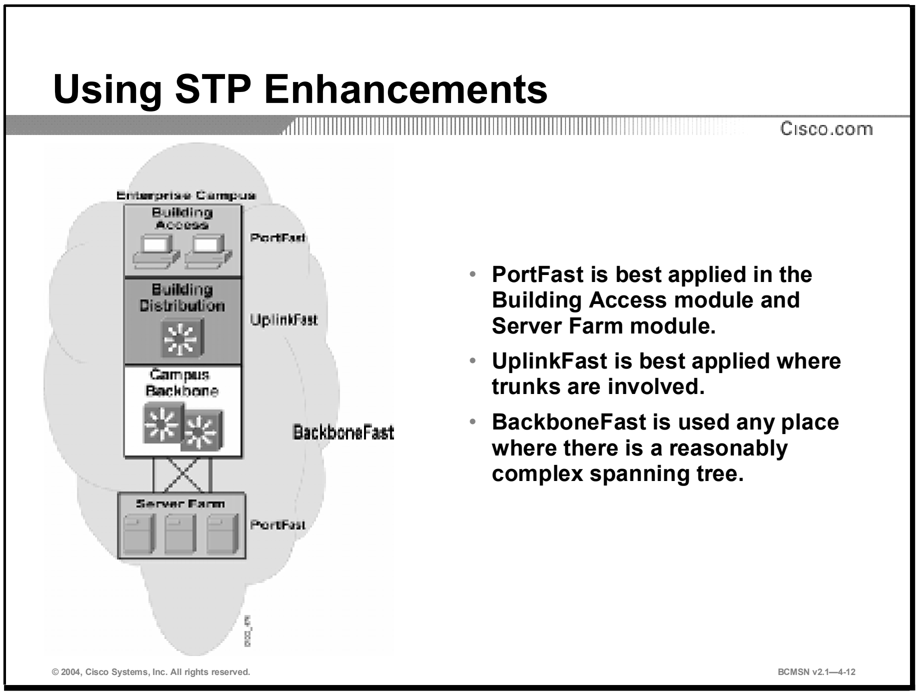

Using STP Enhancements

Cisco.com

- PortFast is best applied in the Building Access module and Server Farm module.
- UplinkFast is best applied where trunks are involved.
- BackboneFast is used any place where there is a reasonably complex spanning tree.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-12

Layer 2 and multilayer switches are used for connectivity in every module within the Enterprise Composite Network model. Layer 2 LAN networks are found throughout the Enterprise Composite Network model, depending on the scale of the network. Typically, Layer 2 LAN networks are found in these modules and submodules:

- Campus Backbone submodule
- Server Farm module
- Network Management module
- E-Commerce module
- Internet Connectivity module
- Remote Access and VPN module

If optimally designed, the Layer 2 LAN networks have a diameter of 1. This value means that there is only one Layer 2 link from the root switch to the furthest leaf node switch. Typically, the leaf node switch has a connection to both the root switch and the backup root switch. The UplinkFast protocol was specifically designed for this type of topology.

In this Layer 2 LAN network design, the access port has two routes to the root switch: A directly connected link and a link through the backup root switch.

UplinkFast is root port optimization. On the access switch, spanning tree blocks the inferior port to the root switch. UplinkFast takes advantage of the fact that this inferior route is known. When the direct link to the root fails, the indirect route to the root switch through the backup root switch is brought up immediately. The switch port for the indirect link enters into the forwarding state, without waiting for the max-age timer to expire and without entering the listening and learning states. UplinkFast should only be configured on leaf node switches.

PortFast is a spanning tree feature designed to optimize switch ports connected to user devices. User devices are turned on and off and connected or disconnected to the network. Each time the user device is connected or turned on, there is a change in the spanning tree topology and topology change notification (TCN) BPDUs are sent. PortFast not only prevents the time delay in forwarding traffic, PortFast also reduces TCNs in the network. However the benefit of spanning tree is lost if there are incorrectly configured switch ports using PortFast. Only ports that will not have end-user devices attached should be configured with PortFast.

Backbone Fast is only required if a Layer 2 LAN network has a diameter of 2 or more. In these Layer 2 LAN networks, BackboneFast would be configured on every switch.

EtherChannel is not a spanning tree feature, but it interacts with spanning tree. EtherChannel is a feature that provides link aggregation, such that the bandwidth between any two devices is increased. Etherchannel can be used on any Layer 2 or Layer 3 uplinks. When EtherChannel is implemented between the access and distribution layers on Layer 2 ports, spanning tree sees the aggregated link as one link. Thus spanning tree will not block any ports.

Summary

This topic summarizes the key points discussed in this lesson.

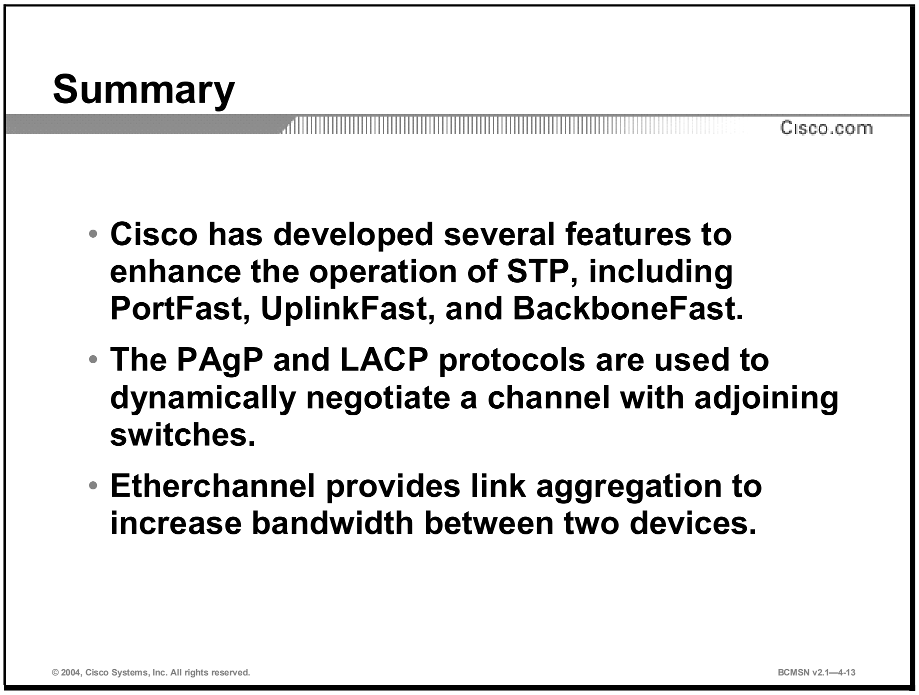

Summary

- Cisco has developed several features to enhance the operation of STP, including PortFast, UplinkFast, and BackboneFast.
- The PAgP and LACP protocols are used to dynamically negotiate a channel with adjoining switches.
- Etherchannel provides link aggregation to increase bandwidth between two devices.

BCMSN v2.1—4-13

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Select the appropriate feature that applies to PortFast.

A) enables an interface to enter the learning state immediately

B) enables an interface to enter the forwarding state immediately

C) enables an interface to enter the blocking state immediately

D) enables an interface to enter the listening state immediately

Q2) Select the appropriate feature that applies to UplinkFast.

A) provides immediate transition to the learning state

B) maintains blocked ports until it receives a BPDU telling it otherwise

C) "backup link" is put into pseudo forwarding state

D) provides fast convergence after a direct link failure

Q3) Select the appropriate feature that applies to BackboneFast.

A) blocked port processes BPDUs when max_age interval expires

B) recovers from an indirect failure in 50 seconds (default)

C) recovers from an indirect failure in 30 seconds (default)

D) recovers from a direct failure in 15 seconds (default)

Q4) Select the characteristics that apply to EtherChannel. (Choose two.)

A) The interfaces in each EtherChannel can be different speeds.

B) It bundles individual Ethernet links into a single logical link.

C) Each EtherChannel has a numbered port channel interface.

D) The configuration applied to the physical interface affects the entire EtherChannel.

Q5) Match the features with the correct aggregation protocol.

_____ 1. PAgP protocol

_____ 2. LACP protocol

A) manages Ethernet channeling with devices that conform to the 802.3ad specification

B) is the default protocol in auto mode

C) uses active and passive channel modes if you want to handle channeling

D) must have system and port priorities

E) groups links into an EtherChannel when links are correctly matched

F) handles channeling in active and passive channel modes

G) uses only auto and desirable modes

Q6) Select the methods in which STP enhances the Enterprise Composite Network model. (Choose two.)

A) Layer 2 and multilayer switches are used for connectivity in every module within the Enterprise Composite Network model.

B) There are only three Layer 2 links between the root switch and the furthest leaf node switch.

C) UplinkFast takes advantage of the fact that an inferior route is known.

D) All end-user devices can be configured with PortFast.

Quiz Answer Key

Q1) B

Relates to: Applying PortFast

Q2) D

Relates to: Applying UplinkFast

Q3) C

Relates to: Applying BackboneFast

Q4) B, C

Relates to: Applying EtherChannel

Q5) 1= B, C, D, G; 2=A, E, F

Relates to: Comparing Port and Link Aggregation Protocols

Q6) A, C

Relates to: Using STP Enhancements in the Enterprise Composite Network Model

Accelerating Spanning Tree Convergence

Overview

Rapid Spanning Tree Protocol (RSTP) is designed to significantly speed the recalculation of the spanning tree. RSTP can be seen as an evolution of the 802.1D standard, and it is the preferred protocol for preventing Layer 2 loops in a switched network environment.

Relevance

Proper configuration of STP versions and enhancements can have a significant impact on overall network performance.

Objectives

Upon completing this lesson, you will be able to:

- Select the correct features of RSTP
- Match the correct RSTP port states with their appropriate function
- Match the correct RSTP port role with the appropriate function
- Match the correct RSTP link type with the appropriate function
- Select the features that are attributed to RSTP BPDUs
- Select the correct features of rapid transition to the forwarding state
- Match the actions that apply to notifying topology changes with RSTP in their correct order
- Select the features that apply to MST
- Select the characteristics that apply to the MST region
- Select the appropriate benefits of applying RSTP in the Enterprise Composite Network model

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Improving Spanning Tree Convergence
- Identifying RSTP Port States
- Identifying RSTP Port Roles
- Identifying RSTP Link Types
- Examining RSTP BPDU
- Transitioning Rapidly to the Forwarding State
- Notifying Topology Changes with RSTP
- Extending STP to Multiple Spanning Trees
- Identifying the Characteristics of an MST Region
- Applying RSTP in the Enterprise Composite Network Model
- Summary
- Quiz

Improving Spanning Tree Convergence

This topic identifies the features of RSTP.

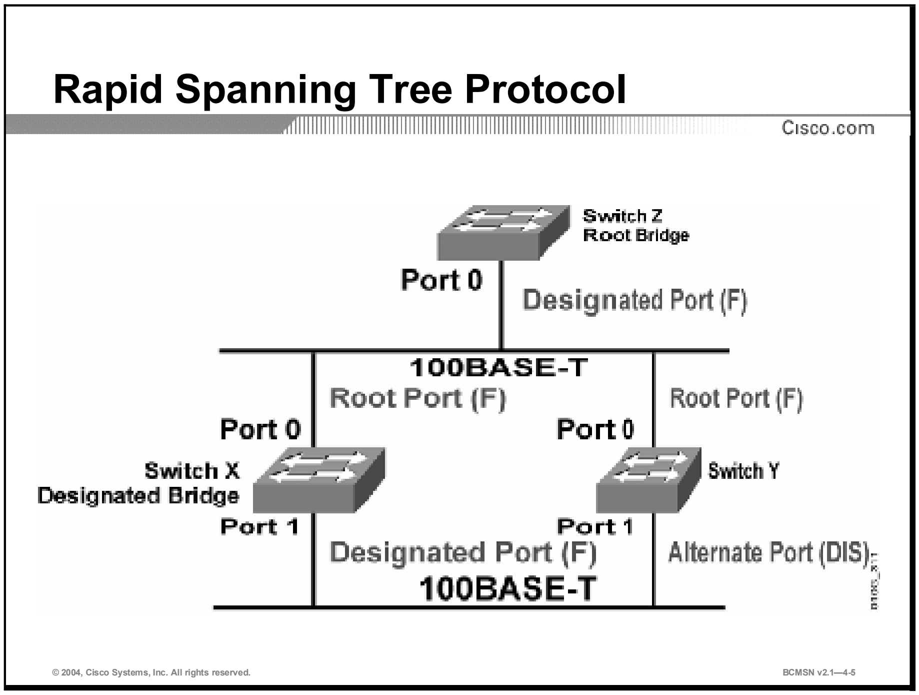

Rapid Spanning Tree Protocol

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-5

RSTP is designed to significantly speed the recalculation of the spanning tree when a Layer 2 network topology changes. RSTP is an IEEE standard that incorporates many of the concepts used in the Cisco-proprietary STP enhancements. RSTP redefines the base operation of STP, the port roles and states, and BPDUs.

RSTP is proactive instead of passive and has completely negated the need for the 802.1D delay timers. RSTP (802.1w) supersedes 802.1D, while still remaining backward compatible. The 802.1D terminology remains primarily the same; most parameters remain unchanged. In addition, 802.1w is capable of reverting back to 802.1D to interoperate with legacy switches on a per-port basis. However, in doing so, the benefits of 802.1w over 802.1D are negated.

In a switched domain, there can be only one forwarding path toward a single reference point, which is the root switch. The RSTP spanning tree algorithm (STA) elects a root switch in the switched domain and calculates the final topology for the spanning tree. This STA uses exactly the same criteria and in the same order as does 802.1D.

However, there are critical differences that make RSTP the preferred protocol for preventing Layer 2 loops in a switched network environment. Many of the differences stem from the Cisco proprietary enhancements. The Cisco-based RSTP enhancements have these characteristics:

- They are integrated into the protocol at a low level.
- They are transparent.
- They require no additional configuration.
- They generally perform better than the Cisco-proprietary STP enhancements.

Because the RSTP and the Cisco-proprietary enhancements are functionally similar, features such as UplinkFast and BackboneFast are not compatible with RSTP.

Identifying RSTP Port States

This topic discusses RSTP port states and their appropriate functions.

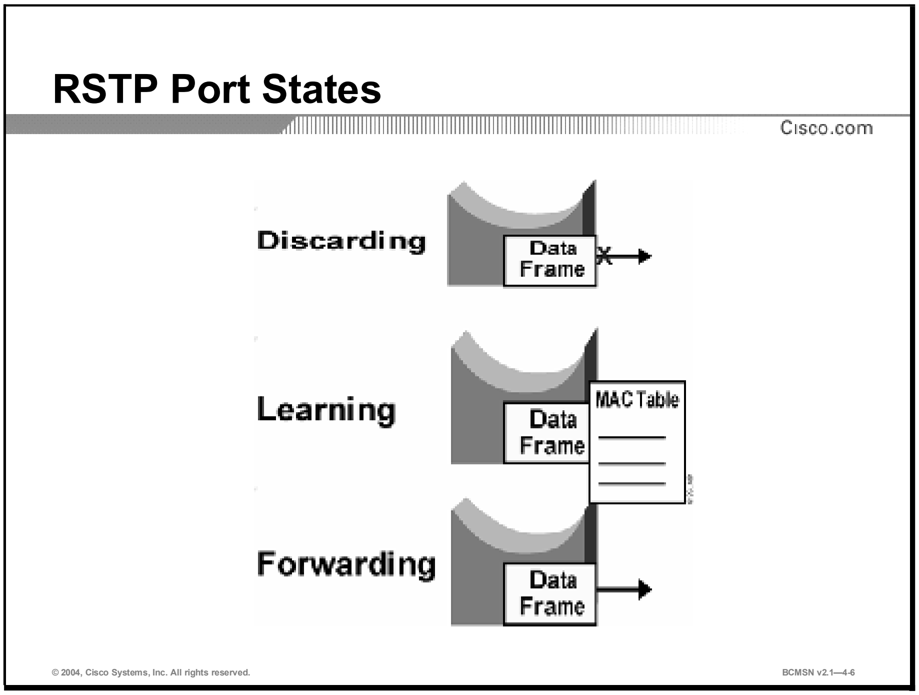

RSTP Port States

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-6

RSTP provides rapid convergence following the failure or re-establishment of a switch, switch port, or link. An RSTP topology change will cause a transition in the appropriate switch ports to the forwarding state through explicit handshakes or a proposal and agreement process and synchronization.

With RSTP, the role of a port is separated from the state of a port. For example, a designated port may be in the discarding state, even though this condition would be temporary. However, in a stable, active topology, the port role still determines the underlying final port state. For example, a designated port will still need to end up in the forwarding state.

The RTSP port states correspond to the three basic operations of a switch port: discarding, learning, and forwarding.

The port states have these characteristics:

- **Discarding:** This state is seen in both a stable active topology and during topology synchronization and changes. The discarding state prevents the forwarding of data frames from “breaking” the continuity of a Layer 2 loop.
- **Learning:** This state is seen in both a stable active topology and during topology synchronization and changes. The learning state accepts data frames to populate the MAC table in an effort to limit flooding of unknown unicast frames.
- **Forwarding:** This state is seen only in stable active topologies. The forwarding switch ports determine the topology. Following a topology change or during synchronization, the forwarding of data frames is only achieved by a proposal and agreement process.

In addition to these characteristics, the port states all accept and receive BPDU frames for processing.

Identifying RSTP Port Roles

This topic discusses the RSTP port role functions.

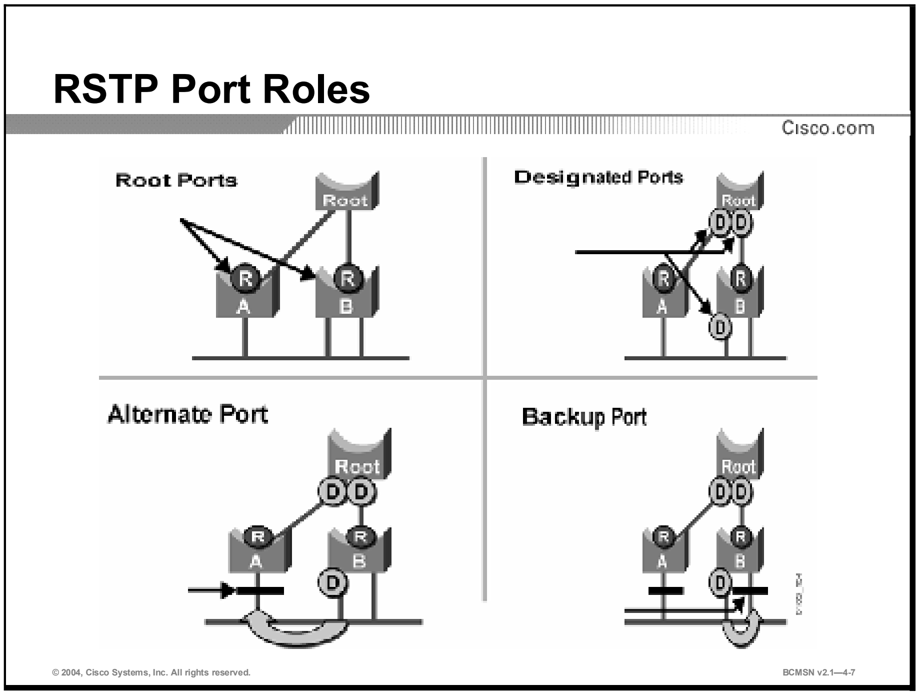

RSTP Port Roles

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-7

The port role defines the ultimate objective state of a switch port and, therefore, the handling of data frames. Yet the port role and the port state are able to transition independently of each other. RSTP uses these definitions for port roles:

- **Root port:** This is the switch port on every nonroot switch that is the closest to the root switch in terms of path cost. There can only be one root port on every switch. The root port assumes the forwarding state in a stable active topology.
- **Designated port:** This is the switch port on every switch that connects the Ethernet segments to the root port toward the root switch. There can only be one designated port per segment. The designated port assumes the forwarding state in a stable active topology. All switches connected to a given segment listen to all BPDUs and determine the switch that will be the designated switch for a particular Ethernet segment.
- **Alternate port:** This is a switch port that offers an alternate path toward the root switch. The alternate port assumes a discarding state in a stable active topology. An alternate port will be present on nondesignated switches and will make a transition to a designated port if the current designated path fails.
- **Backup port:** This is an additional switch port on the designated switch with a redundant link to the segment for which the switch is designated. A backup port has a higher port ID than the designated port on the designated switch. The backup port assumes the discarding state in a stable active topology.
- **Disabled port:** This is a port that has no role within the operation of spanning tree.

Establishing the additional port roles allows RSTP to define a standby switch port before a failure or topology change. The alternate port moves to the forwarding state if there is any failure in the path through the designated port and root port on a designated switch, including the failure of the designated switch itself.

Identifying RSTP Link Types

This topic matches the correct RSTP link type with the appropriate function.

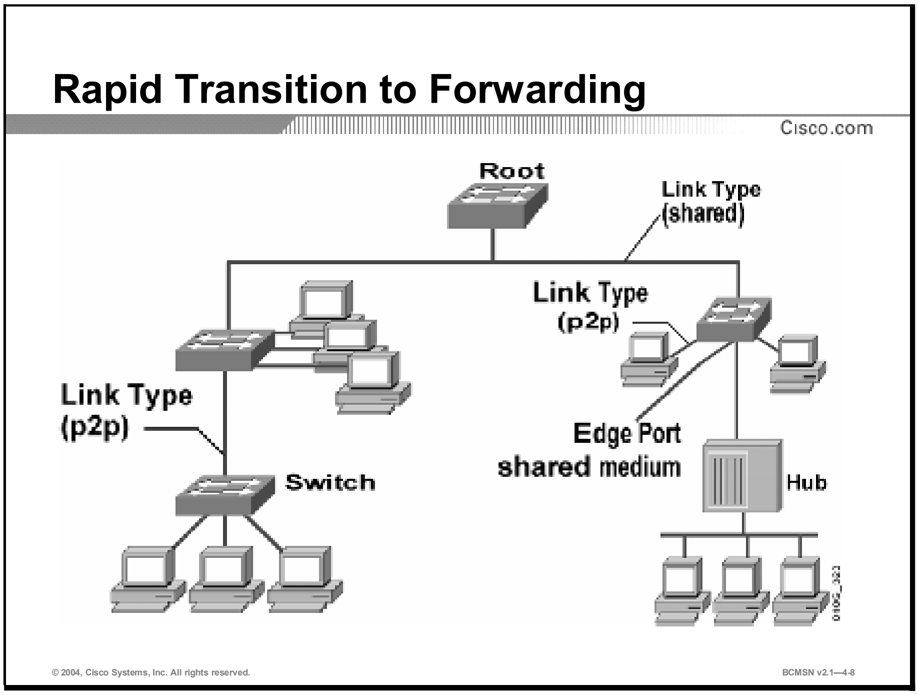

Rapid Transition to Forwarding

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-8

The IEEE 802.1w standard includes parameters that allow for the immediate transition to the forwarding state if certain conditions are met. These parameters fall into two separate and distinct groups, yet they are used to accomplish the same result, which is the transition to the forwarding state. The parameters in the one group determine if the switch port is an edge port. The parameters in the other group determine the type of link to which the switch port is connected. These parameters are configurable.

An edge port is defined as a port that is not connected to another bridge device. If the edge port parameter is “true,” the port is immediately transitioned to the forwarding state. If an edge port receives a topology change notification BDPU, the parameter is immediately changed to “false,” a TCN is sent, and the spanning tree algorithm runs.

The link type parameter is either point-to-point or shared. The value of the parameter can be configured or automatically determined. A port operating in full-duplex mode is point-to-point, while a port operating in half-duplex mode is considered to be on a shared medium by default. You can override the automatic link type setting with an explicit configuration.

Before the link type parameter is used, the spanning tree algorithm must determine the port role.

Root ports do not use the link type parameter. Root ports are able to make a rapid transition to the forwarding state as soon as the port is in “sync.” The dependency for a root port to sync and directly move from the discarding state to the forwarding state is the explicit handshakes that occur between the root port and the upstream designated port.

In most cases, alternate and backup ports do not use the link type parameter RSTP keeps track of ports that provide alternative paths to the root bridge. If a root port fails, RSTP can quickly make an alternative port the new root port. This new root port is placed in the forwarding state without delay.

The link type parameter is most important for the designated port. Rapid transition to the forwarding state for the designated Port occurs only if the link type parameter indicates a point-to-point link.

Examining RSTP BPDU

This topic identifies the features of RSTP BPDUs.

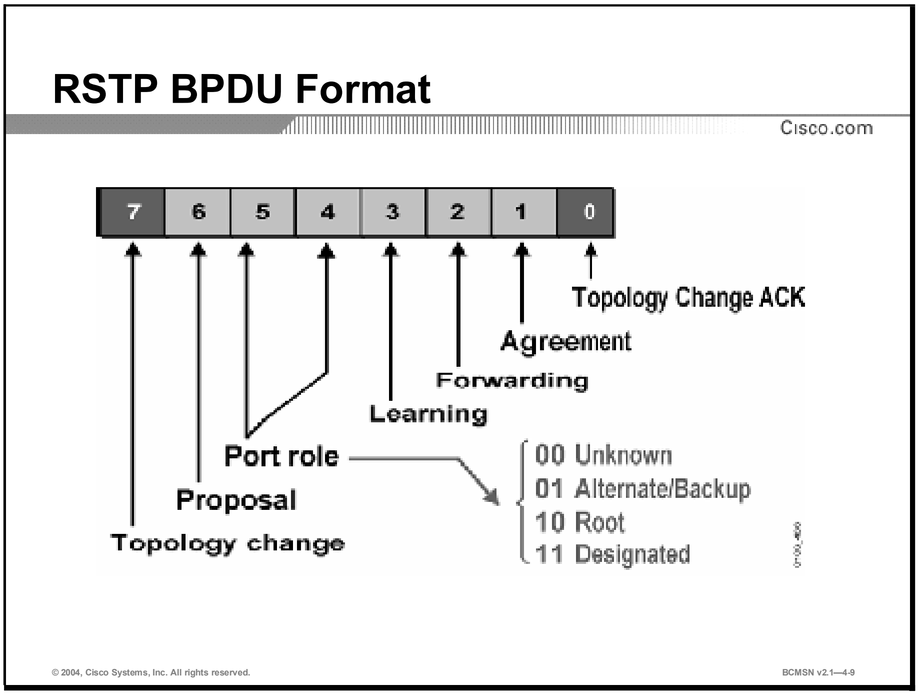

RSTP BPDU Format

Cisco.com

7 6 5 4 3 2 1 0

Topology Change ACK

Agreement

Forwarding

Learning

Port role

- 00 Unknown
- 01 Alternate/Backup
- 10 Root
- 11 Designated

Proposal

Topology change

BCMSN v2.1—4-9

RSTP has these features:

- RSTP uses all 6 remaining bits of the flag byte to do the following:
  - Encode the role and state of the port originating the BPDU
  - Handle the proposal and agreement mechanism

RSTP BPDUs are defined as type 2, version 2 so an 802.1w bridge can detect legacy bridges.

- BPDUs are now sent every hello time.
- With RSTP, a bridge now sends a BPDU with its current information every hello time period (2 seconds by default), even if it does not receive any BPDUs from the root bridge.

Protocol information can be immediately aged on a port if hellos are not received for three consecutive hello times or if the max-age timer expires. Because BPDUs are now used as a "keepalive" mechanism, three consecutively missed BPDUs indicate lost connectivity between a bridge and its neighboring root or designated bridge. This fast aging of the information allows quick failure detection.

Transitioning Rapidly to the Forwarding State

This topic identifies the features of transitioning rapidly to the forwarding state.

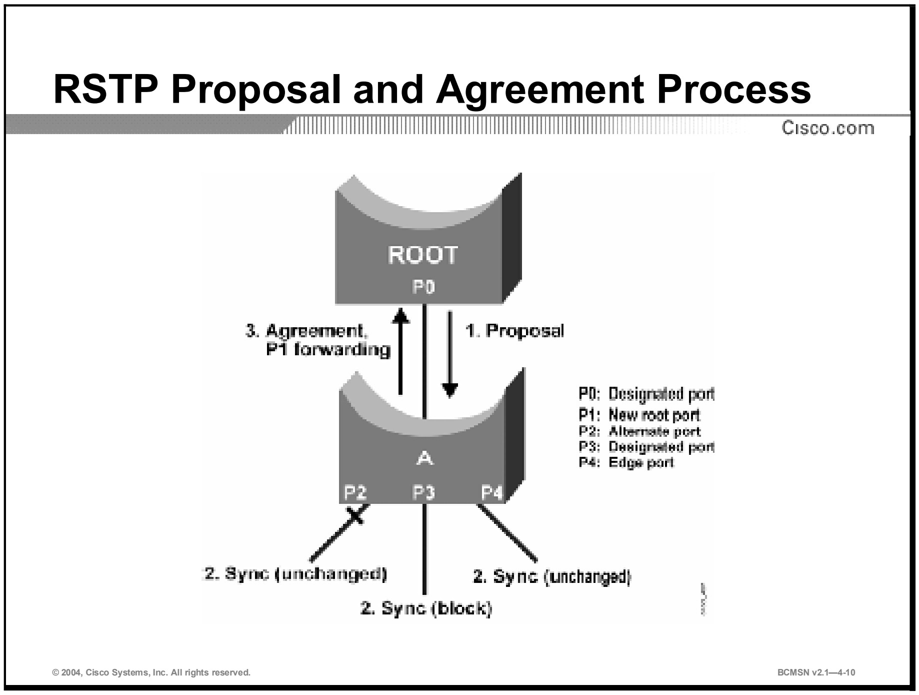

RSTP Proposal and Agreement Process

Cisco.com

BCMSN v2.1—4-10

RSTP significantly speeds up the recalculation process after a topology change occurs in the network. RSTP works by designating an alternate port and a back up port. These ports are allowed to immediately enter the forwarding state rather than passively waiting for the network to converge.

BPDUs are transmitted out all nonedge switch ports. If hellos are not received for three consecutive hello time periods, RSTP protocol information will be immediately aged out and the switch port is out of sync with the rest of the switched domain. Because the BPDUs form neighbor relationships, the switch that is not receiving BPDUs for the three hello time periods can positively identify which connection is lost and that switch port is "marked" as requiring synchronization when the link comes back up. To avoid Layer 2 loops when the link is reinstated, the switch port will default to the designated port and discarding state without any current protocol configuration information.

Note A physical link loss is detected in less than three hello time periods.

The RSTP BPDU is now used to implement the process of an explicit handshake, which quickly propagates through a switched domain and ensures a stable loop-free topology during convergence.

Example: RSTP Proposal and Agreement Process

In the example, the root switch is sending the proposal. Switch A has its P1 switch port in the designated discarding state and is transmitting a proposal BPDU. The BPDU of the root switch is superior to that of switch A. Therefore, switch A immediately moves to discard on all nonedge ports to prevent a Layer 2 loop.

Switch A responds to the root switch with an agreement to forward. Both switch ports rapidly move to the forwarding state with the assurance that there will be no Layer 2 loops.

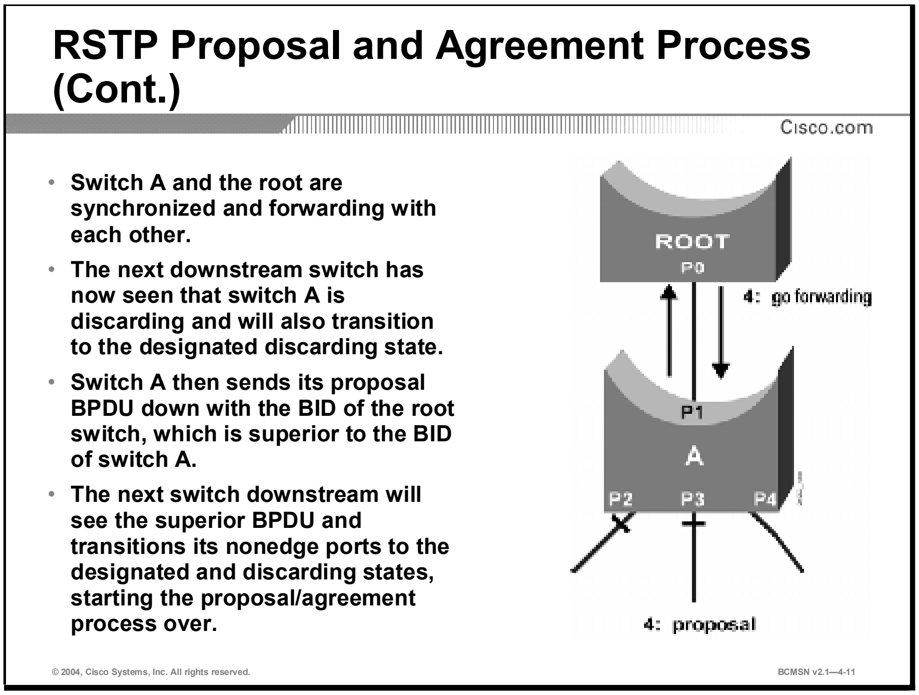

RSTP Proposal and Agreement Process (Cont.)

Cisco.com

- **Switch A and the root are synchronized and forwarding with each other.**
- **The next downstream switch has now seen that switch A is discarding and will also transition to the designated discarding state.**
- **Switch A then sends its proposal BPDU down with the BID of the root switch, which is superior to the BID of switch A.**
- **The next switch downstream will see the superior BPDU and transitions its nonedge ports to the designated and discarding states, starting the proposal/agreement process over.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-11

Notifying Topology Changes with RSTP

This topic identifies the actions that apply to notifying topology changes (TCs) with RSTP in the correct order.

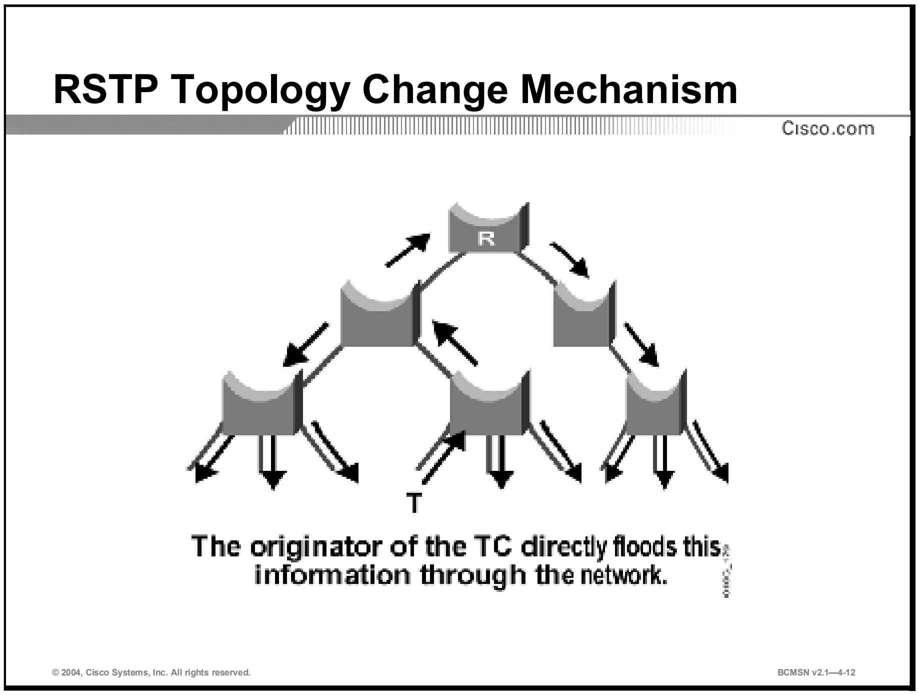

RSTP Topology Change Mechanism

Cisco.com

The originator of the TC directly floods this information through the network.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-12

When an 802.1D bridge detects a topology change, the bridge first notifies the root bridge. After the root bridge is aware of a change in the topology of the network, the root bridge sets the TC flag on the outbound BPDUs. Those BPDUs are then relayed to all the bridges in the network. When a bridge receives a BPDU with the TC flag bit set, the bridge reduces its bridging-table aging time to forward delay seconds. This ensures a relatively quick flushing of noncurrent information.

In RSTP, only nonedge ports moving to the forwarding state cause a topology change. Loss of connectivity is not considered to be a topology change, and, under these conditions, a port moving to the blocking state does not generate a TC BDPU.

When an RSTP bridge detects a TC, it performs these actions:

Step	Action	Notes
1.	The RSTP bridge starts the TC-While timer.	RTSP sets the While timer with a value equal to twice the hello time for all its nonedge designated ports and its root port, if necessary.
2.	The RSTP bridge flushes the MAC addresses associated with all these ports.	—
3.	The TC bit set is set on all outbound BPDUs.	BPDUs are sent on the root port as long as the While timer is active.
4.	The bridge receives a BPDU with the TC bit set from a neighbor and clears the MAC addresses on its ports.	The port that received the TC BPDU retains its learned MAC address.
5.	The bridge starts the TC-While timer and sends BPDUs with a TC bit set out of all its designated ports and root port.	RSTP does not use the specific TCN BPDU unless a legacy bridge needs to be notified.

The topology change notification (TCN) is flooded across the entire network. The topology change propagation is now a one-step process. There is no need for each switch port to wait for the root bridge to be notified and then maintain the TC state for the value of the max_age plus forward delay seconds.

If the port consistently keeps receiving BPDUs that do not correspond to its current operating mode for two periods of hello time, the port switches to the mode indicated by the BPDUs.

Extending STP to Multiple Spanning Trees

This topic identifies the features that apply to Multiple Spanning Tree (MST).

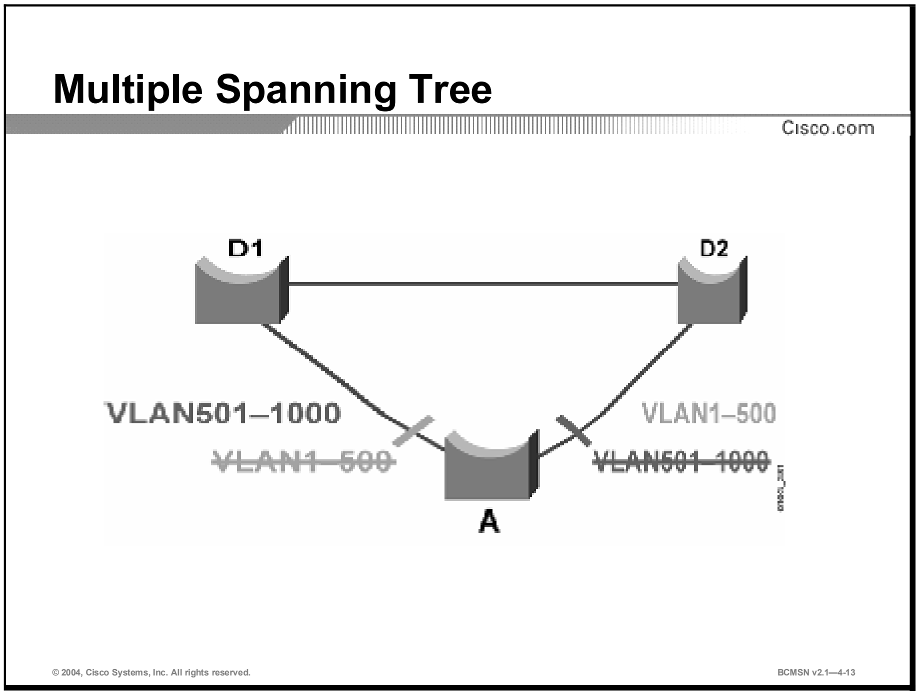

Multiple Spanning Tree

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-13

The main purpose of MST is to reduce the total number of spanning tree instances to match the physical topology of the network and thus reduce the CPU cycles of a switch. The instances of spanning tree are reduced to the number of links that are available, unlike PVST.

Example: Extending STP to Multiple Spanning Trees

In this example, there are two links and 1000 VLANs. The 1000 VLANs are mapped to two MST instances (MSTI). Rather than maintaining 1000 spanning trees, each switch needs to maintain only two. This reduces the need for switch resources. If there were 2000 VLANs, they would be reduced to two spanning tree instances also. MST converges faster than Per VLAN Spanning Tree+ (PVST+) and is backward-compatible with 802.1D STP, 802.1w (RSTP), and the Cisco PVST+ architecture.

MST allows you to build multiple spanning trees over trunks by grouping and associating VLANs to spanning tree instances. Each instance can have a topology independent of other spanning tree instances. This architecture provides multiple forwarding paths for data traffic and enables load balancing. Network fault tolerance is improved because a failure in one instance (forwarding path) does not affect other instances.

In large networks, you can more easily administer the network and use redundant paths by locating different VLAN and spanning tree assignments in different parts of the network. A spanning tree instance can exist only on bridges that have compatible VLAN instance assignments. You must configure a set of bridges with the same MST configuration information, which allows them to participate in a specific set of spanning tree instances. Interconnected bridges that have the same MST configuration are referred to as an “MST region.”

The example shows a very common design featuring access switch A with 1000 VLANs redundantly connected to two Building Distribution switches, D1 and D2. In this setup, users connect to switch A. The network administrator typically seeks to achieve load balancing on the access switch uplinks based on even or odd VLANs, or any other scheme deemed appropriate.

In a Cisco PVST+ environment, the spanning tree parameters are tuned so that half of the VLANs are forwarding on each uplink trunk. This is easily achieved by electing bridge D1 to be the root for VLAN501–1000, and bridge D2 to be the root for VLAN1–500. In this configuration, the following is true:

- Optimum load balancing results.
- One spanning tree instance for each VLAN is maintained, which means 1000 instances for only two different logical topologies. This consumes resources for all the switches in the network (in addition to the bandwidth used by each instance sending its own BPDUs).

MST can be used to achieve the same purpose with less resource and bandwidth requirements. In MST, two instances are configured, each assigned a root switch. Half the VLANs mapped to one instance and the other half mapped to the second instance. With only two instead of 1000 instances running, less resources and bandwidth are consumed in the network.

Identifying the Characteristics of an MST Region

This topic identifies the characteristics that apply to the MST region.

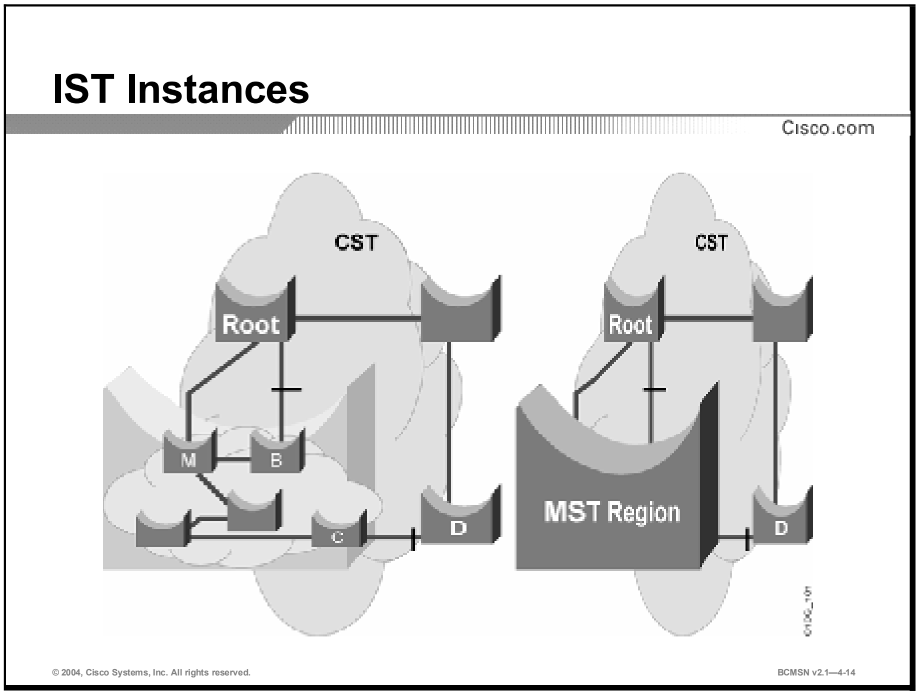

IST Instances

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-14

According to the IEEE 802.1s specification, an MST switch must be able to handle at least one Internal Spanning Tree (IST). The IST instance receives and sends BPDUs to the Common Spanning Tree (CST). The IST is capable of representing the entire MST region as a CST virtual bridge to switched networks outside the MST region.

- The MST region appears as a single virtual bridge to the adjacent CST and MST regions. The MST region uses 802.1w port roles and operation.
- MST switches run IST, which augments CST information with internal information about the MST region.
- IST connects all the MST switches in the region and any CST switched domain.
- MST establishes and maintains additional spanning trees within each MST region. These spanning trees are termed MST instances (MSTIs). The IST is numbered 0, and the MSTIs are numbered 1, 2, 3, and so on to 15. Any MSTI is local to the MST region and is independent of MSTIs in another region, even if the MST regions are interconnected. MST instances combine with the IST at the boundary of MST regions to become the CST as follows:
  - — M-records are always encapsulated within MST BPDUs. The original spanning trees are called “M-trees,” which are active only within the MST region. M-trees merge with the IST at the boundary of the MST region and form the CST.

- MST supports some of the PVST+ extensions as follows:
  - UplinkFast and BackboneFast are not available in MST mode; they are part of RSTP.
  - PortFast is supported.
  - BPDU filter and BPDU guard are supported in MST mode.
  - Loop guard and root guard are supported in MST.
  - MST switches operate as if MAC address reduction is enabled. MAC address reduction is used to enable extended-range VLAN identification.
  - For PVLANs, you must map a secondary VLAN to the same instance as the primary.

Example: Identifying the Characteristics of an IST Instance

The example shows two functionally equivalent diagrams.

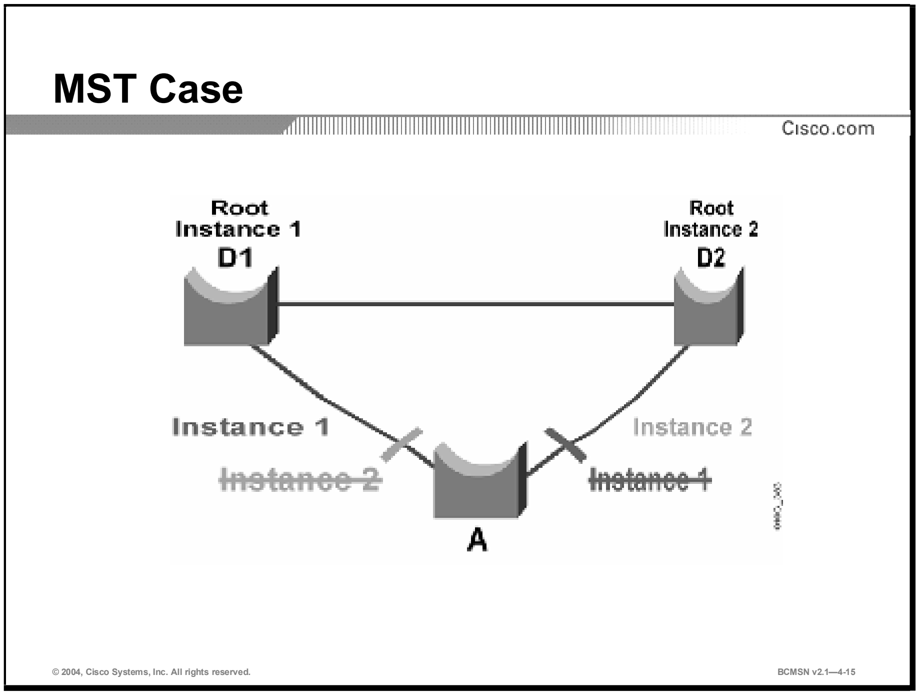

MST Case

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-15

Notice the location of the different blocked ports. In a typically switched network, you would expect to see a blocked port between switches M and B. This is because the other ports are root ports and designated ports. Instead of blocking on D, you would expect to have the second loop broken by a blocked port somewhere in the middle of the MST region. However, because of the IST, the entire region appears as one virtual bridge that runs a CST. This makes it possible to understand that the virtual bridge blocks an alternate port on B. Also, that virtual switch is on the C to D segment, thus leading switch D to block its port.

Applying RSTP in the Enterprise Composite Network Model

This topic identifies the benefits of applying RSTP in the Enterprise Composite Network model.

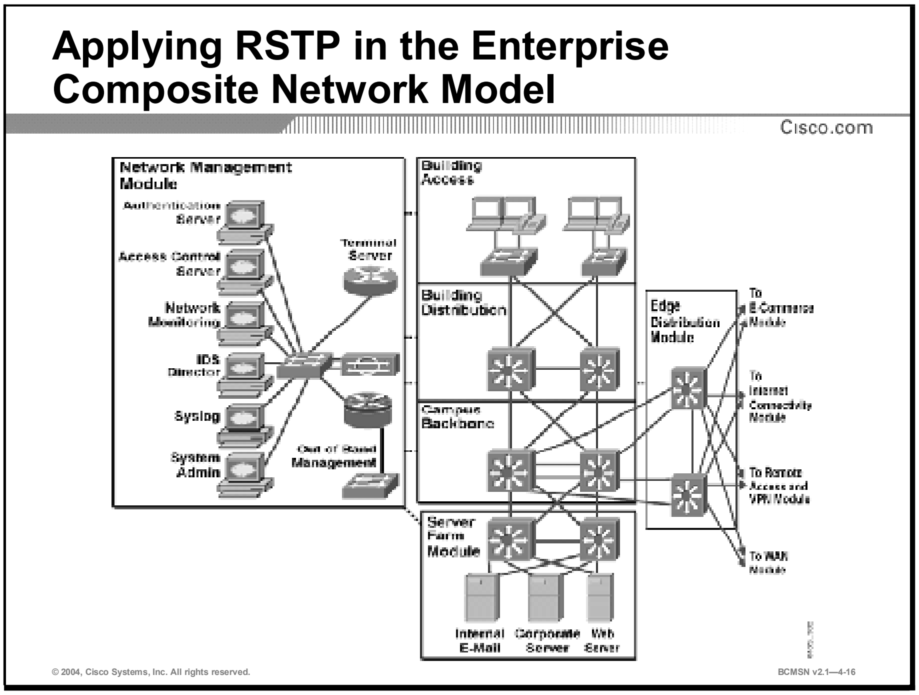

Applying RSTP in the Enterprise Composite Network Model

Cisco.com

BCMSN v2.1—4-16

The benefits of applying RSTP in the Enterprise Composite Network model include the following:

- RSTP provides an improved mode of bridge operation while still retaining the plug-and-play benefits of 802.1D STP. RSTP discards the significant time taken for 802.1D STP to restore service after a link failure.
- RSTP improves the operation of STP while maintaining backward compatibility with the following:
  - 802.1D STP
  - Cisco PVST+
  - Cisco-proprietary Multi-Instance STP (MISTP)
- RSTP natively includes most of Cisco proprietary enhancements to the 802.1D spanning tree, such as BackboneFast, UplinkFast, and PortFast.
- RSTP can achieve much faster convergence in a properly configured network, sometimes in the order of a few hundred milliseconds.
- Classic 802.1D timers, such as forward delay and max_age, are almost only used as a backup. These times should not be necessary if point-to-point links and edge ports are properly set by the administrator, and if there is no interaction with legacy 802.1D bridges.
- In a stable topology with consistent port roles throughout the network, RSTP ensures that every root port and designated port immediately transition to the forwarding state, while all alternate and backup ports are in the discarding state (equivalent to the blocking state in 802.1D).

Summary

This topic summarizes the key points discussed in this lesson.

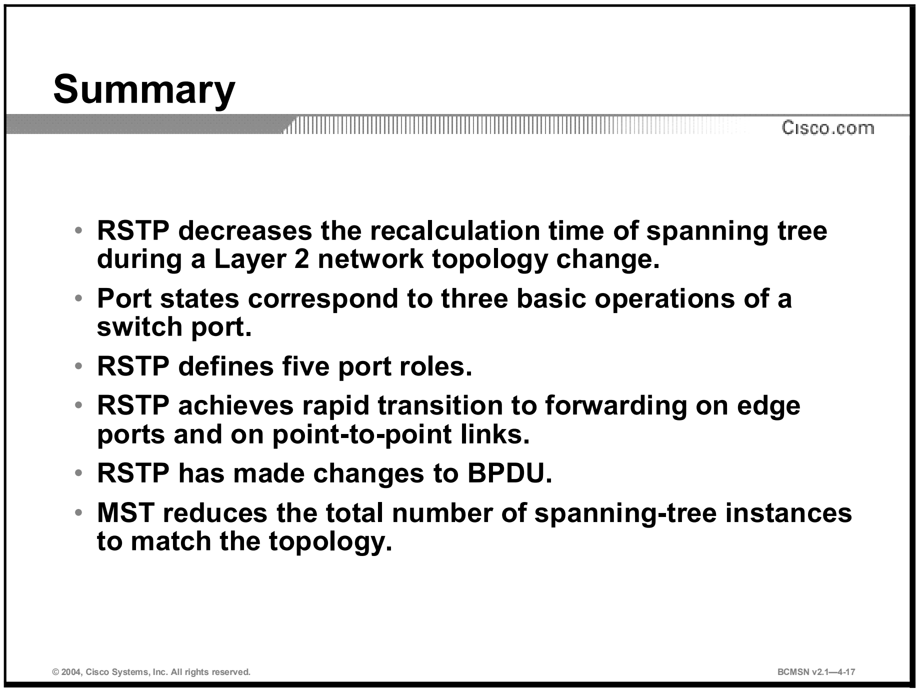

Summary

- RSTP decreases the recalculation time of spanning tree during a Layer 2 network topology change.
- Port states correspond to three basic operations of a switch port.
- RSTP defines five port roles.
- RSTP achieves rapid transition to forwarding on edge ports and on point-to-point links.
- RSTP has made changes to BPDU.
- MST reduces the total number of spanning-tree instances to match the topology.

BCMSN v2.1—4-17

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Select the correct features of RSTP. (Choose two.)

A) It speeds up the recalculation of the spanning tree when there are any network topology changes.

B) It works smoothly with UplinkFast, PortFast, and BackboneFast.

C) It is backward compatible.

D) It has negated the need for 802.1D delay timers.

Q2) Match the correct port state with the appropriate function.

_____ 1. discarding

_____ 2. learning

_____ 3. forwarding

A) It accepts data frames to populate the MAC table.

B) It prevents the forwarding of data frames to break the continuity of a Layer 2 loop.

C) The forwarding of frames is achieved by a proposal and agreement process.

Q3) Match the correct port role with the appropriate function.

_____ 1. root port

_____ 2. disabled port

_____ 3. designated port

_____ 4. non-designated port

_____ 5. alternate path port

A) has no role in spanning tree

B) is not active in the spanning tree

C) has the least path cost to the root switch and is not a root port

D) blocks data frames and does not populate the MAC table on any given segment

E) exists on nonroot switches and is the switch port with the least path cost to the root switch

Q4) Match the correct parameter with its appropriate function.

_____ 1. point-to-point link type

_____ 2. shared link type

_____ 3. edgeport parameter

A) is defined as a port that is not connected to another bridge device

B) operates in full duplex mode

C) does not generate a topology change notification BPDU on any change

D) operates in half duplex mode

Q5) Select the features that are attributed to RSTP BPDUs. (Choose two.)

A) They are type 2, version 3.

B) BPDUs are sent every hello time.

C) Protocol information is immediately discarded if hellos are not received for four consecutive hello times.

D) RSTP BPDUs enable bridges to detect a failure quickly.

Q6) Which statement best describes a feature of the rapid transition to the forwarding state. (Choose two.)

A) It results from the process of an explicit handshake that propagates fast through a switched domain.

B) If hellos are not received at a switch port for four consecutive hello times, RSTP protocol ages out.

C) Switch ports default to the forwarding state to avoid any Layer 2 loops.

D) The designated port in a discarding state sets the proposal bit on the BPDU that it sends out.

Q7) Match the actions that apply to notifying topology changes with RSTP with their correct order.

_____ 1. Step 1

_____ 2. Step 2

_____ 3. Step 3

_____ 4. Step 4

_____ 5. Step 5

A) The bridge clears MAC addresses learned on all its ports, except the one that received the topology change.

B) The RSTP bridge flushes the MAC addresses associated with all the ports.

C) The bridge starts the TC-while timer and sends BPDUs with TC set on all its designated ports.

D) The RSTP bridge starts the TC-while timer with a value equal to twice the hello time for all its nonedge designated ports.

E) While the TC-while timer is running on a port, the BPDUs sent out of that port have the TC bit set.

Q8) Select the features that apply to multiple spanning tree. (Choose two.)

A) It can group and associate VLANs to spanning tree instances.

B) It prevents redundant paths.

C) A failure in one instance can affect a failure in another instance.

D) It reduces the total number of spanning tree instances to match the physical topology of the network.

Q9) Select the characteristics that apply to the MST region. (Choose two.)

A) uses 802.1D port roles

B) establishes and maintains additional spanning trees within each MST region

C) supports PortFast, UplinkFast, and BackboneFast

D) supports loop guard and root guard

Q10) Select the benefits of applying RSTP in the Enterprise Composite Network model. (Choose two.)

A) RSTP natively includes most of Cisco proprietary enhancements to the 802.1D spanning tree such as BackboneFast, UplinkFast, and PortFast.

B) RSTP can achieve fast convergence in a properly configured network, typically in the order of 30 to 50 seconds.

C) With the PVST+ and MISTP, RSTP improves the operation of the Spanning Tree Protocol while maintaining backward compatibility with the 802.1D Spanning Tree Protocol.

D) In a stable topology with consistent port roles throughout the network, RSTP ensures that every root port and designated port immediately transition to the forwarding state while all alternate and backup ports are in the learning state.

Quiz Answer Key

Q1) C, D

Relates to: Improving Spanning Tree Convergence

Q2) 1=B, 2=A, 3=C

Relates to: Identifying RSTP Port States

Q3) 1=E, 2=B, 3=C, 4=D, 5=A

Relates to: Identifying RSTP Port Roles

Q4) 1=B, 2=D, 3=A, C

Relates to: Identifying RSTP Link Types

Q5) B, D

Relates to: Examining RSTP BPDU

Q6) A, D

Relates to: Transitioning Rapidly to the Forwarding State

Q7) 1=D, 2=B, 3=E, 4=A, 5=C

Relates to: Notifying Topology Changes with RSTP

Q8) A, D

Relates to: Extending STP to Multiple Spanning Trees

Q9) B, D

Relates to: Identify the Characteristics of an MST Region

Q10) A, C

Relates to: Applying RSTP in the Enterprise Composite Network Model

Configuring Spanning Tree Enhancements

Overview

To enhance the scalability of the spanning tree, UplinkFast also reduces TCN messages. EtherChannel balances traffic load across the links in a channel. EtherChannel does so by reducing part of the binary pattern formed from the addresses in the frame to a numerical value that selects one of the links in the channel.

Relevance

Spanning tree is a critical feature of a multilayer switched network, and the configuration of spanning tree can have a huge impact on network performance. Understanding the proper configuration methods helps ensure that your network operates at full efficiency.

Objectives

Upon completing this lesson, you will be able to:

- Identify the command to configure port-level tuning with PortFast
- Identify the command to configure UplinkFast
- Identify the command to configure BackboneFast
- Identify the command to enable MST
- Identify the command to verify MST
- Select the appropriate guidelines that apply to configuring EtherChannel
- Identify the command to configure EtherChannel
- Identify the command to verify EtherChannel
- Identify the command to configure PAgP and LACP
- Identify the command to verify PAgP and LACP
- Identify the command to configure EtherChannel load balancing

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Configuring Port-Level Tuning with PortFast
- Configuring UplinkFast
- Configuring BackboneFast
- Enabling Multiple Spanning Tree
- Verifying Multiple Spanning Tree
- Applying Guidelines to Configuring EtherChannel
- Configuring EtherChannel
- Verifying EtherChannel
- Configuring PAgP and LACP
- Verifying PAgP and LACP
- Balancing Ethernet Traffic Load
- Summary
- Quiz

Configuring Port-Level Tuning with PortFast

This topic identifies the commands used to configure port-level tuning with PortFast.

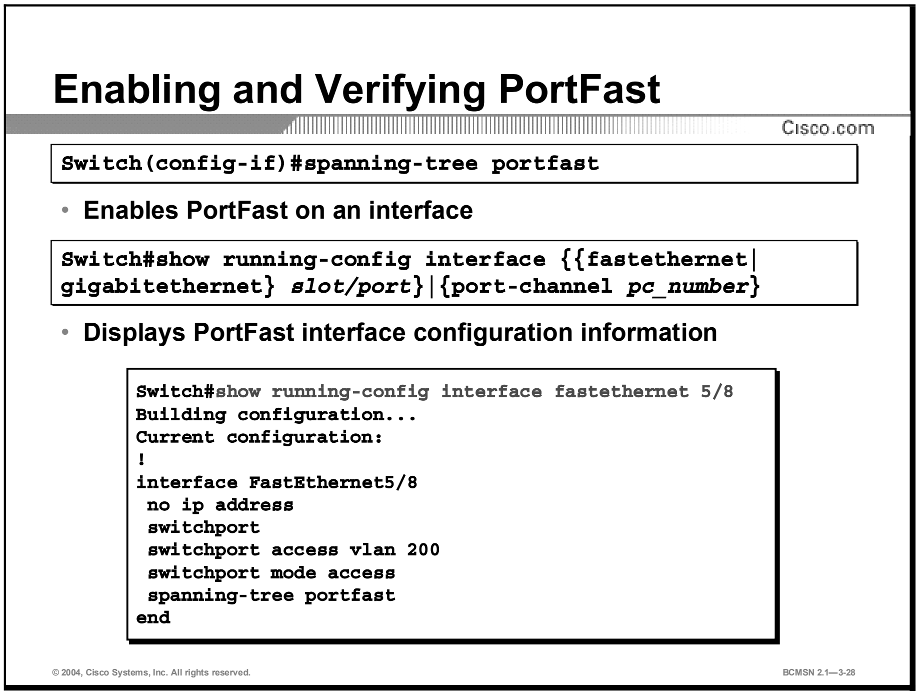

Enabling and Verifying PortFast

Cisco.com

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

- **Enables PortFast on an interface**

```
Switch#show running-config interface {{fastethernet|
gigabitethernet} slot/port}|{port-channel pc_number}
```

- **Displays PortFast interface configuration information**

```
Switch#show running-config interface fastethernet 5/8
Building configuration...
Current configuration:
!
interface FastEthernet5/8
 no ip address
 switchport
 switchport access vlan 200
 switchport mode access
 spanning-tree portfast
end
```

BCMSN 2.1—3-28

Spanning tree PortFast causes an interface configured as a Layer 2 access port to enter the forwarding state immediately, bypassing the listening and learning states. You enable PortFast on access ports connected to an end device.

To enable PortFast on a Layer 2 access port to force it to enter the forwarding state immediately, or to disable PortFast, use this command:

```
Switch(config-if)#[no] spanning-tree portfast
```

Caution Use PortFast when connecting a single end station to a Layer 2 access port only. Otherwise, you might create a network loop.

To display interface information, including PortFast configuration, use this command:

```
Switch#show running-config interface {{fastethernet |
gigabitethernet} slot/port} | {port-channel
port_channel_number}
```

Example: Enabling and Verifying PortFast

This example shows how to enable PortFast on Fast Ethernet interface 5/8:

```
Switch(config)#interface fastethernet 5/8
Switch(config-if)#spanning-tree portfast
```

This example shows how to verify the configuration.

```
Switch#show running-config interface fastethernet 5/8
Building configuration...

Current configuration:
!
interface FastEthernet5/8
no ip address
switchport
switchport access vlan 200
switchport mode access
spanning-tree portfast
end
```

Configuring UplinkFast

This topic identifies the command used to configure UplinkFast.

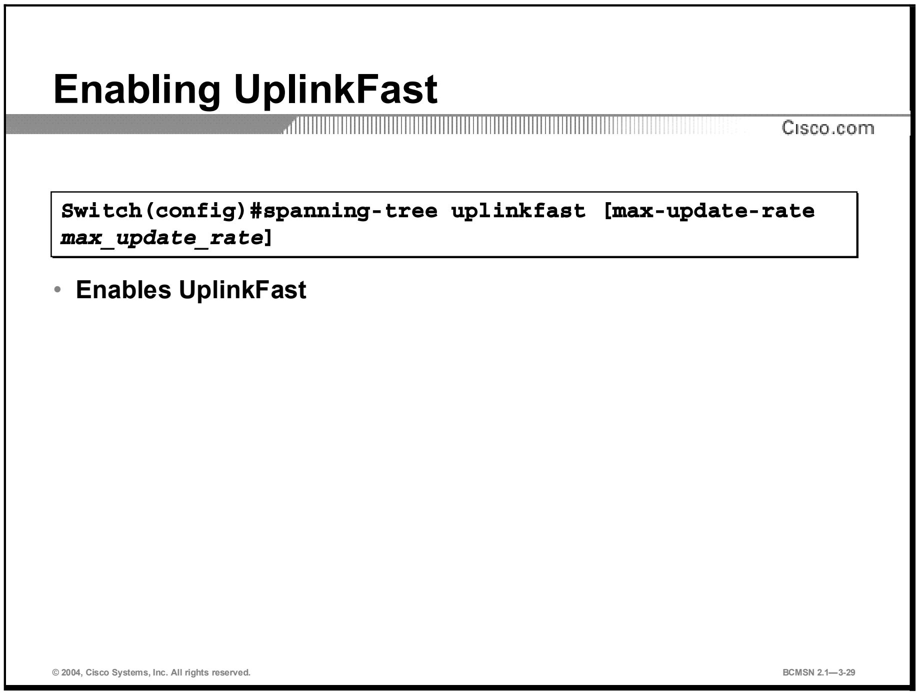

Enabling UplinkFast

Cisco.com

```
Switch(config)#spanning-tree uplinkfast [max-update-rate
max_update_rate]
```

- **Enables UplinkFast**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-29

UplinkFast increases the bridge priority to 49,152 and adds a value of 3000 to the spanning tree port cost of all interfaces on the switch. In this case, it is unlikely that the switch will become the root switch. UplinkFast does not increase the bridge priority or increment the port cost for spanning tree with nondefault bridge priority value and nondefault port cost.

UplinkFast reduces TCN messages that would normally be transmitted toward the root bridge. Instead, the access-layer switch transmits dummy multicast frames for each workstation attached to the switch. The transmission of these dummy multicast frames has the effect of flushing the old forwarding entries for the failed path and reinstalling forwarding entries for the new path. This mechanism updates the forwarding entries for the failed path only, without changing the unaffected entries. UplinkFast cannot be enabled on VLANs configured for bridge priority. To enable UplinkFast on a VLAN with bridge priority configured, set the bridge priority on the VLAN to the default value by entering the **no spanning-tree vlan** *vlan_ID* **priority** command in global configuration mode.

Note Enabling UplinkFast affects all VLANs on the switch. You cannot configure UplinkFast on an individual VLAN.

To enable or disable UplinkFast, use this command:

```
Switch(config)#[no] spanning-tree uplinkfast [max-update-rate
max_update_rate]
```

The *max_update_rate* value represents the number of multicast packets transmitted per second. The default value is 150 packets per second (pps).

Example: Enabling UplinkFast

This example shows how to enable UplinkFast with a maximum update rate of 400 pps:

```
Switch(config)#spanning-tree uplinkfast max-update-rate 400
```

Note This command should be used on access switches only.

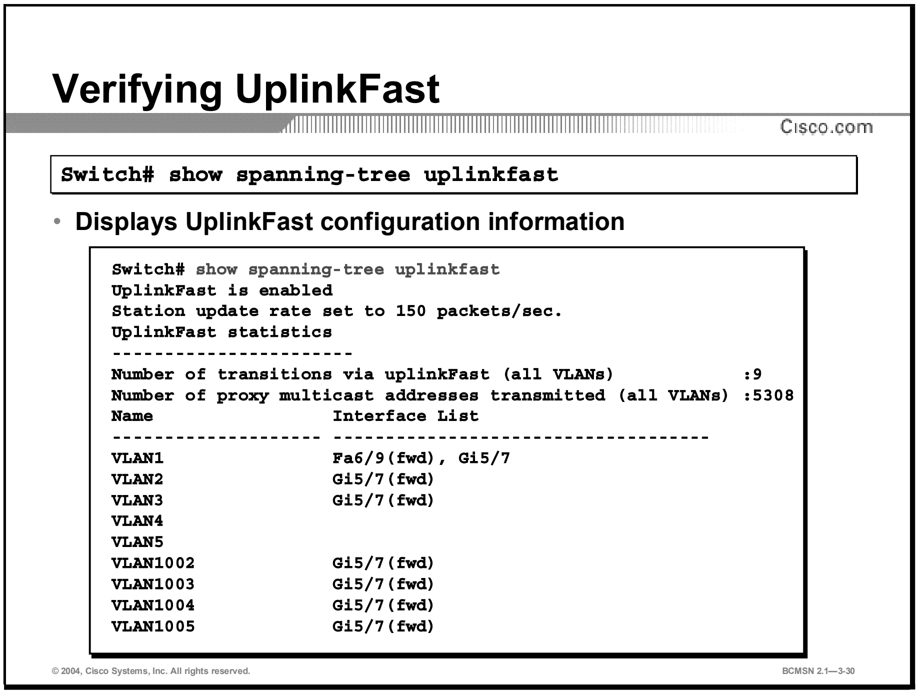

Verifying UplinkFast

Cisco.com

```
Switch# show spanning-tree uplinkfast
```

- **Displays UplinkFast configuration information**

```
Switch# show spanning-tree uplinkfast
UplinkFast is enabled
Station update rate set to 150 packets/sec.
UplinkFast statistics
-----------------------
Number of transitions via uplinkFast (all VLANs)            :9
Number of proxy multicast addresses transmitted (all VLANs) :5308
Name                 Interface List
-------------------- ------------------------------------
VLAN1                Fa6/9(fwd), Gi5/7
VLAN2                Gi5/7(fwd)
VLAN3                Gi5/7(fwd)
VLAN4
VLAN5
VLAN1002             Gi5/7(fwd)
VLAN1003             Gi5/7(fwd)
VLAN1004             Gi5/7(fwd)
VLAN1005             Gi5/7(fwd)
```

BCMSN 2.1—3-30

Use the **show spanning-tree uplinkfast** command to verify UplinkFast configuration.

Example: Verifying UplinkFast

This example shows how to identify which VLANS have UplinkFast enabled:

```
Switch#show spanning-tree uplinkfast

UplinkFast is enabled

Station update rate set to 150 packets/sec.

UplinkFast statistics

-----------------------

Number of transitions via uplinkFast (all VLANs)
:14

Number of proxy multicast addresses transmitted (all VLANs)
:5308

Name                 Interface List

-------------------- ------------------------------------

VLAN1                Fa6/9(fwd), Gi5/7

VLAN2                Gi5/7(fwd)

VLAN3                Gi5/7(fwd)

VLAN4

VLAN5

VLAN6

VLAN7

VLAN8

VLAN10

VLAN15

VLAN1002             Gi5/7(fwd)

VLAN1003             Gi5/7(fwd)

VLAN1004             Gi5/7(fwd)

VLAN1005             Gi5/7(fwd)
```

Configuring BackboneFast

This topic identifies the command used to configure BackboneFast.

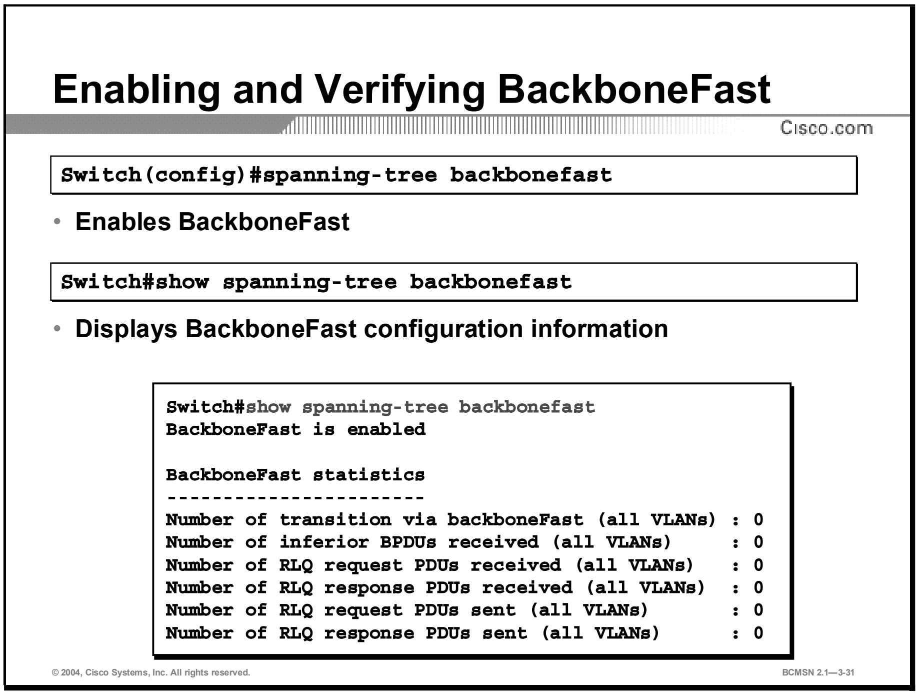

Enabling and Verifying BackboneFast

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

- Enables BackboneFast

```
Switch#show spanning-tree backbonefast
```

- Displays BackboneFast configuration information

```
Switch#show spanning-tree backbonefast
BackboneFast is enabled

BackboneFast statistics
-----------------------
Number of transition via backboneFast (all VLANs) : 0
Number of inferior BPDUs received (all VLANs)     : 0
Number of RLQ request PDUs received (all VLANs)   : 0
Number of RLQ response PDUs received (all VLANs)  : 0
Number of RLQ request PDUs sent (all VLANs)       : 0
Number of RLQ response PDUs sent (all VLANs)      : 0
```

BCMSN 2.1—3-31

For BackboneFast to work, you must enable it on all switches in the network. BackboneFast is supported for use with third-party switches, but it is not supported on Token Ring VLANs.

To enable or disable BackboneFast, use this command:

```
Switch(config)#[no] spanning-tree backbonefast
```

To verify BackboneFast configuration, use this command:

```
Switch#show spanning-tree backbonefast
```

Example: Configuring BackboneFast

This example shows how to enable and verify BackboneFast on a switch:

```
Switch(config)#spanning-tree backbonefast
Switch(config)#end
Switch#show spanning-tree backbonefast
BackboneFast is enabled

BackboneFast statistics
-----------------------
Number of transition via backboneFast (all VLANs) : 0
Number of inferior BPDUs received (all VLANs)     : 0
Number of RLQ request PDUs received (all VLANs)   : 0
Number of RLQ response PDUs received (all VLANs)  : 0
Number of RLQ request PDUs sent (all VLANs)       : 0
Number of RLQ response PDUs sent (all VLANs)      : 0
```

Enabling Multiple Spanning Tree

This topic identifies the command used to enable MST.

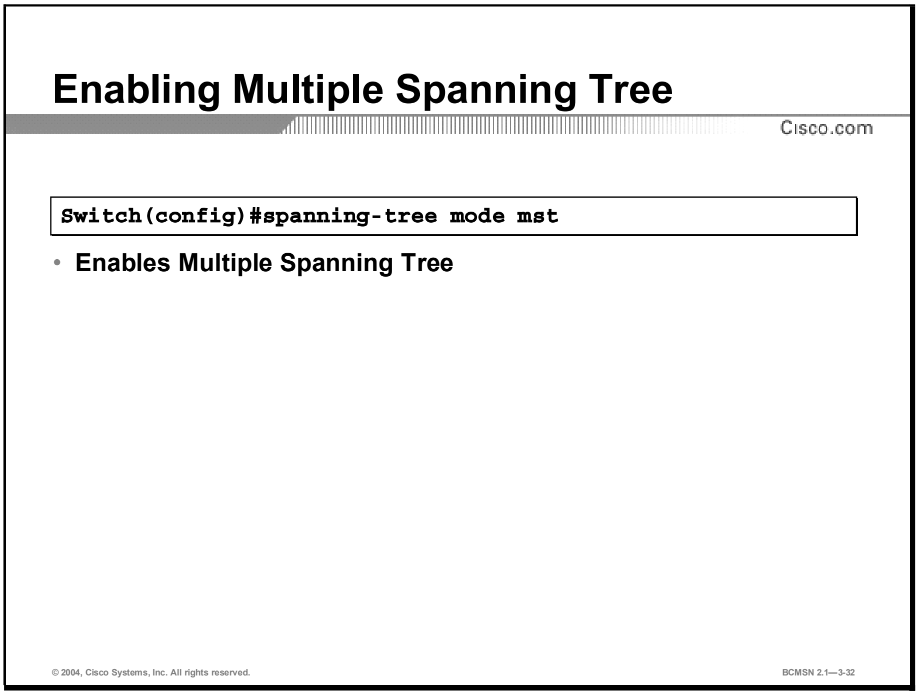

Enabling Multiple Spanning Tree

```
Switch(config)#spanning-tree mode mst
```

- **Enables Multiple Spanning Tree**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-32

Because MST applies to multiple VLANs, it requires some additional configuration beyond that needed for PVST+ or Rapid PVST+. After you have enabled MST with the command **spanning-tree mode mst**, you must configure the regions and instances with additional configuration commands.

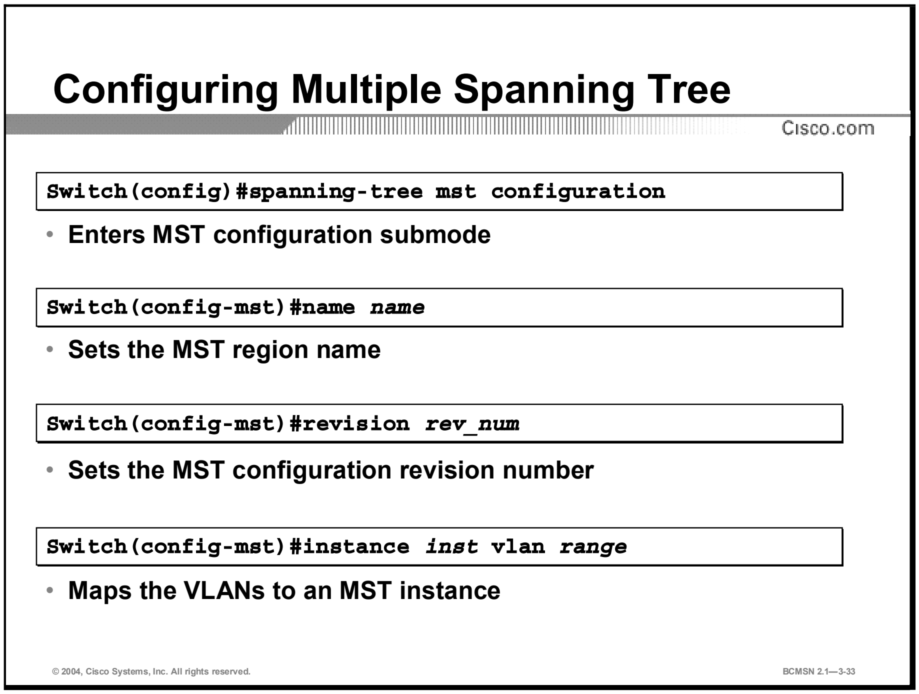

Configuring Multiple Spanning Tree

```
Switch(config)#spanning-tree mst configuration
```

- Enters MST configuration submode

```
Switch(config-mst)#name name
```

- Sets the MST region name

```
Switch(config-mst)#revision rev_num
```

- Sets the MST configuration revision number

```
Switch(config-mst)#instance inst vlan range
```

- Maps the VLANs to an MST instance

Use these commands to configure MST:

Step	Description	Notes and Comments
1.	Enter MST configuration submode. `Switch(config)#spanning-tree mst configuration`	You can use the **no** keyword to clear the MST configuration.
2.	Display the current MST configuration. `Switch(config-mst)#show current`	
3.	Set the MST region name. `Switch(config-mst)#name name`	
4.	Set the MST configuration revision number. `Switch(config-mst)#revision revision_number`	The revision number can be any unassigned 16-bit integer. It is not incremented automatically when you commit a new MST configuration.
5.	Map the VLANs to an MSTI. `Switch(config-mst)#instance instance_number vlan vlan_range`	If you do not specify the **vlan** keyword, you can use the **no** keyword to unmap all the VLANs that were mapped to an MSTI. If you specify the **vlan** keyword, you can use the **no** keyword to unmap a specified VLAN from an MSTI.
6.	Display the new MST configuration to be applied. `Switch(config-mst)#show pending`	
7.	Apply the configuration and exit MST configuration submode. `Switch(config-mst)#end`	

Example: Enabling MST

This example shows how to enable two instances of MST:

```
Switch#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)#spanning-tree mode mst
Switch(config)#spanning-tree mst configuration
Switch(config-mst)#show current
Current MST configuration
Name      []
Revision  0
Instance  Vlans mapped
--------  -----------------------------------------------------
-------
0         1-4094
---------------------------------------------------------------
-------
Switch(config-mst)#name cisco
Switch(config-mst)#revision 2
Switch(config-mst)#instance 1 vlan 1
Switch(config-mst)#instance 2 vlan 1-1000
Switch(config-mst)#show pending
Pending MST configuration
Name      [cisco]
Revision  2
Instance  Vlans mapped
--------  -----------------------------------------------------
-------
0         1001-4094
2         1-1000
---------------------------------------------------------------
-------
Switch(config-mst)#no instance 2
Switch(config-mst)#show pending
Pending MST configuration
Name      [cisco]
Revision  2
Instance  Vlans mapped
--------  -----------------------------------------------------
-------
0         1-4094
```

```
--------------------------------------------------------------
-------
Switch(config-mst)#instance 1 vlan 2000-3000
Switch(config-mst)#no instance 1 vlan 2500
Switch(config-mst)#show pending
Pending MST configuration
Name      [cisco]
Revision  2
Instance  Vlans mapped
--------  ----------------------------------------------------
-------
0         1-1999,2500,3001-4094
1         2000-2499,2501-3000
--------------------------------------------------------------
-------
Switch(config-mst)#end
```

Verifying Multiple Spanning Tree

This topic identifies the command used to verify MST.

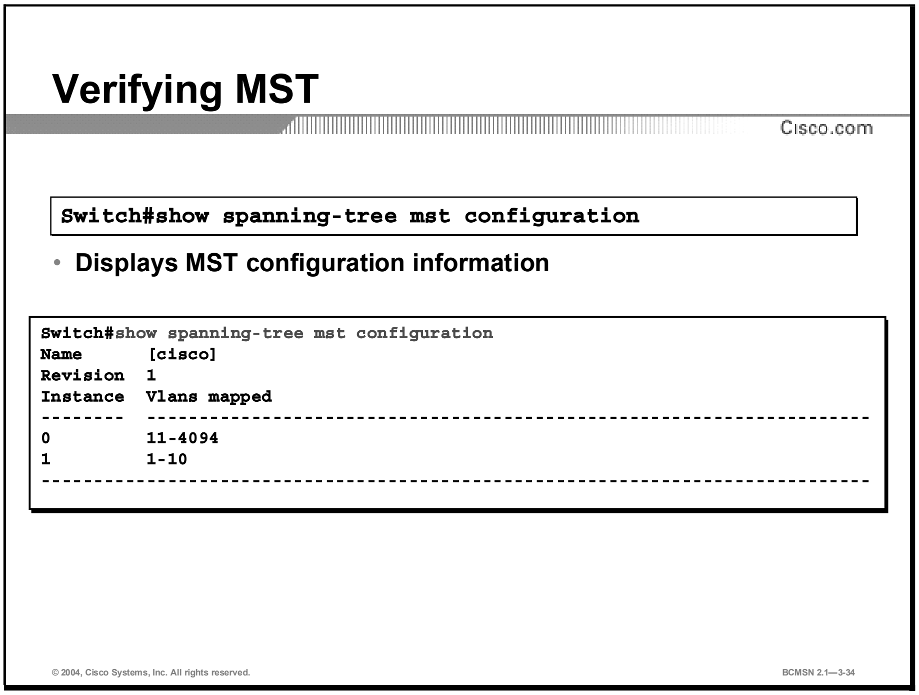

Verifying MST

```
Switch#show spanning-tree mst configuration
```

- **Displays MST configuration information**

```
Switch#show spanning-tree mst configuration
Name      [cisco]
Revision  1
Instance  Vlans mapped
--------  ---------------------------------------------------------------------
0         11-4094
1         1-10
-------------------------------------------------------------------------------
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-34

Use the **show spanning-tree mst** command to display MST information.

Example: Displaying MST Configuration Information

This example shows how to display MST configuration information:

```
Switch#show spanning-tree mst configuration
Name      [cisco]
Revision  1
Instance  Vlans mapped
--------  ----------------------------------------------------
-------
0         11-4094
1         1-10
--------------------------------------------------------------
-------
```

Example: Displaying General MST Information

This example shows how to display general MST information:

```
Switch#show spanning-tree mst

 ###### MST00        vlans mapped:  11-4094

Bridge       address 00d0.00b8.1400  priority  32768 (32768
sysid 0)

Root         address 00d0.004a.3c1c  priority  32768 (32768
sysid 0)

             port    Fa4/48          path cost 203100

IST master   this switch

Operational hello time 2, forward delay 15, max age 20, max
hops 20

Configured  hello time 2, forward delay 15, max age 20, max
hops 20


Interface        Role Sts Cost      Prio.Nbr Status

---------------- ---- --- --------- -------- -----------------
-------

Fa4/4            Back BLK 1000      240.196  Point-to-point

Fa4/5            Desg FWD 200000    128.197  Point-to-point

Fa4/48           Root FWD 200000    128.240  Point-to-point
Bound(STP)


###### MST01        vlans mapped:  1-10

Bridge       address 00d0.00b8.1400  priority  32769 (32768
sysid 1)

Root         this switch for MST01


Interface        Role Sts Cost      Prio.Nbr Status


Fa4/4            Back BLK 1000      240.196  Point-to-point

Fa4/5            Desg FWD 200000    128.197  Point-to-point

Fa4/48           Boun FWD 200000    128.240  Point-to-point
Bound(STP)
```

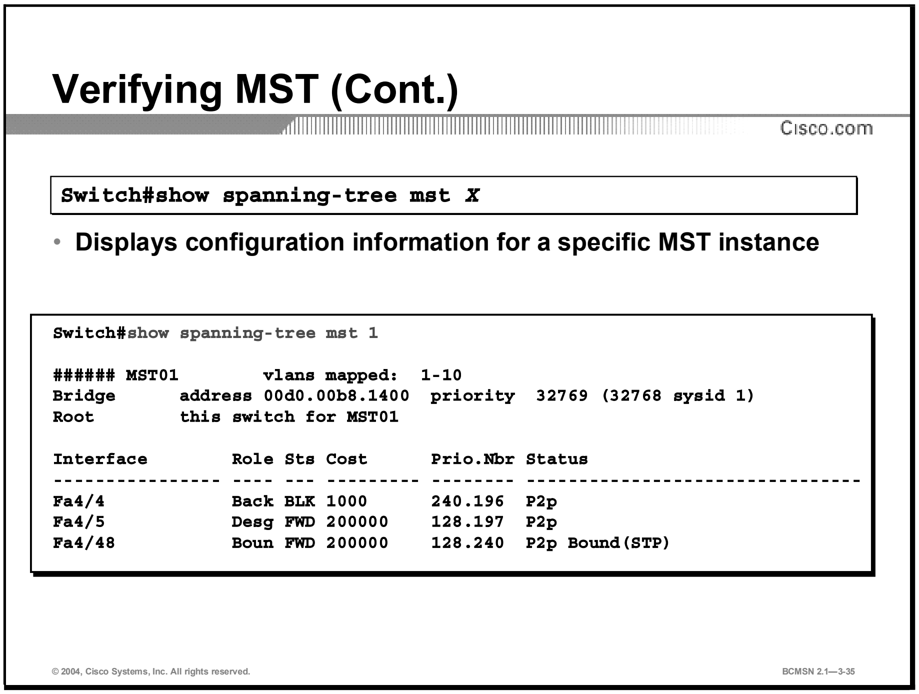

Verifying MST (Cont.)

Cisco.com

```
Switch#show spanning-tree mst X
```

- **Displays configuration information for a specific MST instance**

```
Switch#show spanning-tree mst 1

###### MST01        vlans mapped:  1-10
Bridge      address 00d0.00b8.1400  priority  32769 (32768 sysid 1)
Root        this switch for MST01

Interface        Role Sts Cost      Prio.Nbr Status
---------------- ---- --- --------- -------- --------------------------------
Fa4/4            Back BLK 1000      240.196  P2p
Fa4/5            Desg FWD 200000    128.197  P2p
Fa4/48           Boun FWD 200000    128.240  P2p Bound(STP)
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-35

Example: Displaying MST Information for a Specific Instance

This example displays spanning tree information for a specific MSTI:

```
Switch#show spanning-tree mst 1


###### MST01        vlans mapped:  1-10

Bridge      address 00d0.00b8.1400  priority  32769 (32768
sysid 1)

Root        this switch for MST01


Interface        Role Sts Cost      Prio.Nbr Status

---------------- ---- --- --------- -------- -----------------
-------

Fa4/4            Back BLK 1000      240.196  Point-to-point

Fa4/5            Desg FWD 200000    128.197  Point-to-point

Fa4/48           Boun FWD 200000    128.240  Point-to-point
Bound(STP)
```

Example: Displaying MST Information for a Specific Interface

This example displays MST information for a specific interface:

```
Switch#show spanning-tree mst interface fastethernet 4/4

FastEthernet4/4 of MST00 is backup blocking
Edge port:no             (default)        port guard :none
(default)
Link type:point-to-point (auto)           bpdu filter:disable
(default)
Boundary :internal                        bpdu guard :disable
(default)
Bpdus sent 2, received 368

Instance Role Sts Cost      Prio.Nbr Vlans mapped
-------- ---- --- --------- -------- -------------------------
------
0        Back BLK 1000      240.196  11-4094
1        Back BLK 1000      240.196  1-10
```

Example: Displaying MST Information for a Specific Instance and Interface

This example displays MST information for a specific interface and a specific MSTI:

```
Switch#show spanning-tree mst 1 interface fastethernet 4/4

FastEthernet4/4 of MST01 is backup blocking
Edge port:no             (default)        port guard :none
(default)
Link type:point-to-point (auto)           bpdu filter:disable
(default)
Boundary :internal                        bpdu guard :disable
(default)
Bpdus (MRecords) sent 2, received 364

Instance Role Sts Cost      Prio.Nbr Vlans mapped
-------- ---- --- --------- -------- -------------------------
------
1        Back BLK 1000      240.196  1-10
```

Example: Displaying Detailed MST Information

This example displays detailed MST information for a specific instance:

```
Switch#show spanning-tree mst 1 detail


###### MST01        vlans mapped:  1-10

Bridge      address 00d0.00b8.1400  priority  32769 (32768
sysid 1)

Root        this switch for MST01


FastEthernet4/4 of MST01 is backup blocking

Port info            port id        240.196  priority    240
cost        1000

Designated root      address 00d0.00b8.1400  priority  32769
cost           0

Designated bridge    address 00d0.00b8.1400  priority  32769
port id  128.197

Timers:message expires in 5 sec, forward delay 0, forward
transitions 0

Bpdus (MRecords) sent 123, received 1188


FastEthernet4/5 of MST01 is designated forwarding

Port info            port id        128.197  priority    128
cost      200000

Designated root      address 00d0.00b8.1400  priority  32769
cost           0

Designated bridge    address 00d0.00b8.1400  priority  32769
port id  128.197

Timers:message expires in 0 sec, forward delay 0, forward
transitions 1

Bpdus (MRecords) sent 1188, received 123


FastEthernet4/48 of MST01 is boundary forwarding

Port info            port id        128.240  priority    128
cost      200000

Designated root      address 00d0.00b8.1400  priority  32769
cost           0

Designated bridge    address 00d0.00b8.1400  priority  32769
port id  128.240

Timers:message expires in 0 sec, forward delay 0, forward
transitions 1

Bpdus (MRecords) sent 78, received 0
```

Applying Guidelines to Configuring EtherChannel

This topic identifies the appropriate guidelines that apply to configuring EtherChannel.

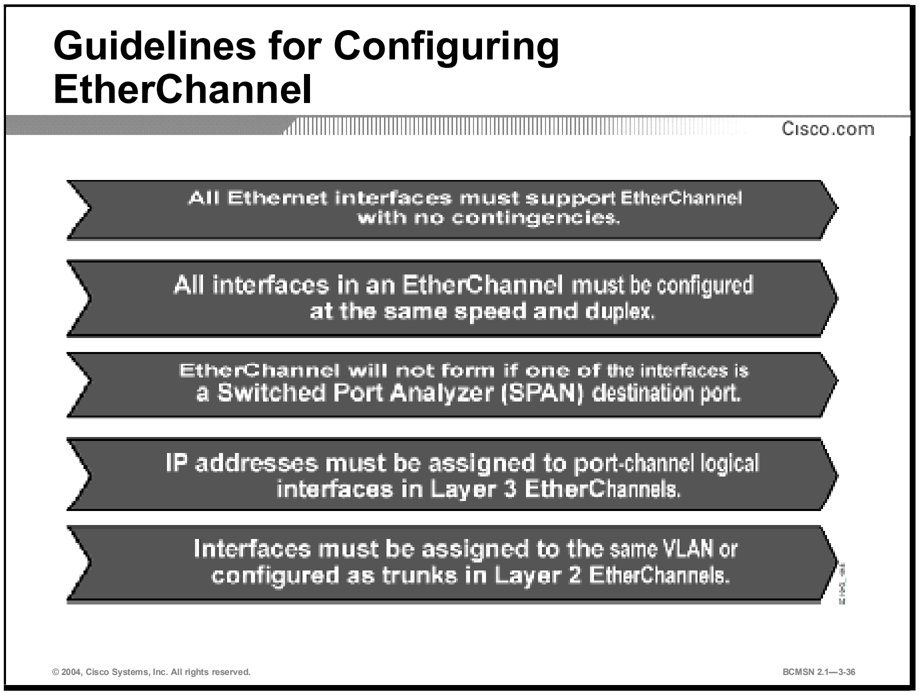

Guidelines for Configuring EtherChannel

Cisco.com

- All Ethernet interfaces must support EtherChannel with no contingencies.
- All interfaces in an EtherChannel must be configured at the same speed and duplex.
- EtherChannel will not form if one of the interfaces is a Switched Port Analyzer (SPAN) destination port.
- IP addresses must be assigned to port-channel logical interfaces in Layer 3 EtherChannels.
- Interfaces must be assigned to the same VLAN or configured as trunks in Layer 2 EtherChannels.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-36

Follow these guidelines and restrictions when configuring EtherChannel interfaces:

- All Ethernet interfaces on all modules support EtherChannel (maximum of eight interfaces), with no requirement that interfaces be physically contiguous or on the same module.
- Configure all interfaces in an EtherChannel to operate at the same speed and in the same duplex mode.
- Enable all interfaces in an EtherChannel. If you shut down an interface in an EtherChannel, it is treated as a link failure and its traffic is transferred to one of the remaining interfaces in the EtherChannel.
- An EtherChannel will not form if one of the interfaces is a Switched Port Analyzer (SPAN) destination port.
- For Layer 3 EtherChannels, assign Layer 3 addresses to the port-channel logical interface, not to the physical interfaces in the channel.

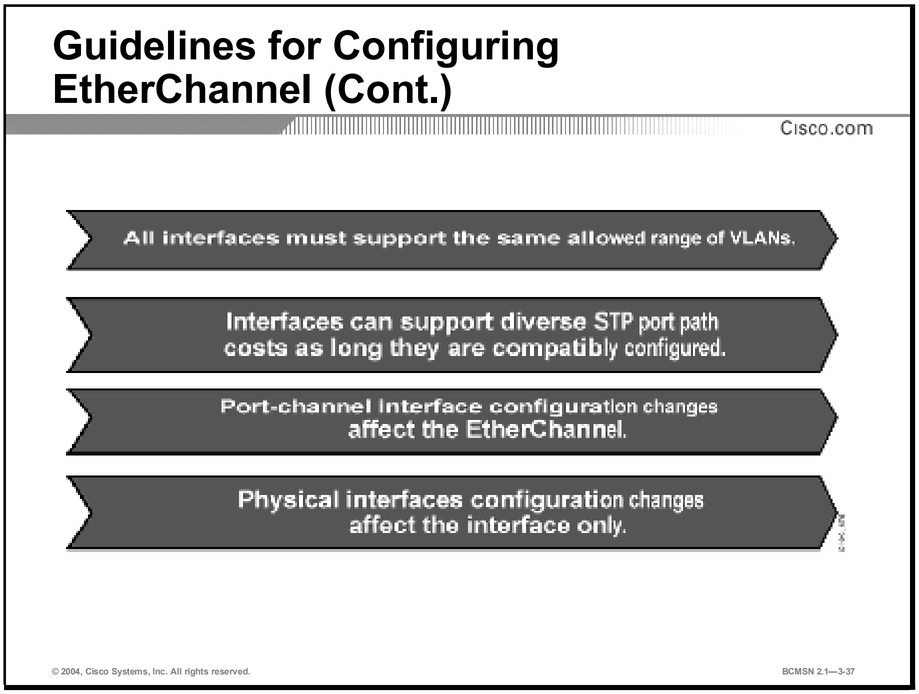

Guidelines for Configuring EtherChannel (Cont.)

Cisco.com

- All interfaces must support the same allowed range of VLANs.
- Interfaces can support diverse STP port path costs as long they are compatibly configured.
- Port-channel interface configuration changes affect the EtherChannel.
- Physical interfaces configuration changes affect the interface only.

BCMSN 2.1—3-37

- For Layer 2 EtherChannels, either assign all interfaces in the EtherChannel to the same VLAN or configure them as trunks. If you configure an EtherChannel from trunk interfaces, verify that the trunking mode is the same on all the trunks. Interfaces in an EtherChannel with different trunk modes can have unexpected results.
- An EtherChannel supports the same allowed range of VLANs on all the interfaces in a trunking Layer 2 EtherChannel. If the allowed range of VLANs is not the same, the interfaces do not form an EtherChannel, even when set to the auto or desirable mode.
- Interfaces with different STP port path costs can form an EtherChannel as long they are otherwise compatibly configured. Setting different STP port path costs does not, by itself, make interfaces incompatible for the formation of an EtherChannel.
- After you configure an EtherChannel, any configuration you apply to the port-channel interface affects the EtherChannel. Any configuration you apply to the physical interfaces affects only the specific interface you configured.

Configuring EtherChannel

This topic identifies the commands used to configure EtherChannel.

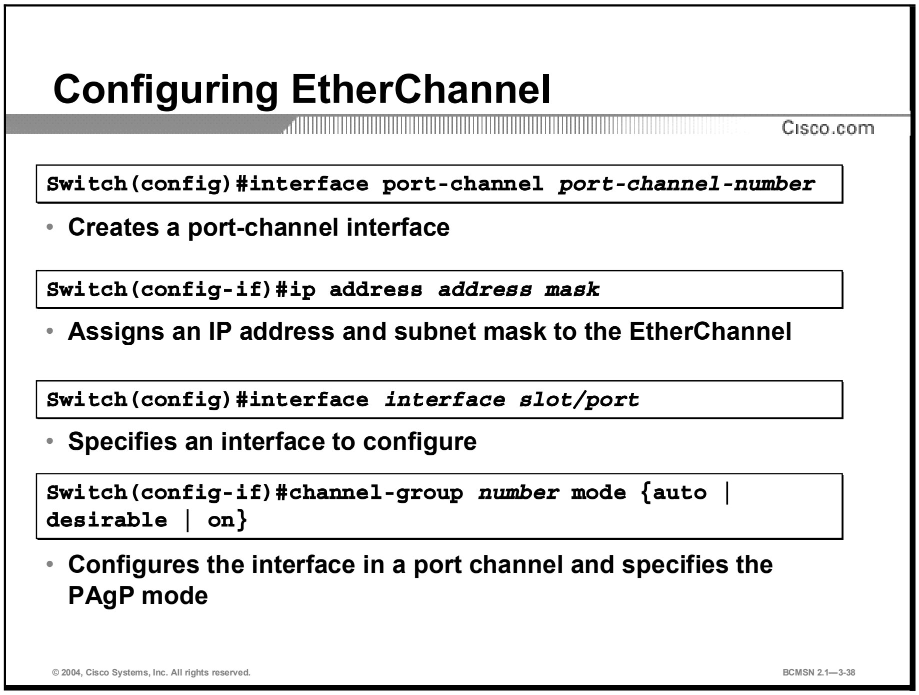

Configuring EtherChannel

Cisco.com

```
Switch(config)#interface port-channel port-channel-number
```

- **Creates a port-channel interface**

```
Switch(config-if)#ip address address mask
```

- **Assigns an IP address and subnet mask to the EtherChannel**

```
Switch(config)#interface interface slot/port
```

- **Specifies an interface to configure**

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

- **Configures the interface in a port channel and specifies the PAgP mode**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-38

To configure Layer 3 EtherChannels, create the port-channel logical interface and then put the Ethernet interfaces into the port channel. To configure physical interfaces as Layer 3 EtherChannels, configure each interface with the **channel-group** command.

Note	To move an IP address from a physical interface to an EtherChannel, you must delete the IP address from the physical interface before configuring it on the port-channel interface.

To configure Layer 2 EtherChannels, configure the Ethernet interfaces with the **channel-group** command. This creates the port-channel logical interface.

Note	IOS creates port-channel interfaces for Layer 2 EtherChannels when you configure Layer 2 Ethernet interfaces with the **channel-group** command.

Example: Creating a Layer 3 EtherChannel

This example creates a port-channel interface for a Layer 3 EtherChannel and then configures Fast Ethernet interfaces 5/4 and 5/5 into port channel 1 with Port Aggregation Protocol (PAgP) mode desirable:

```
Switch#configure terminal
Switch(config)#interface port-channel 1
Switch(config-if)#ip address 172.32.52.10 255.255.255.0
Switch(config-if)#end
Switch(config)#interface range fastethernet 5/4 - 5 (Note:
Space is mandatory.)
Switch(config-if-range)#no switchport
Switch(config-if-range)#no ip address
Switch(config-if-range)#channel-group 1 mode desirable
Switch(config-if-range)#end
```

Example: Creating a Layer 2 EtherChannel

This example shows how to configure Fast Ethernet interfaces 5/6 and 5/7 into port channel 2 with PAgP mode desirable:

```
Switch#configure terminal
Switch(config)#interface range fastethernet 5/6 - 7 (Note:
Space is mandatory.)
Switch(config-if-range)#channel-group 2 mode desirable
Switch(config-if-range)#end
```

Verifying EtherChannel

This topic identifies the commands used to verify EtherChannel.

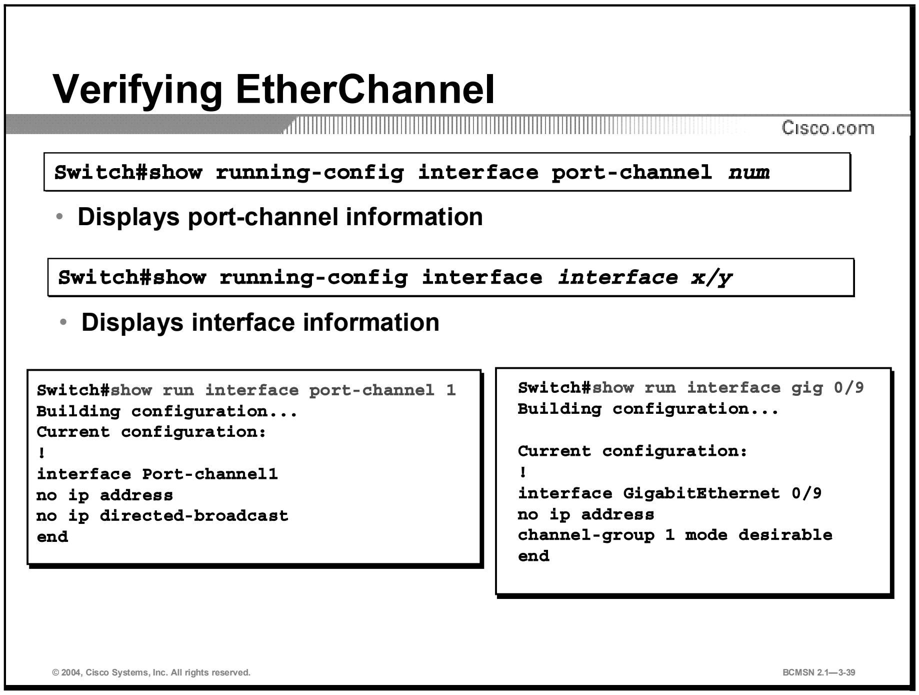

Verifying EtherChannel

Cisco.com

```
Switch#show running-config interface port-channel num
```

- **Displays port-channel information**

```
Switch#show running-config interface interface x/y
```

- **Displays interface information**

```
Switch#show run interface port-channel 1
Building configuration...
Current configuration:
!
interface Port-channel1
no ip address
no ip directed-broadcast
end
```

```
Switch#show run interface gig 0/9
Building configuration...

Current configuration:
!
interface GigabitEthernet 0/9
no ip address
channel-group 1 mode desirable
end
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-39

Use the **show running-config** command to display information about the port channel and the specific EtherChannel interfaces.

Example: Verifying the Configuration of a Port Channel

This example shows how to verify the configuration of port-channel interface 1:

```
Switch#show running-config interface port-channel 1
Building configuration...

Current configuration:
!
interface Port-channel1
ip address 172.32.52.10 255.255.255.0
no ip directed-broadcast
end
```

Example: Verifying the Configuration of a Layer 3 EtherChannel

The following two examples show how to verify the configuration of Fast Ethernet interface 5/4:

```
Switch#show running-config interface fastethernet 5/4
Building configuration...

Current configuration:
!
interface FastEthernet5/4
no ip address
no switchport
no ip directed-broadcast
channel-group 1 mode desirable
end

Switch#show interfaces fastethernet 5/4 etherchannel
Port state    = EC-Enbld Up In-Bndl Usr-Config
Channel group = 1           Mode = Desirable     Gcchange = 0
Port-channel  = Po1         GC   = 0x00010001    Pseudo-port-
channel = Po1
Port indx     = 0           Load = 0x55

Flags:  S - Device is sending Slow hello.  C - Device is in
Consistent state.
        A - Device is in Auto mode.        P - Device learns
on physical port.
Timers: H - Hello timer is running.        Q - Quit timer is
running.
        S - Switching timer is running.    I - Interface timer
is running.

Local information:
                                Hello    Partner  PAgP
Learning  Group
Port      Flags State   Timers  Interval Count   Priority
Method  Ifindex
Fa5/4     SC    U6/S7           30s      1        128
Any      55

Partner's information:
```

```
          Partner              Partner         Partner
Partner Group
Port      Name                 Device ID       Port       Age
Flags   Cap.
Fa5/4     JAB031301            0050.0f10.230c   2/45
1s SAC     2D

Age of the port in the current state: 00h:54m:52s

Switch#
```

Example: Verifying the Configuration of a Layer 2 EtherChannel

The following two examples show how to verify the configuration of Fast Ethernet interface 5/6:

```
Switch#show running-config interface fastethernet 5/6
Building configuration...

Current configuration:
!
interface FastEthernet5/6
switchport access vlan 10
switchport mode access
channel-group 2 mode desirable
end

Switch#show interfaces fastethernet 5/6 etherchannel
Port state    = EC-Enbld Up In-Bndl Usr-Config
Channel group = 1           Mode = Desirable     Gcchange = 0
Port-channel  = Po1         GC   = 0x00010001
Port indx     = 0           Load = 0x55

Flags:  S - Device is sending Slow hello.  C - Device is in
Consistent state.
        A - Device is in Auto mode.        P - Device learns
on physical port.
Timers: H - Hello timer is running.        Q - Quit timer is
running.
        S - Switching timer is running.    I - Interface timer
is running.

Local information:
```

```
                                Hello    Partner  PAgP
Learning  Group
Port      Flags State   Timers  Interval Count   Priority
Method  Ifindex
Fa5/6     SC    U6/S7           30s      1        128
Any      56
Partner's information:


          Partner              Partner          Partner
Partner Group
Port      Name                 Device ID        Port       Age
Flags   Cap.
Fa5/6     JAB031301            0050.0f10.230c   2/47
18s SAC     2F


Age of the port in the current state: 00h:10m:57s
```

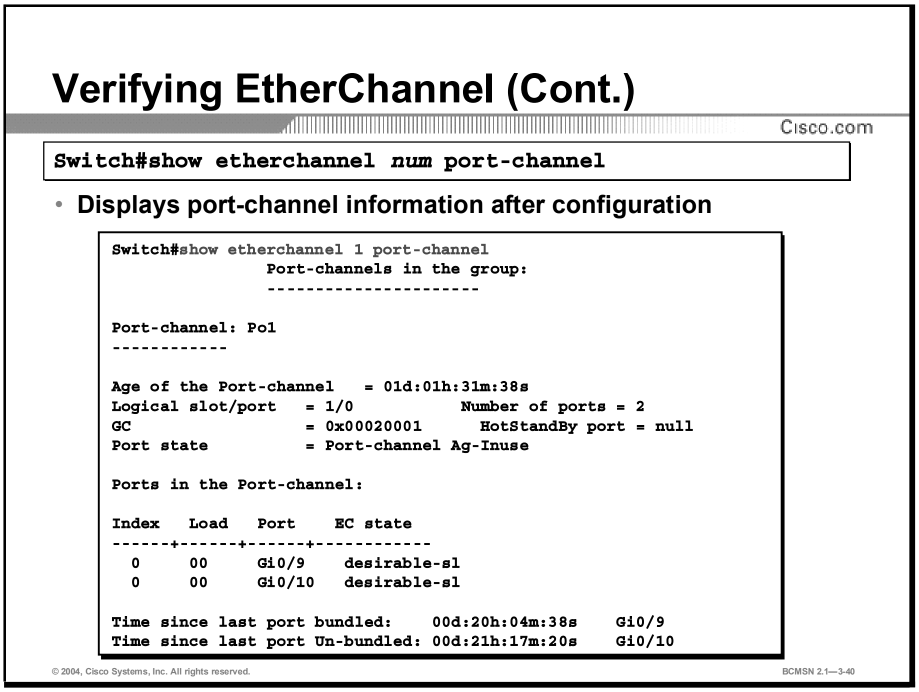

Verifying EtherChannel (Cont.)

Cisco.com

```
Switch#show etherchannel num port-channel
```

- **Displays port-channel information after configuration**

```
Switch#show etherchannel 1 port-channel
                Port-channels in the group:
                ----------------------

Port-channel: Po1
------------

Age of the Port-channel   = 01d:01h:31m:38s
Logical slot/port   = 1/0           Number of ports = 2
GC                  = 0x00020001      HotStandBy port = null
Port state          = Port-channel Ag-Inuse

Ports in the Port-channel:

Index   Load   Port    EC state
------+------+------+------------
  0     00     Gi0/9    desirable-sl
  0     00     Gi0/10   desirable-sl

Time since last port bundled:    00d:20h:04m:38s    Gi0/9
Time since last port Un-bundled: 00d:21h:17m:20s    Gi0/10
```

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN 2.1—3-40

Use the **show etherchannel** command to display port-channel information after configuration.

Example: Verifying Port-Channel Configuration

This example shows how to verify the configuration of port-channel interface 1 after the interfaces have been configured:

```
Switch#show etherchannel 1 port-channel

            Channel-group listing:
            ----------------------
Group: 1
------------

         Port-channels in the group:
         ----------------------
Port-channel: Po1
------------

Age of the Port-channel   = 01h:56m:20s
Logical slot/port   = 10/1           Number of ports = 2
GC                  = 0x00010001      HotStandBy port = null
Port state          = Port-channel L3-Ag Ag-Inuse

Ports in the Port-channel:

Index   Load   Port
-------------------
1     00      Fa5/6
0     00      Fa5/7

Time since last port bundled:    00h:23m:33s    Fa5/6

Switch#
```

This example shows how to verify the configuration of port-channel interface 2 (a Layer 2 EtherChannel) after the interfaces have been configured:

```
Switch#show etherchannel 2 port-channel
Port-channels in the group:
----------------------


Port-channel: Po2
------------


Age of the Port-channel   = 00h:23m:33s
Logical slot/port   = 10/2           Number of ports in agport
= 2
GC                  = 0x00020001      HotStandBy port = null
Port state          = Port-channel Ag-Inuse

Ports in the Port-channel:

Index   Load   Port
-------------------
1     00      Fa5/6
0     00      Fa5/7

Time since last port bundled:    00h:23m:33s    Fa5/6
```

Configuring PAgP and LACP

This topic identifies the commands used to configure PAgP and LACP.

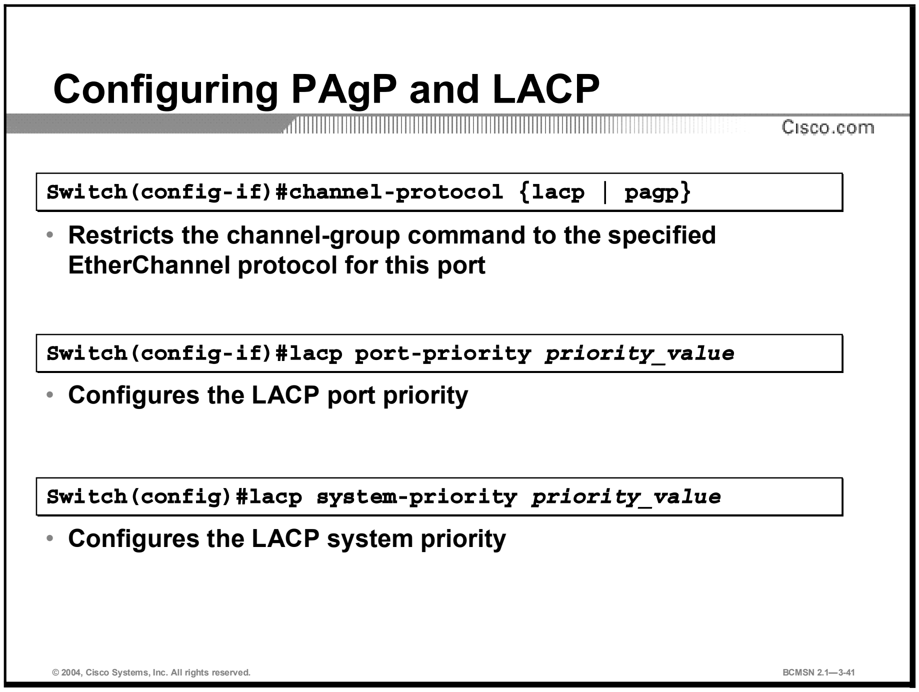

Configuring PAgP and LACP

Cisco.com

```
Switch(config-if)#channel-protocol {lacp | pagp}
```

- Restricts the channel-group command to the specified EtherChannel protocol for this port

```
Switch(config-if)#lacp port-priority priority_value
```

- Configures the LACP port priority

```
Switch(config)#lacp system-priority priority_value
```

- Configures the LACP system priority

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN 2.1—3-41

To configure LACP or PAgP (configured by default), complete the steps in the table:

Step	Description	Notes and Comments
1.	Select a port to configure. `Switch(config)#interface type1 slot/port`	
2.	Restrict the **channel-group** command to the EtherChannel protocol. `Switch(config-if)#channel-protocol {lacp | pagp}`	On the selected LAN port, restrict the **channel-group** command to the EtherChannel protocol configured with the **channel-protocol** command.
3.	Configure the LACP port priority. `Switch(config-if)#lacp port-priority priority_value`	Optional. Valid values are 1 through 65535. Higher numbers have lower priority. The default is 32768.
4.	Exit configuration mode. `Switch(config-if)#end`	
5.	Configure the LACP system priority. `Switch(config)#lacp system-priority priority_value`	The LACP system ID is the combination of the LACP system priority value and the MAC address of the switch. Valid values are 1 through 65535. Higher numbers have lower priority. The default is 32768.

Verifying PAgP and LACP

This topic identifies the commands used to verify PAgP and LACP.

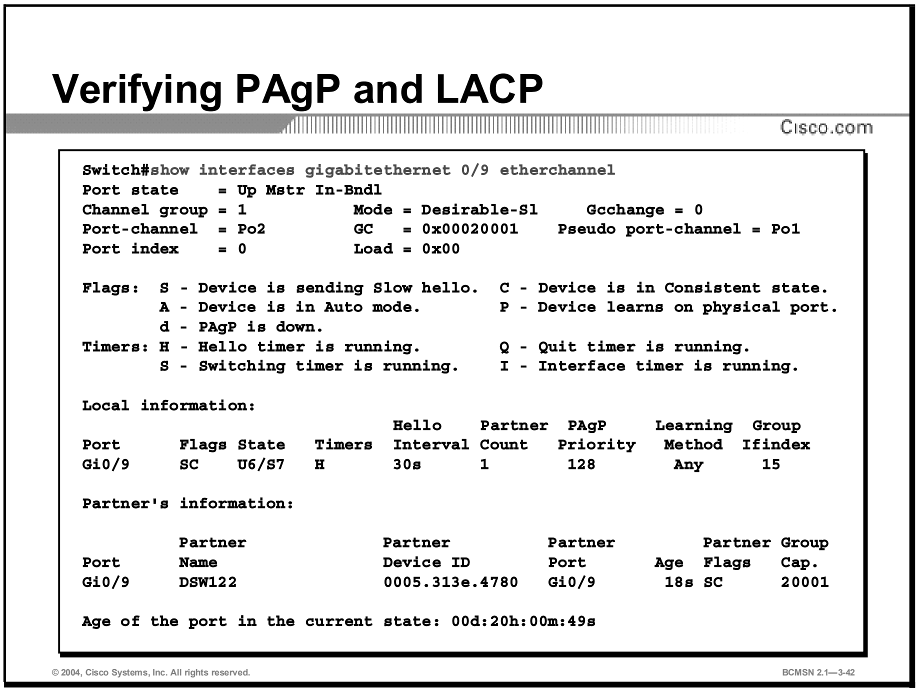

Verifying PAgP and LACP

```
Switch#show interfaces gigabitethernet 0/9 etherchannel
Port state    = Up Mstr In-Bndl
Channel group = 1           Mode = Desirable-Sl     Gcchange = 0
Port-channel  = Po2         GC   = 0x00020001    Pseudo port-channel = Po1
Port index    = 0           Load = 0x00

Flags:  S - Device is sending Slow hello.  C - Device is in Consistent state.
        A - Device is in Auto mode.        P - Device learns on physical port.
        d - PAgP is down.
Timers: H - Hello timer is running.        Q - Quit timer is running.
        S - Switching timer is running.    I - Interface timer is running.

Local information:
                                Hello    Partner  PAgP     Learning  Group
Port      Flags State   Timers  Interval Count   Priority   Method  Ifindex
Gi0/9     SC    U6/S7   H       30s      1        128        Any      15

Partner's information:

          Partner              Partner          Partner         Partner Group
Port      Name                 Device ID        Port       Age  Flags   Cap.
Gi0/9     DSW122               0005.313e.4780   Gi0/9       18s SC      20001

Age of the port in the current state: 00d:20h:00m:49s
```

To verify LACP or PAgP configuration, use one of these commands:

```
Switch#show running-config interface type1 slot/port
Switch#show interfaces type1 slot/port etherchannel
```

Balancing Ethernet Traffic Load

This topic identifies the commands to configure EtherChannel load balancing.

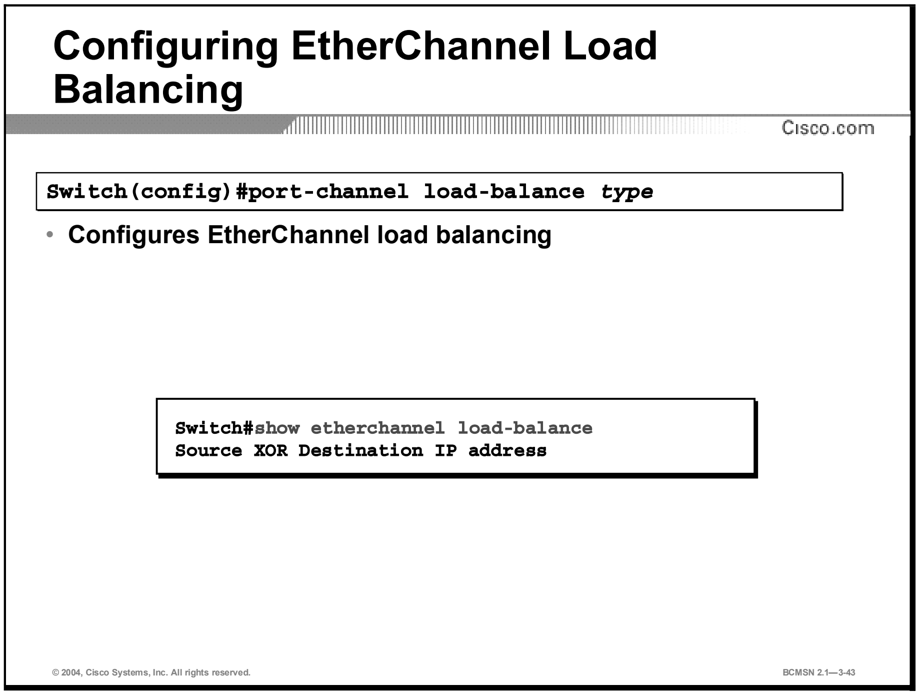

Configuring EtherChannel Load Balancing

Cisco.com

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

- Configures EtherChannel load balancing

```
Switch#show etherchannel load-balance
Source XOR Destination IP address
```

BCMSN 2.1—3-43

EtherChannel balances traffic load across the links in a channel by reducing part of the binary pattern formed from the addresses in the frame to a numerical value that selects one of the links in the channel.

EtherChannel load balancing can use MAC addresses, IP addresses, or Layer 4 port numbers; either source, destination, or both source and destination.

Load balancing operates at the switch level rather than per channel, applying globally for all channels on the switch. To configure EtherChannel load balancing, use the **port-channel load-balance** *type* command. The load-balancing keywords are as follows:

- **src-mac:** Source MAC addresses
- **dst-mac:** Destination MAC addresses
- **src-dst-mac:** Source and destination MAC addresses
- **src-ip:** Source IP addresses
- **dst-ip:** Destination IP addresses
- **src-dst-ip:** Source and destination IP addresses (default)
- **src-port:** Source TCP/User Datagram Protocol (UDP) port
- **dst-port:** Destination TCP/UDP port
- **src-dst-port:** Source and destination TCP/UDP port

Example: Configuring EtherChannel Load Balancing

This example shows how to configure EtherChannel to use source and destination IP addresses:

```
Switch# configure terminal
Switch(config)# port-channel load-balance src-dst-ip
```

This example shows how to verify the configuration:

```
Switch# show etherchannel load-balance
Source XOR Destination IP address
```

Summary

This topic summarizes the key points discussed in this lesson.

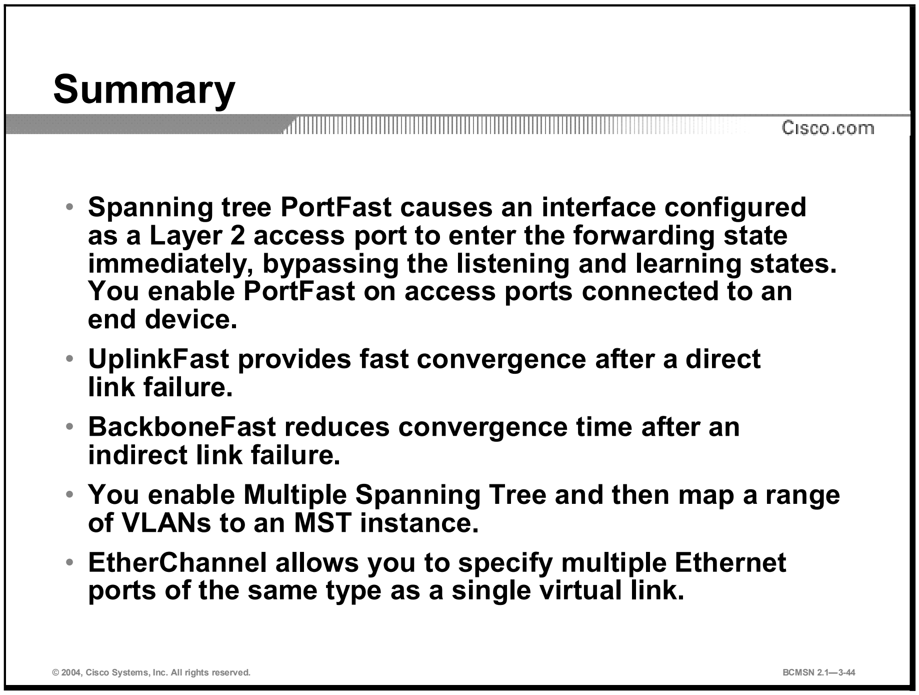

Summary

Cisco.com

- Spanning tree PortFast causes an interface configured as a Layer 2 access port to enter the forwarding state immediately, bypassing the listening and learning states. You enable PortFast on access ports connected to an end device.
- UplinkFast provides fast convergence after a direct link failure.
- BackboneFast reduces convergence time after an indirect link failure.
- You enable Multiple Spanning Tree and then map a range of VLANs to an MST instance.
- EtherChannel allows you to specify multiple Ethernet ports of the same type as a single virtual link.

BCMSN 2.1—3-44

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Select the command that correctly configures port-level tuning with PortFast.

A) switch(config)#**spanning-tree portfast enable**

B) switch(config-if)#**[no] spanning-tree portfast**

C) switch(config-if)#**spanning-tree portfast enable**

D) switch(config)#**spanning-tree portfast interface** *interface x/y*

Q2) Select the command that correctly configures UplinkFast.

A) switch(config)# **spanning-tree uplinkfast**

B) switch(config-if)# **spanning-tree uplinkfast**

C) switch(config)# **spanning-tree uplinkfast enable**

D) switch(config)# **spanning-tree uplinkfast vlan** *vlan-id*

Q3) Select the command that correctly configures BackboneFast.

A) switch(config)# **spanning-tree backbonefast**

B) switch(config-if)# **spanning-tree backbonefast**

C) switch(config)# **spanning-tree backbonefast enable**

D) switch(config)# **spanning-tree backbonefast vlan** *vlan-id*

Q4) Select the commands that correctly configure multiple spanning tree. (Choose two.)

A) switch(config-mst)#**name**

B) switch#**show mst 2 interface fastethernet 5/7**

C) switch(config-mst)#**instance** *instance_number* **vlan** *vlan_range*

D) switch(config-if)#**spanning-tree enable**

Q5) Select the command that correctly displays the STP configuration information for MSTI 2 on interface Fast Ethernet 5/7.

A) switch#**show mst 2 interface fastethernet 5/7**

B) switch#**show stp mst 2 interface fastethernet 5/7**

C) switch#**show mst 2 spanning-tree interface fastethernet 5/7**

D) switch#**show spanning-tree mst 2 interface fastethernet 5/7**

Q6) Select the correct guidelines that apply to configuring EtherChannel. (Choose two.)

A) A maximum of ten interfaces can support EtherChannel.

B) For Layer 2 EtherChannels, assign Layer 2 addresses to the port-channel logical interface.

C) All interfaces should be configured to operate at the same speed and in the same duplex mode.

D) All interfaces should be enabled.

Q7) Select the command used to create the port-channel interface.

A) switch(config)#**interface port-channel** *port-channel-number*

B) switch (config)#**interface address** *address mask*

C) switch (configiif)#**interface port-channel** *port-channel-number*

D) switch(config-if)#**channel-group**

Q8) Select the command that verifies the configuration of port-channel interface 2 before the interfaces have been configured.

A) switch#**show running-config interface channel 2**

B) switch#**show running-config interface fastethernet 5/4**

C) switch#**show running-config interface port-channel 2**

D) switch#**show etherchannel 2 port-channel**

Q9) Select the command that correctly configures the LACP system priority with a high priority.

A) switch(config-if)#**lacp port-priority 2**

B) switch(config)#**lacp port-priority 2**

C) switch(config-if)#**lacp port-priority 5**

D) switch(config)#**lacp port-priority 5**

Q10) Select the command that verifies the LACP configuration.

A) switch#**show running-config interface** *type1 slot/port*

B) switch#**show interfaces** *type1 slot/port* **interface**

C) switch#**show-if interfaces** *type1 slot/port* **etherchannel**

D) switch#**show interfaces** *type1 slot/port* **fastethernet**

Q11) Select the command that configures EtherChannel load balancing.

A) switch#**show etherchannel load-balance**

B) switch(config)#**port-channel load-balance** *type*

C) switch(config)# **port-channel load-balance src-dst-ip**

D) switch#**show-if interfaces** *type1 slot/port* **etherchannel**

Quiz Answer Key

Q1) B

Relates to: Configuring Port-Level Tuning with PortFast

Q2) A

Relates to: Configuring UplinkFast

Q3) A

Relates to: Configuring BackboneFast

Q4) A, C

Relates to: Enabling Multiple Spanning Tree

Q5) A

Relates to: Verifying Multiple Spanning Tree

Q6) C, D

Relates to: Applying Guidelines to Configuring EtherChannel

Q7) A

Relates to: Configuring EtherChannel

Q8) D

Relates to: Verifying EtherChannel

Q9) A

Relates to: Configuring PAgP and LACP

Q10) A

Relates to: Verifying PAgP and LACP

Q11) B

Relates to: Balancing Ethernet Traffic Load

Tuning the Spanning Tree Protocol

Overview

By enhancing spanning tree performance, you can reduce convergence time and failover recovery, help prevent bridging loops, and prevent specific switches from becoming the root bridge.

Relevance

Because spanning tree is so critical to the operation of a multilayer-switched network, being able to tune and troubleshoot this feature is a primary consideration for network administrators.

Objectives

Upon completing this lesson, you will be able to:

- Identify the command to configure BPDU guard
- Identify the command to enable BPDU filtering
- Identify the possible causes of BPDU skewing
- Identify the function of root guard
- Identify the command to configure root guard
- Identify the correct features for UDLD
- Identify the command to configure UDLD
- Identify the features of loop guard
- Identify the command to configure loop guard
- Match the correct feature with UDLD or loop guard
- Identify the protection features used in the Enterprise Composite Network model

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Configuring BPDU Guard
- Enabling BPDU Filtering
- Enabling BPDU Skewing Detection
- Using Root Guard
- Configuring Root Guard
- Shutting Down Unidirectional Links
- Configuring UDLD
- Preventing Layer 2 Forwarding Loops
- Configuring Loop Guard
- Differentiating Between Loop Guard and UDLD
- Applying Protection Features to the Enterprise Composite Network Model
- Summary
- Quiz

Configuring BPDU Guard

This topic identifies the command to configure BPDU guard.

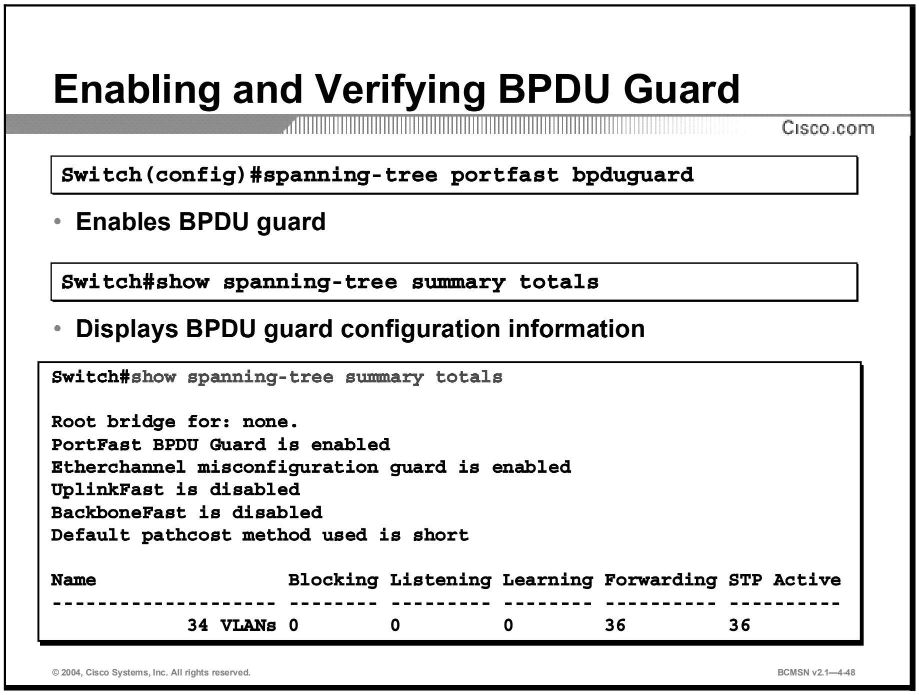

Enabling and Verifying BPDU Guard

```
Switch(config)#spanning-tree portfast bpduguard
```

- Enables BPDU guard

```
Switch#show spanning-tree summary totals
```

- Displays BPDU guard configuration information

```
Switch#show spanning-tree summary totals

Root bridge for: none.
PortFast BPDU Guard is enabled
Etherchannel misconfiguration guard is enabled
UplinkFast is disabled
BackboneFast is disabled
Default pathcost method used is short

Name                 Blocking Listening Learning Forwarding STP Active
-------------------- -------- --------- -------- ---------- ----------
            34 VLANs 0        0         0        36         36
```

BCMSN v2.1—4-48

Spanning tree BPDU guard shuts down PortFast-configured interfaces that receive BPDUs, rather than putting them into the spanning tree blocking state (the default behavior). In a valid configuration, PortFast-configured interfaces do not receive BPDUs. Reception of a BPDU by a PortFast-configured interface signals an invalid configuration, such as connection of an unauthorized device. BPDU guard provides a secure response to invalid configurations, because the administrator must manually put the interface back in service.

Note When the BPDU guard feature is enabled, spanning tree applies BPDU guard to all PortFast-configured interfaces.

To enable BPDU guard to shut down PortFast-configured interfaces that receive BPDUs, or to disable BPDU guard, use this command:

```
Switch(config)# [no] spanning-tree portfast bpduguard
```

Example: Enabling and Verifying BPDU Guard

This example shows how to enable BPDU guard:

```
Switch(config)# spanning-tree portfast bpduguard
Switch(config)# end
```

This example shows how to verify the BPDU configuration:

```
Switch#show spanning-tree summary totals

Root bridge for: none.
PortFast BPDU Guard is enabled
Etherchannel misconfiguration guard is enabled
UplinkFast is disabled
BackboneFast is disabled
Default pathcost method used is short

Name                 Blocking Listening Learning Forwarding
STP Active
-------------------- -------- --------- -------- ---------- --
-------
            34 VLANs 0        0         0        36         36
```

Enabling BPDU Filtering

This topic identifies the command used to enable BPDU filtering.

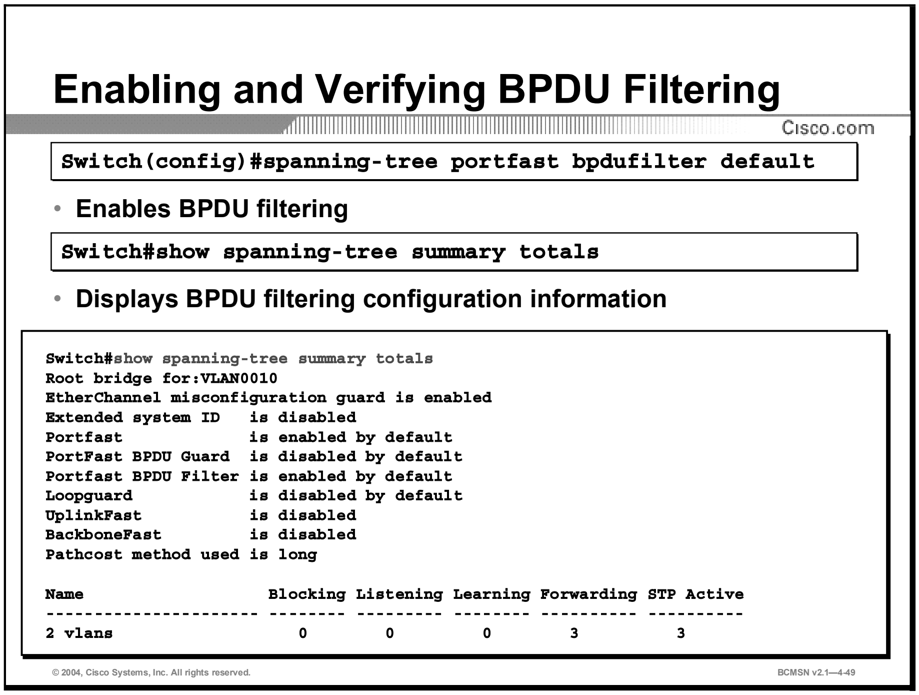

Enabling and Verifying BPDU Filtering

```
Switch(config)#spanning-tree portfast bpdufilter default
```

- **Enables BPDU filtering**

```
Switch#show spanning-tree summary totals
```

- **Displays BPDU filtering configuration information**

```
Switch#show spanning-tree summary totals
Root bridge for:VLAN0010
EtherChannel misconfiguration guard is enabled
Extended system ID   is disabled
Portfast             is enabled by default
PortFast BPDU Guard  is disabled by default
Portfast BPDU Filter is enabled by default
Loopguard            is disabled by default
UplinkFast           is disabled
BackboneFast         is disabled
Pathcost method used is long

Name                   Blocking Listening Learning Forwarding STP Active
---------------------- -------- --------- -------- ---------- ----------
2 vlans                   0        0         0        3          3
```

BCMSN v2.1—4-49

PortFast BPDU filtering allows the administrator to prevent the system from sending or even receiving BPDUs on specified ports. When configured globally, PortFast BPDU filtering applies to all operational PortFast ports. Ports in an operational PortFast state are supposed to be connected to hosts that typically drop BPDUs. If an operational PortFast port receives a BPDU, it immediately loses its operational PortFast status. In that case, PortFast BPDU filtering is disabled on this port, and STP resumes sending BPDUs on this port.

PortFast BPDU filtering can also be configured on a per-port basis. When PortFast BPDU filtering is explicitly configured on a port, it does not send any BPDUs and drops all BPDUs it receives.

When you enable PortFast BPDU filtering globally and set the port configuration as the default for PortFast BPDU filtering, PortFast enables or disables PortFast BPDU filtering.

Caution Explicit configuration of PortFast BPDU filtering on a port that is not connected to a host can result in bridging loops. This is because the port ignores any incoming BPDU and changes its role to the forwarding state.

If the port configuration is not set to default, the PortFast configuration will not affect PortFast BPDU filtering. PortFast BPDU filtering allows access ports to move directly to the forwarding state as soon as the end hosts are connected. The table lists all the possible PortFast BPDU filtering combinations.

Per-Port Configuration	Global Configuration	PortFast State	PortFast BPDU Filtering State
Default	Enable	Enable	Enable
Default	Enable	Disable	Disable
Default	Disable	Not applicable	Disable
Disable	Not applicable	Not applicable	Disable
Enable	Not applicable	Not applicable	Enable

The port transmits at least ten BPDUs. If this port receives any BPDUs, PortFast and PortFast BPDU filtering are disabled.

To enable PortFast BPDU filtering globally on the switch, enter this command:

```
Switch(config)#spanning-tree portfast bpdufilter default
```

To verify the configuration, enter this command:

```
Switch#show spanning-tree summary totals
```

Enabling BPDU Skewing Detection

This topic identifies the possible causes of BPDU skewing.

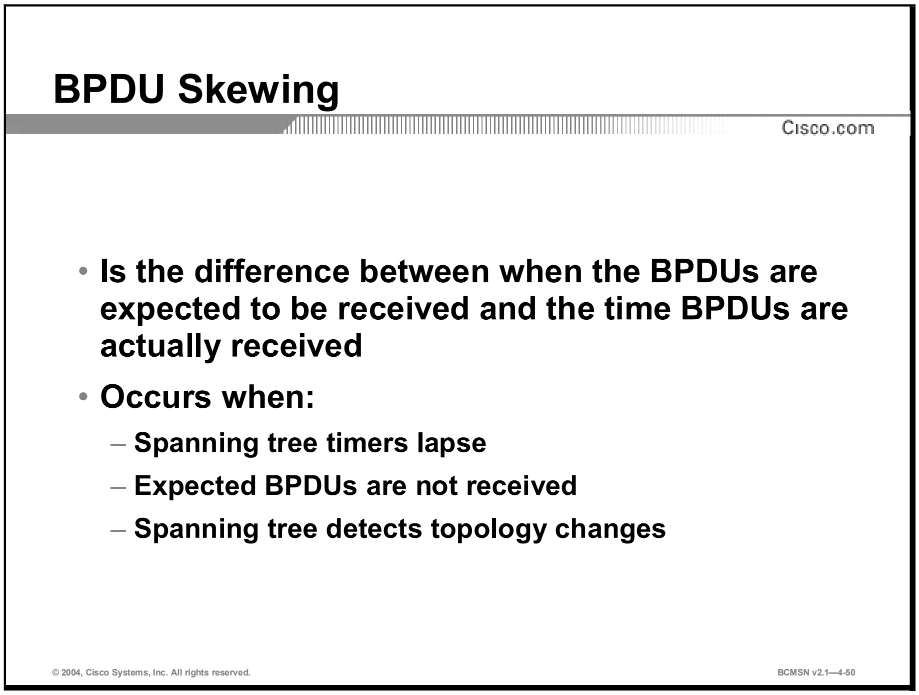

BPDU Skewing

Cisco.com

- Is the difference between when the BPDUs are expected to be received and the time BPDUs are actually received
- Occurs when:
  - Spanning tree timers lapse
  - Expected BPDUs are not received
  - Spanning tree detects topology changes

BCMSN v2.1—4-50

BPDU skewing is the difference between when the BPDUs are expected to be received and the time BPDUs are actually received. Skewing occurs when the following occurs:

- Spanning tree timers lapse
- Expected BPDUs are not received
- Spanning tree detects topology changes

The skewing causes BPDUs to reflood the network to keep the spanning tree topology database current. The root switch advertises its presence by sending out BPDUs for the configured hello time interval. The nonroot switches receive and process one BPDU during each configured time period. A VLAN might not receive the BPDU as scheduled. If the BPDU is not received on a VLAN at the configured time interval, the BPDU is skewed.

Spanning tree uses the hello time to detect when a connection to the root switch exists through a port and when that connection is lost. In MST, the skew detection is on a per-instance basis.

BPDU skewing detects BPDUs that are not processed in a regular time frame on the nonroot switches in the network. If BPDU skewing occurs, a syslog message is displayed.

The number of syslog messages that are generated may have an impact on the convergence of the network and the CPU utilization of the switch. New syslog messages are not generated as individual messages for every VLAN. This is because the higher the number of syslog messages that are reported, the slower the switching process will be.

To reduce the impact on the switch, the syslog messages are generated 50 percent of the max_age time and are limited at the rate of one for every 60 seconds.

Commands that support the spanning tree BPDU skewing feature perform these functions:

- Allow you to enable or disable BPDU skewing. The default is disabled.
- Modify the **show spantree summary** output to show if the skew detection is enabled and for which VLANs or STP instances the skew was detected.
- Provide a display of the VLAN or STP instance and the port affected by the skew, including the following:
  - The duration (in absolute time) of the last skew
  - The duration (in absolute time) of the worst skew
  - The date and time of the worst duration

To change how spanning tree performs BPDU skewing statistics gathering, enter the **set spantree bpdu-skewing** command. The **bpdu-skewing** command is disabled by default.

Using Root Guard

This topic identifies the function of root guard.

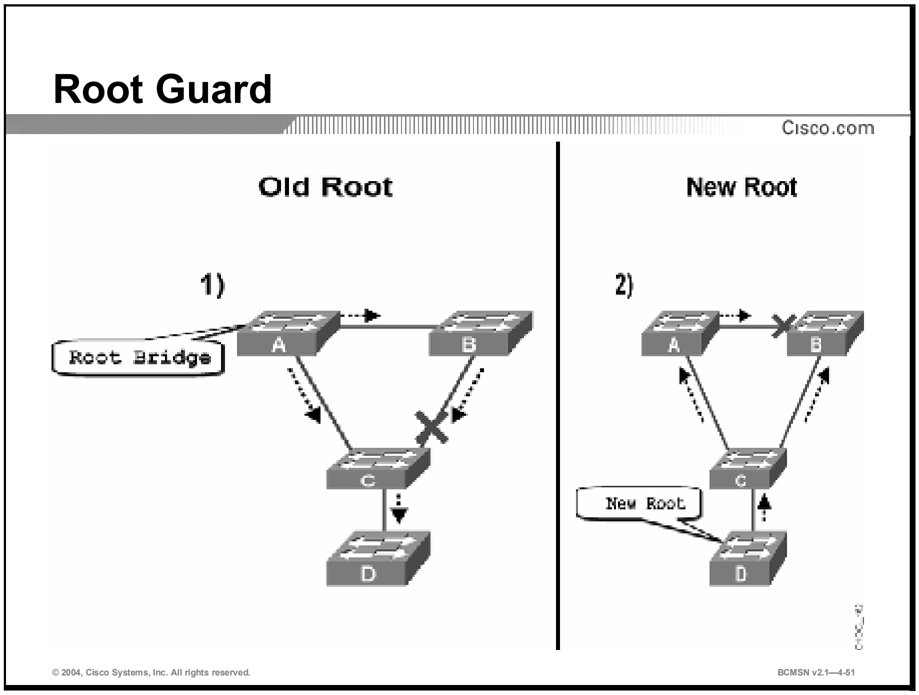

Root Guard

Cisco.com

BCMSN v2.1—4-51

Root guard is designed to provide a way to enforce the root bridge placement in the network. The root guard ensures that the port on which it is enabled is the designated port. (Normally, root bridge ports are all designated, unless two or more ports of the root bridge are connected together.) If the bridge receives superior STP BPDUs on a root guard-enabled port, this port will be moved to a root-inconsistent STP state (effectively equal to a listening state), and no traffic will be forwarded across this port. The position of the root bridge will be enforced.

Example: Using Root Guard

The example shows why a rogue root bridge can cause problems on the network and how root guard can help. Switches A and B are the core of the network (see Figure 1). A is the root bridge for a VLAN. Switch C is an access layer switch. The link between B and C is blocking on the C side. The flow of STP BPDUs is shown with arrows.

On the left, device D begins to participate in STP (for example, software-based bridge applications launched on PCs or other switches connected by a customer to a service provider network). If the priority of bridge D is zero or any value lower than that of the root bridge, the software bridge will be elected as a root bridge for this VLAN. This will cause the gigabit link connecting the two core switches to block, thus causing all the data in that particular VLAN to flow via a 100 Mbps link across the access layer. If this link cannot accommodate the data flow, some frames will be dropped, causing performance loss or a connectivity outage.

The root guard feature has been designed to protect the network against such issues. Root guard is configured on a per-port basis, and it does not allow the port to become an STP root port. This means that the port is always STP-designated. If there is a better BPDU received on this port, root guard disables the port, rather than taking the BPDU into account and electing a new STP root.

Root guard must be enabled on all ports where the root bridge should not appear. In the example, the root guard on the left should be enabled on switches A, B, and C on the ports as follows:

- Switch A: port connecting to switch C
- Switch B: port connecting to switch C
- Switch C: port connecting to switch D

In the configuration, switch C will block the port connecting to switch D after switch C receives a superior BPDU. The port is put in a special STP state (root-inconsistent), which is the same as the listening state. No traffic will pass through the port in this state.

When switch D stops sending superior BPDUs, the port will be unblocked again and will move, via STP, into states of listening, learning, and eventually transition to the forwarding state. Recovery is automatic; no human intervention is required.

The following message is printed when root guard blocks a port:

```
%SPANTREE-2-ROOTGUARDBLOCK: Port 1/1 tried to become non-
designated in VLAN 77. Moved to root-inconsistent state
```

Configuring Root Guard

This topic identifies the command used to configure root guard.

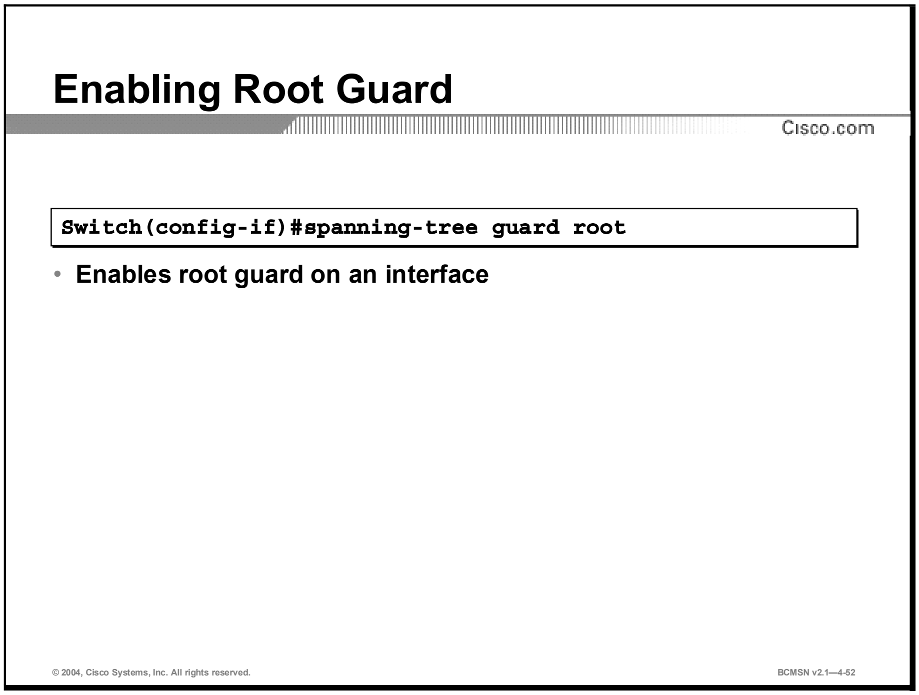

Enabling Root Guard

Cisco.com

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

- **Enables root guard on an interface**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-52

To enable root guard on a Layer 2 access port (to force it to become a designated port), or to disable root guard, use this command:

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

Example: Enabling Root Guard

This example shows how to enable root guard on Fast Ethernet interface 5/8:

```
Switch(config)#interface fastethernet 5/8
Switch(config-if)#spanning-tree guard root
Switch(config-if)#end
```

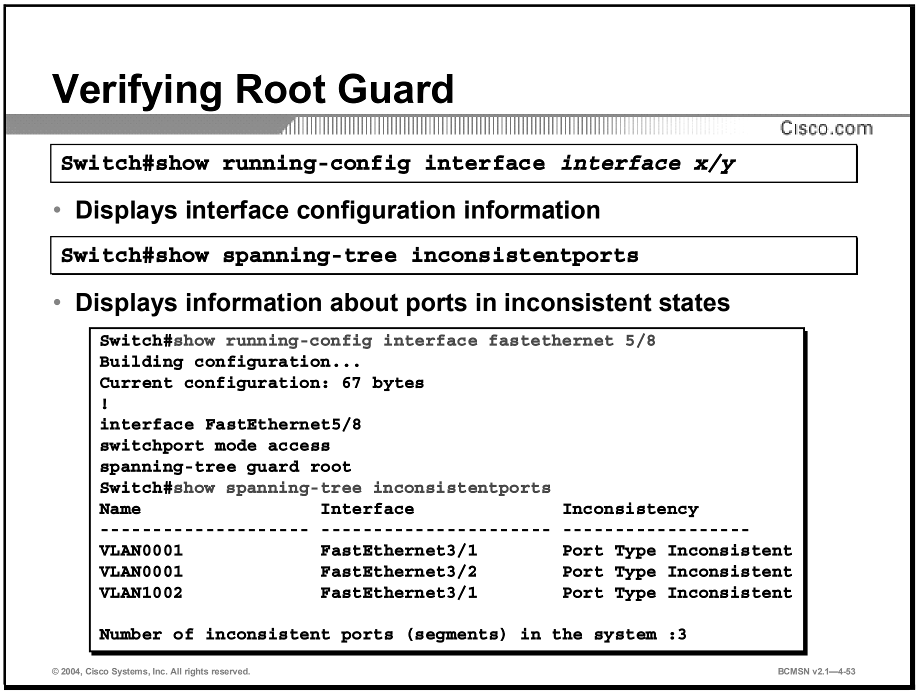

Verifying Root Guard

Cisco.com

```
Switch#show running-config interface interface x/y
```

- Displays interface configuration information

```
Switch#show spanning-tree inconsistentports
```

- Displays information about ports in inconsistent states

```
Switch#show running-config interface fastethernet 5/8
Building configuration...
Current configuration: 67 bytes
!
interface FastEthernet5/8
switchport mode access
spanning-tree guard root
Switch#show spanning-tree inconsistentports
Name                 Interface              Inconsistency
-------------------- ---------------------- ------------------
VLAN0001             FastEthernet3/1        Port Type Inconsistent
VLAN0001             FastEthernet3/2        Port Type Inconsistent
VLAN1002             FastEthernet3/1        Port Type Inconsistent

Number of inconsistent ports (segments) in the system :3
```

BCMSN v2.1—4-53

Example: Verifying Root Guard

This example shows how to verify root guard configuration:

```
Switch#show running-config interface fastethernet 5/8
Building configuration...
Current configuration: 67 bytes
!
interface FastEthernet5/8
switchport mode access
spanning-tree guard root
```

This example shows how to determine whether any ports are in a root-inconsistent state:

```
Switch#show spanning-tree inconsistentports
Name                 Interface              Inconsistency
-------------------- ---------------------- ------------------
VLAN0001             FastEthernet3/1        Port Type
Inconsistent
VLAN0001             FastEthernet3/2        Port Type
Inconsistent
VLAN1002             FastEthernet3/1        Port Type
Inconsistent
VLAN1002             FastEthernet3/2        Port Type
Inconsistent
VLAN1003             FastEthernet3/1        Port Type
Inconsistent
VLAN1003             FastEthernet3/2        Port Type
Inconsistent
```

```
VLAN1004            FastEthernet3/1        Port Type
Inconsistent
VLAN1004            FastEthernet3/2        Port Type
Inconsistent
VLAN1005            FastEthernet3/1        Port Type
Inconsistent
VLAN1005            FastEthernet3/2        Port Type
Inconsistent
Number of inconsistent ports (segments) in the system :10
```

Shutting Down Unidirectional Links

This topic identifies the features associated with Unidirectional Link Detection (UDLD) protocol.

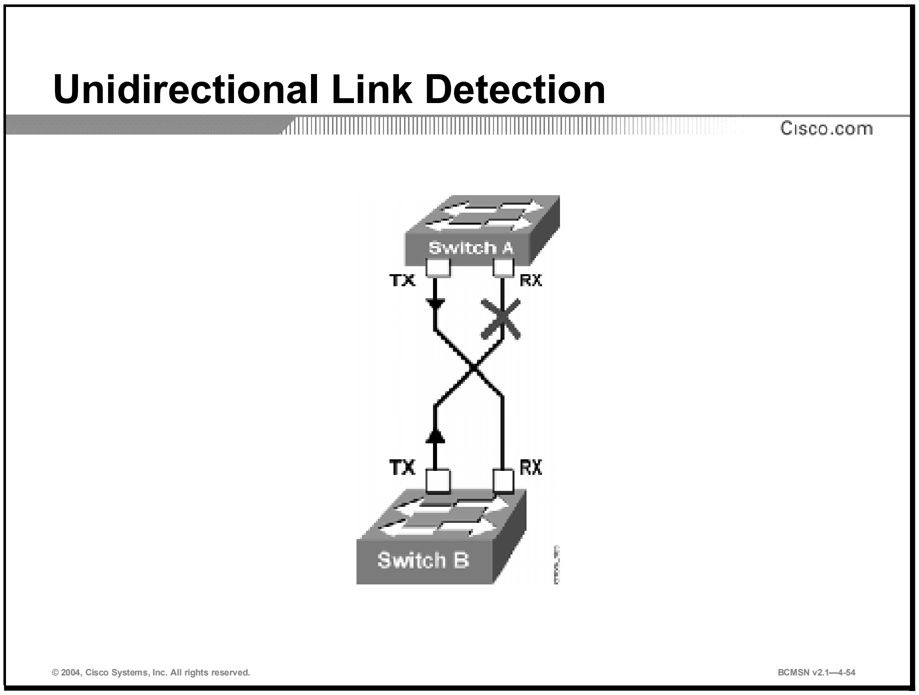

Unidirectional Link Detection

Cisco.com

BCMSN v2.1—4-54

UDLD protocol detects and shuts down unidirectional links. The loop guard STP feature is intended to improve stability of Layer 2 networks.

UDLD allows devices connected through fiber-optic or copper Ethernet cables (for example, Category 5 cabling) to monitor the physical configuration of the cables and detect when a unidirectional link exists. A unidirectional link occurs when traffic transmitted by the local device over a link is received by the neighbor, but traffic transmitted from the neighbor is not received by the local device. When a unidirectional link is detected, UDLD shuts down the affected interface and alerts the user. Unidirectional links can cause a variety of problems, including spanning tree topology loops.

UDLD is a Layer 2 protocol that works with the Layer 1 mechanisms to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected interfaces. When you enable both autonegotiation and UDLD, Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

If one of the fiber strands in a pair is disconnected, as long as autonegotiation is active, the link does not stay up. In this case, the logical link is undetermined, and UDLD does not take any action. If both fibers are working normally from a Layer 1 perspective, UDLD at Layer 2 determines whether those fibers are connected correctly and whether traffic is flowing bidirectionally between the right neighboring devices. Autonegotiation cannot perform this check because autonegotiation operates at Layer 1.

The switch periodically transmits UDLD packets to neighbor devices on interfaces with UDLD enabled. If the packets are echoed back within a specific time frame and they lack a specific acknowledgment (echo), the link is flagged as unidirectional and the interface is shut down. Devices on both ends of the link must support UDLD for the protocol to successfully identify and disable unidirectional links.

Note By default, UDLD is locally disabled on copper interfaces to avoid sending unnecessary control traffic.

The example shows a unidirectional link condition. Switch B successfully receives traffic from switch A on the interface. However, switch A does not receive traffic from switch B on the same interface. UDLD detects the problem and disables the interface.

The table describes the default status for the UDLD states.

Feature	Default Status
UDLD global enable state	Globally disabled
UDLD per-interface enable state for fiber-optic media	Enabled on all Ethernet fiber-optic interfaces
UDLD per-interface enable state for twisted-pair (copper) media	Disabled on all Ethernet 10/100 and 1000BASE-TX interfaces

Configuring UDLD

This topic identifies the command to configure UDLD.

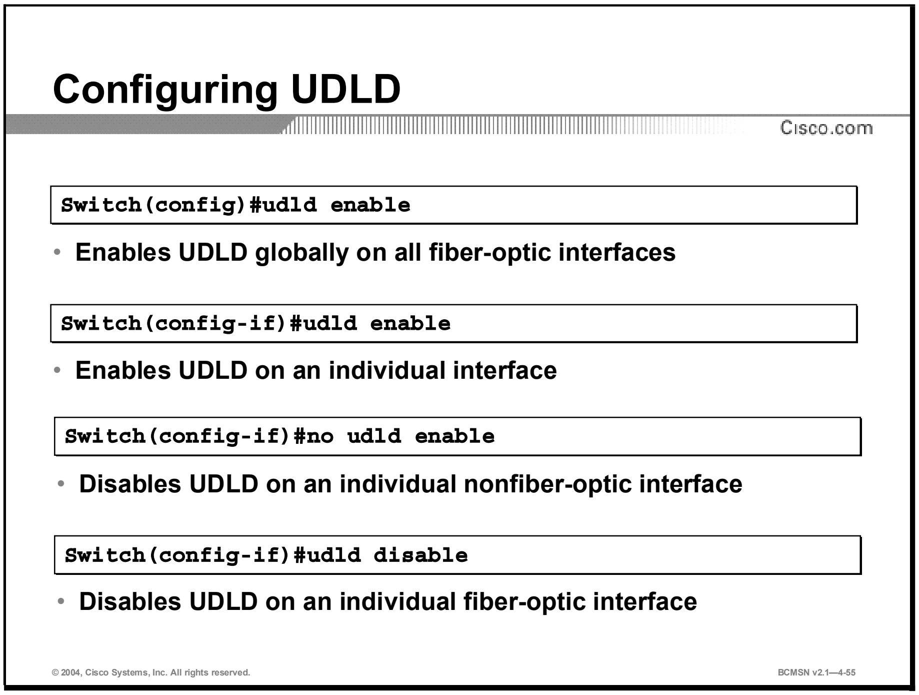

Configuring UDLD

Cisco.com

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

- **Enables UDLD globally on all fiber-optic interfaces**

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

- **Enables UDLD on an individual interface**

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

- **Disables UDLD on an individual nonfiber-optic interface**

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

- **Disables UDLD on an individual fiber-optic interface**

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—4-55

To enable or disable UDLD globally on all fiber-optic interfaces on the switch, enter this command:

```
Switch(config)#[no] udld enable
```

To enable UDLD on individual interfaces, enter this command:

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

To disable UDLD on individual nonfiber-optic interfaces, enter this command:

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

To disable UDLD on individual fiber-optic interfaces, enter this command:

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

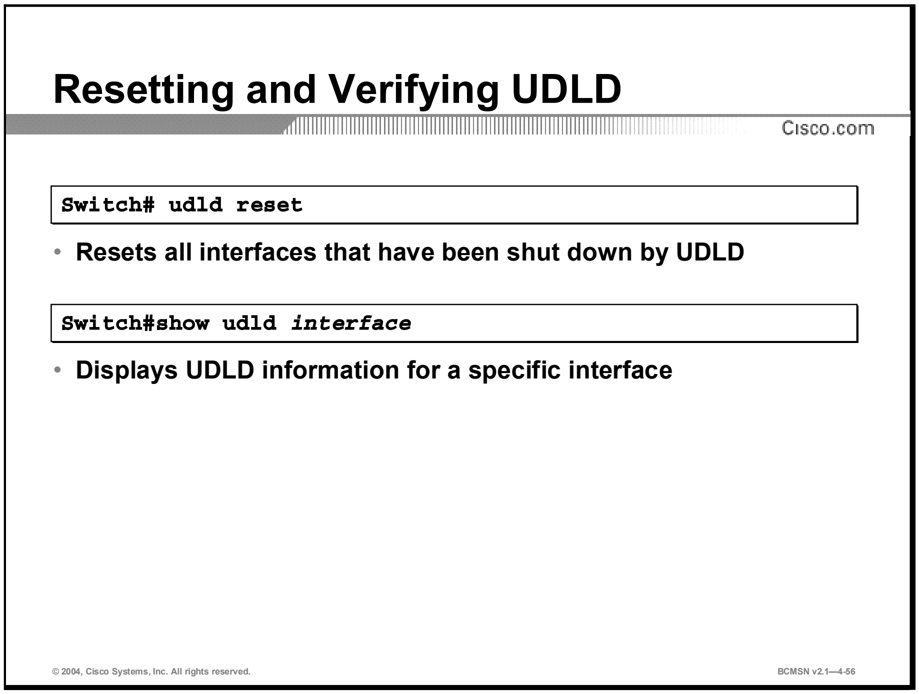

Resetting and Verifying UDLD

Cisco.com

```
Switch# udld reset
```

- **Resets all interfaces that have been shut down by UDLD**

```
Switch#show udld interface
```

- **Displays UDLD information for a specific interface**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-56

To reset all interfaces that have been shut down by UDLD, enter this command:

```
Switch#udld reset
```

To verify the UDLD configuration for an interface, enter this command:

```
Switch#show udld interface
```

Example: Displaying the UDLD State

This example shows how to display the UDLD state for a single interface:

```
Switch#show udld GigabitEthernet2/2

Interface Gi2/2

---

Port enable administrative configuration setting: Follows
device default

Port enable operational state: Enabled

Current bidirectional state: Bidirectional

Current operational state: Advertisement

Message interval: 60

Time out interval: 5

No multiple neighbors detected

    Entry 1

    ---

    Expiration time: 146

    Device ID: 1

    Current neighbor state: Bidirectional

    Device name: 0050e2826000

    Port ID: 2/1

    Neighbor echo 1 device: SAD03160954

    Neighbor echo 1 port: Gi1/1


    Message interval: 5

    CDP Device name: 066527791
```

Preventing Layer 2 Forwarding Loops

This topic identifies the features of loop guard.

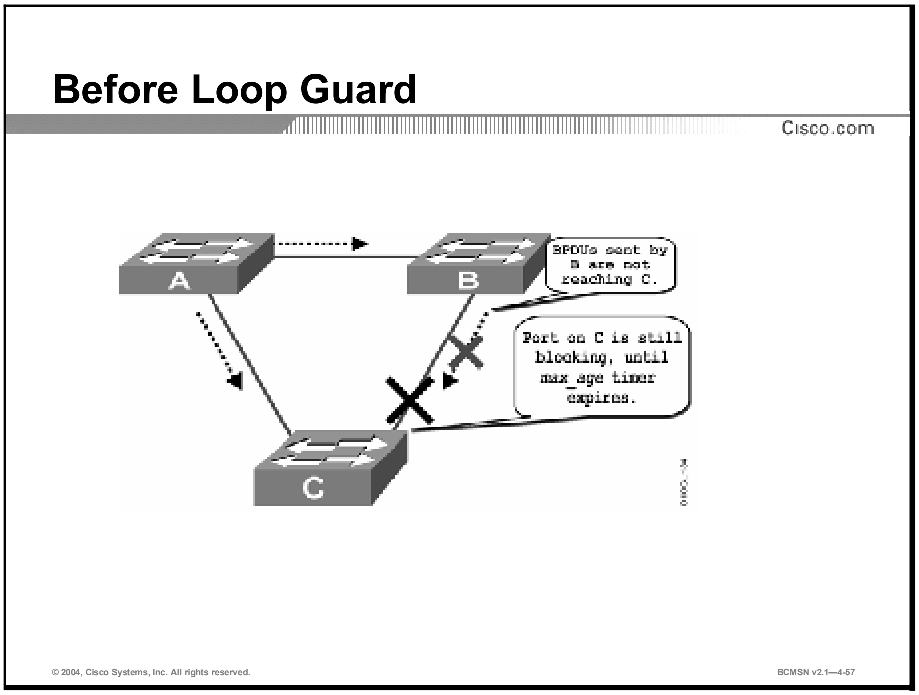

Before Loop Guard

Cisco.com

A
B
C
BPDUs sent by B are not reaching C.
Port on C is still blocking, until max_age timer expires.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-57

Loop guard is intended to provide additional protection against Layer 2 forwarding loops (STP loops). An STP loop is created when an STP blocking port in a redundant topology erroneously transitions to the forwarding state. This usually happens because one of the ports of a physically redundant topology (not necessarily the STP blocking port) stops receiving STP BPDUs. In its operation, STP relies on continuous reception or transmission of BPDUs, depending on the port role (a designated port transmits BPDUs, a nondesignated port receives BPDUs).

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 from the alternate or backup port becomes designated and moves to the forwarding state, thus creating a loop.

With loop guard, an additional check is made. If BPDUs are not received on a nondesignated port and loop guard is enabled, that port will be moved into the STP loop-inconsistent blocking state instead of moving to the forwarding state. Without the loop guard, the port would assume the designated port role. The port would move to STP forwarding state, and thus create a loop.

When loop guard blocks an inconsistent port, this message is logged:

```
SPANTREE-2-LOOPGUARDBLOCK: No BPDUs were received on port 3/2
in vlan 3. Moved to loop-inconsistent state.
```

When the BPDU is received on a port in a loop-inconsistent STP state, the port will transition into another STP state. According to the received BPDU, this means that the recovery is automatic and no intervention is necessary. After the recovery, this message is logged:

```
SPANTREE-2-LOOPGUARDUNBLOCK: port 3/2 restored in vlan 3.
```

Example: Before Loop Guard

In this example, Switch A is the root switch. Due to unidirectional link failure on the link between switch B and switch C, switch C is not receiving BPDUs from B.

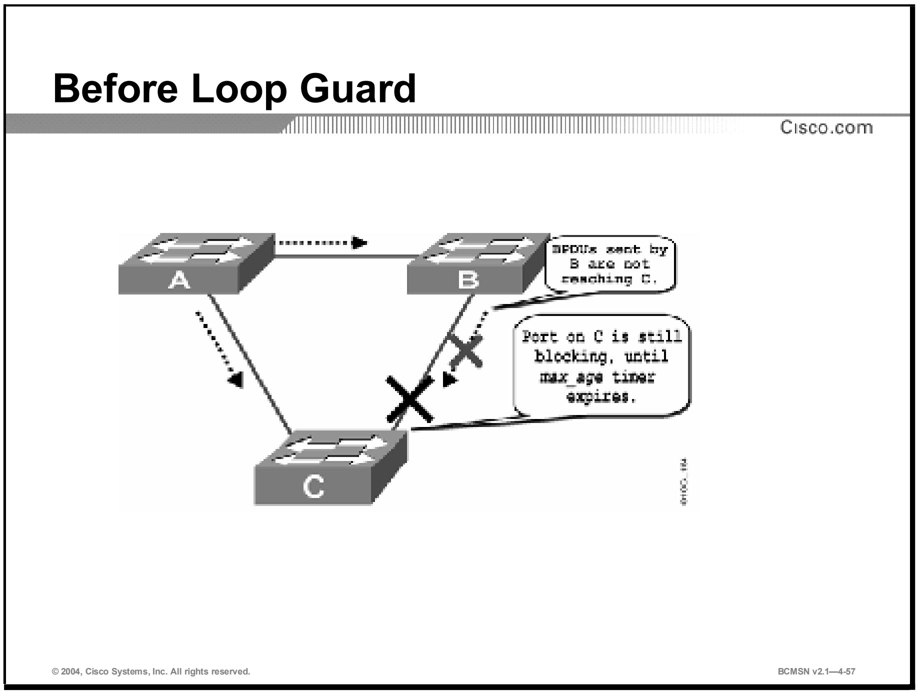

Before Loop Guard

Cisco.com

BPDUs sent by B are not reaching C.

Port on C is still blocking, until max_age timer expires.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-57

Example: With Loop Guard

Without loop guard, the STP blocking port on C will transition to the STP listening state upon max_age timer expiration and then to the forwarding state in two times the forward delay time. A loop will be created.

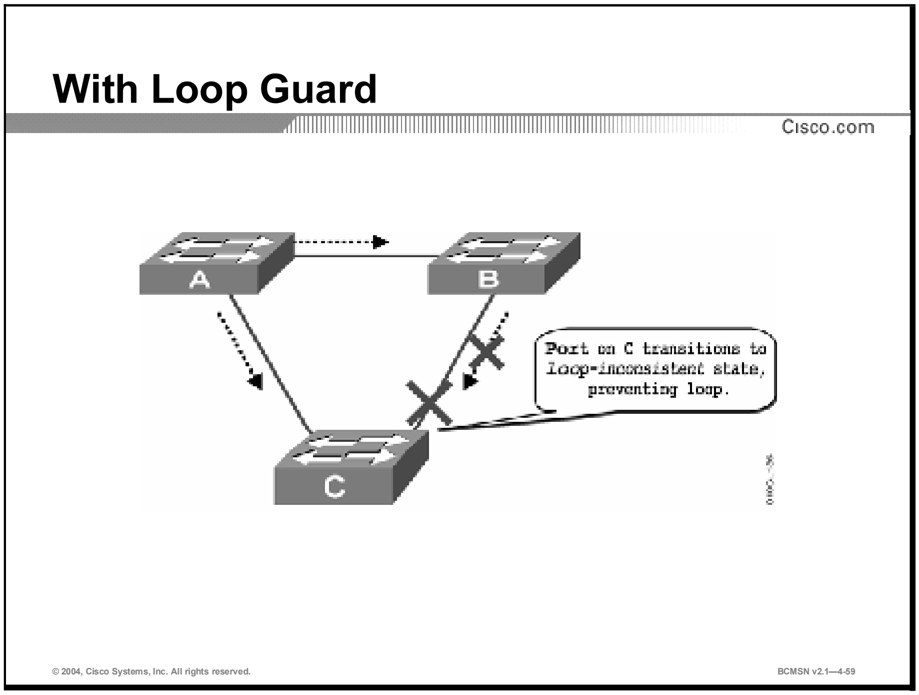

With Loop Guard

Cisco.com

Port on C transitions to loop-inconsistent state, preventing loop.

BCMSN v2.1—4-59

With loop guard enabled, the blocking port on switch C will transition into the STP loop-inconsistent state upon expiration of the max_age timer. Because a port in the STP loop-inconsistent state will not pass user traffic, no loop is created. The loop-inconsistent state is effectively equal to the blocking state.

Configuring Loop Guard

This topic identifies the command to configure loop guard.

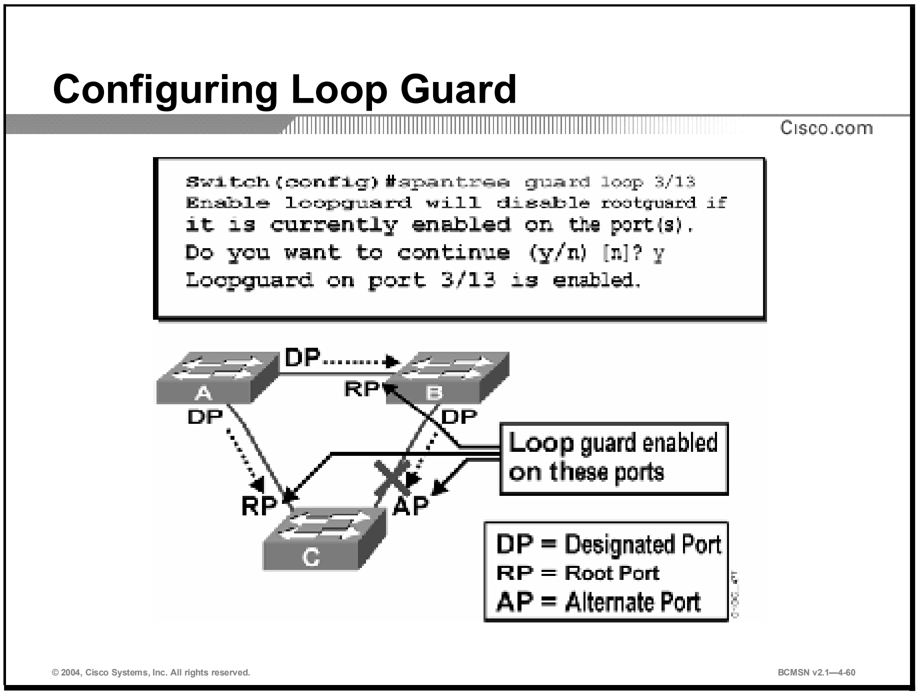

Configuring Loop Guard

Cisco.com

```
Switch(config)#spantree guard loop 3/13
Enable loopguard will disable rootguard if
it is currently enabled on the port(s).
Do you want to continue (y/n) [n]? y
Loopguard on port 3/13 is enabled.
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-60

Loop guard is enabled on a per-port basis. As long as loop guard blocks the port on the STP level, loop guard blocks inconsistent ports on a per-VLAN basis. If loop guard is enabled on an EtherChannel interface, the entire channel will be blocked for a particular VLAN. This is because EtherChannel is regarded as one logical port from an STP point of view.

Loop guard must be enabled on the root and alternate ports for all possible combinations of active topologies. As long as loop guard is not a per-VLAN feature, the same port might be designated for one VLAN and nondesignated for another VLAN. The possible failover scenarios should also be taken into account.

Example: Configuring Loop Guard

In the example, loop guard is disabled by default. This command is used to enable loop guard:

```
Switch(config)#spantree guard loop mod/port
```

For example:

```
Switch(config)#spantree guard loop 3/13
```

Enabling loop guard will disable root guard, if root guard is currently enabled on the ports.

You can enable loop guard globally on all ports. Loop guard is enabled on all point-to-point links. The point-to-point link is detected by the duplex status of the link. If the link is full-duplex, then the link is considered point-to-point. It is still possible to configure global settings on a per-port basis.

To enable loop guard globally, issue this command:

```
Switch(config)#spantree global-default loopguard enable
```

To disable loop guard, issue this command:

```
Switch(config)#spantree guard none mod/port
```

To globally disable loop guard, issue this command:

```
Switch(config)#spantree global-default loopguard disable
```

To verify the loop guard status, issue this command:

```
Switch#show spantree guard mod/port | vlan
```

Example: Verifying the Loop Guard Status

This example shows how to verify the loop guard status:

```
Switch#show spantree guard 3/13

Port                     VLAN Port-State    Guard Type

                       ------------------------ ---- ---------
---- ----------

                       3/13                     2
forwarding          loop

Switch#
```

Differentiating Between Loop Guard and UDLD

This topic discusses loop guard and UDLD features.

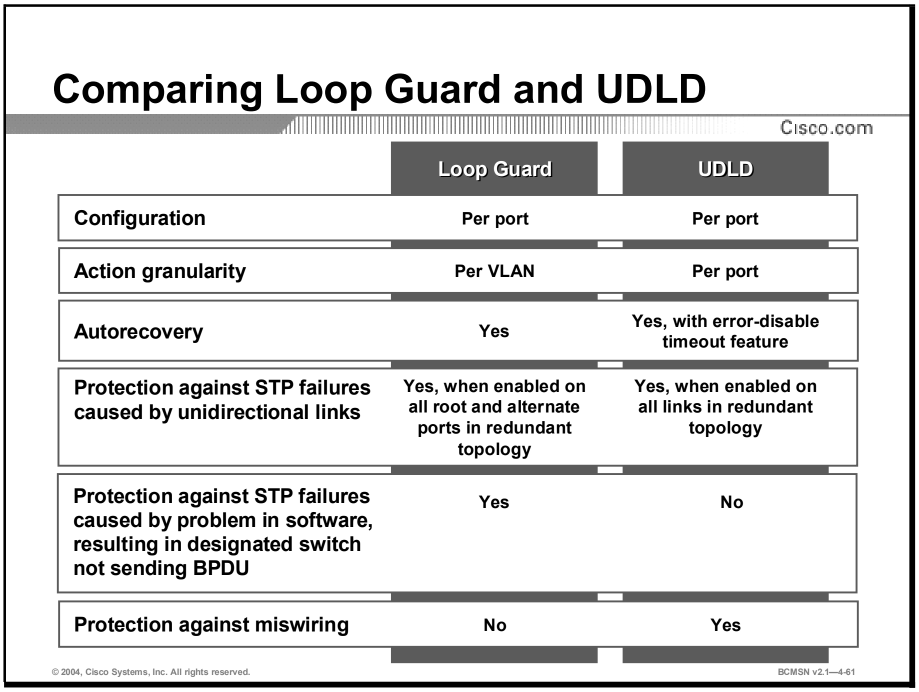

Comparing Loop Guard and UDLD

	Loop Guard	UDLD
Configuration	Per port	Per port
Action granularity	Per VLAN	Per port
Autorecovery	Yes	Yes, with error-disable timeout feature
Protection against STP failures caused by unidirectional links	Yes, when enabled on all root and alternate ports in redundant topology	Yes, when enabled on all links in redundant topology
Protection against STP failures caused by problem in software, resulting in designated switch not sending BPDU	Yes	No
Protection against miswiring	No	Yes

The loop guard and UDLD functionality partially overlap in that both protect against STP failures caused by unidirectional links. These two features are different in their approach to the problem and also in functionality.

Example: Differentiating Between Loop Guard and UDLD

Depending on various design considerations, you can choose either UDLD or loop guard. In terms of STP, the most noticeable difference between the two features is the absence of protection in UDLD against STP failures caused by problems in software that result in the designated switch not sending BPDUs. These types of failure are much less common than those caused by unidirectional links. In return, UDLD might be more flexible when there are unidirectional links on EtherChannel. In this case, UDLD will only disable failed links, and the channel should remain functional with remaining links. Loop guard will block the whole channel in such a failure by putting it into a loop-inconsistent state.

Additionally, loop guard does not work on shared links or in situations in which the link has been unidirectional since the linkup. In this last case, the port never receives BPDUs and is, therefore, a designated port. Because this could be normal behavior, loop guard does not apply in this case. UDLD does provide protection against such a scenario.

Enabling both UDLD and loop guard provides the highest level of protection.

Applying Protection Features to the Enterprise Composite Network Model

This topic identifies the appropriate protection features used in the Enterprise Composite Network model.

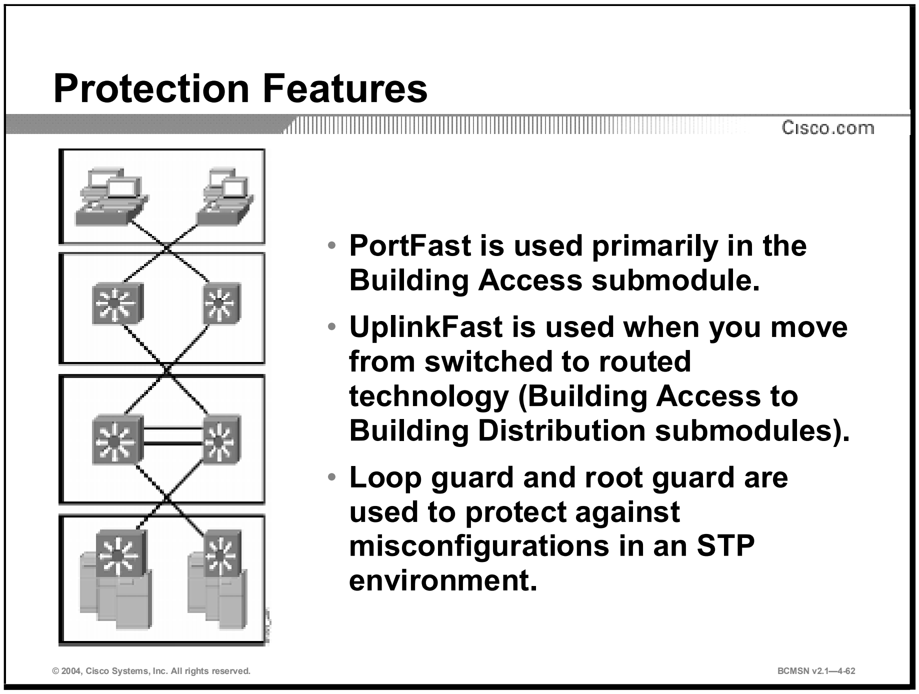

Protection Features

Cisco.com

- PortFast is used primarily in the Building Access submodule.
- UplinkFast is used when you move from switched to routed technology (Building Access to Building Distribution submodules).
- Loop guard and root guard are used to protect against misconfigurations in an STP environment.

BCMSN v2.1—4-62

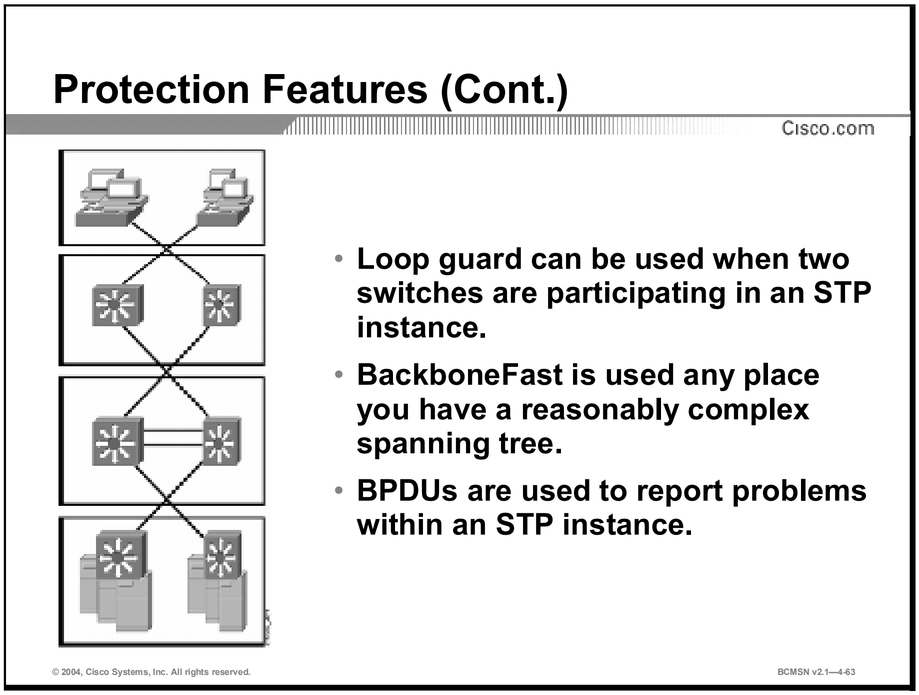

Protection Features (Cont.)

Cisco.com

- Loop guard can be used when two switches are participating in an STP instance.
- BackboneFast is used any place you have a reasonably complex spanning tree.
- BPDUs are used to report problems within an STP instance.

BCMSN v2.1—4-63

BPDU guard and BPDU filtering are used in conjunction with PortFast on switch ports that connect to those edge devices not participating in the STP instance. These ports are typically found in data link and multilayer switches participating in the Building Access submodules of the Enterprise Composite Network model. The use of multilayer switching (routing) in the Building Distribution and Campus Backbone submodules eliminates the need for STP and STP extensions in these modules.

BPDU skewing enables enhanced reporting of STP performance. Current best design practice limits a VLAN instance to a single switch. This eliminates the need for STP and STP extensions.

For network installations in which all ports in a VLAN instance are confined to a single device, it is not necessary to use loop guard and root guard to enforce a specific STP topology or to use loop guard and UDLD to detect media failures that affect STP.

Summary

This topic summarizes the key points discussed in this lesson.

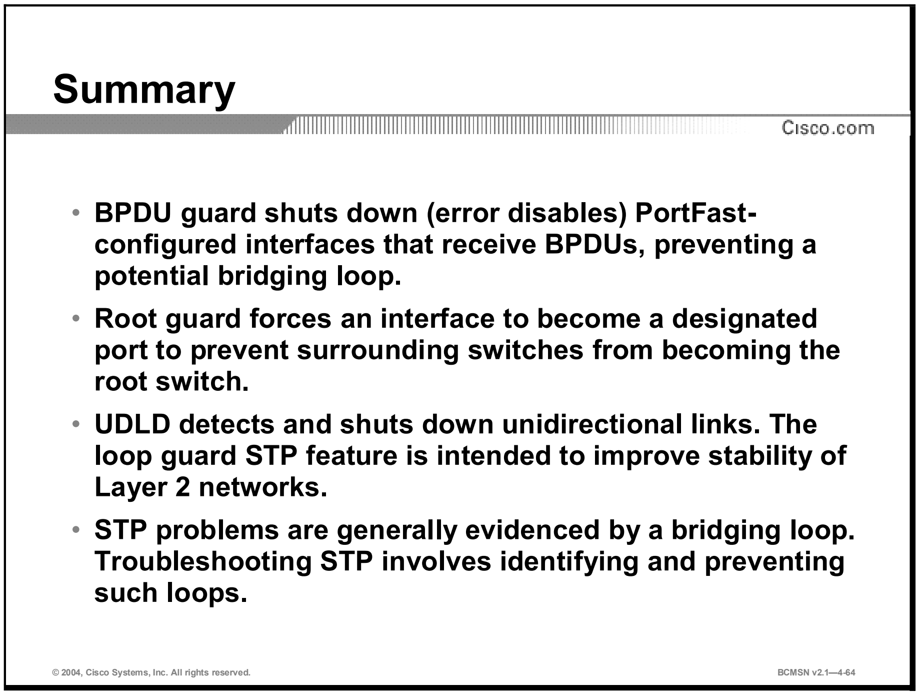

Summary

- BPDU guard shuts down (error disables) PortFast-configured interfaces that receive BPDUs, preventing a potential bridging loop.
- Root guard forces an interface to become a designated port to prevent surrounding switches from becoming the root switch.
- UDLD detects and shuts down unidirectional links. The loop guard STP feature is intended to improve stability of Layer 2 networks.
- STP problems are generally evidenced by a bridging loop. Troubleshooting STP involves identifying and preventing such loops.

BCMSN v2.1—4-64

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are in the Quiz Answer Key.

Q1) Select the command that correctly enables BPDU guard.

A) switch(config)#**spanning-tree bpduguard**

B) switch(config-if)#**spanning-tree bpduguard**

C) switch(config)#**spanning-tree portfast bpduguard**

D) switch(config-if)#**spanning-tree portfast bpduguard**

Q2) Select the command that correctly enables BPDU filtering.

A) switch(config)#**spanning-tree bpdufilter**

B) switch(config)#**spanning-tree default bpdufilter**

C) switch(config-if)#**spanning-tree bpdufilter default**

D) switch(config-if)#**spanning-tree default**

Q3) Select the possible causes of BPDU skewing. (Choose two.)

A) unidirectional links

B) failed EtherChannel link

C) unexpected BPDUs received

D) lapsed spanning tree timers

E) expected BPDUs not received

F) STP detected topology changes

Q4) Select the correct features of root guard. (Choose two.)

A) configured for every two ports

B) forwards across a port if the bridge receives superior STP BPDUs

C) enforces root bridge placement in the network

D) does not allow the port to become an STP root port

Q5) Select the correct command to configure root guard.

A) switch(config)#**spanning-tree guard root**

B) switch(config-if)#**spanning-tree guard root**

C) switch(config-if)#**spanning-tree root guard**

D) switch(config)#**spanning-tree root guard**

Q6) Select the correct features of Unidirectional Link Detection Protocol. (Choose two.)

A) locally enabled on copper interfaces

B) Layer 2 protocol that works with the Layer 1 mechanisms to determine the physical status of a link

C) detects the identities of neighbors and shuts down misconnected interfaces

D) takes action when a logical link is undetermined

Q7) Select the correct command to configure UDLD on individual interfaces.

A) switch(config-if)#**udld enable**

B) switch(config)#**[no] udld enable**

C) switch(config-if)#**[no] udld enable**

D) switch(config)#**udld enable**

Q8) Select the correct features of loop guard. (Choose two.)

A) provides additional protection against Layer 2 forwarding loops (STP loops)

B) allows blocking port in a physically redundant topology to stop receiving BPDUs

C) enables the blocking port to move to a forwarding state

D) moves ports into the STP loop-inconsistent blocking if BPDUs are not received on a non-designated port

Q9) Select the correct command to configure loop guard.

A) switch(config)#**spantree loop guard** *mod/port*

B) switch(config)#**spantree guard loop** *mod/port*

C) switch(config-if)#**spantree guard loop** *mod/port*

D) switch(config-if)#**spantree loop guard** *mod/port*

Q10) Match the correct features with either loop guard or UDLD.

_____ 1. loop guard

_____ 2. UDLD

A) protects against STP failures caused by software problems

B) protects against miswiring

C) takes action per port

D) protects against STP failures when enabled on all root and alternate ports

Q11) Select the correct protection features that are used in the Enterprise Composite Network model. (Choose two.)

A) BPUD guard is used in conjunction with BackboneFast to eliminate the need for STP.

B) BPDU skewing enhances reporting of STP performance by limiting a VLAN instance to a single switch.

C) Loop guard detects media failures.

D) UDLD eliminates the need for STP extensions in the Building Distribution and Campus Backbone submodules.

Quiz Answer Key

Q1) C

Relates to: Configuring BPDU Guard

Q2) C

Relates to: Enabling BPDU Filtering

Q3) D, E, F

Relates to: Enabling BPDU Skewing Detection

Q4) C, D

Relates to: Using Root Guard

Q5) B

Relates to: Configuring Root Guard

Q6) B, C

Relates to: Shutting Down Unidirectional Links

Q7) A

Relates to: Configuring UDLD

Q8) A, D

Relates to: Preventing Layer 2 Forwarding Loops

Q9) B

Relates to: Configuring Loop Guard

Q10) 1=A, D; 2=B, C

Relates to: Differentiating Between Loop Guard and UDLD

Q11) B, C

Relates to: Applying Protection Features to the Enterprise Composite Network Model

Selecting a Troubleshooting Approach

Overview

Troubleshooting STP involves identifying and preventing bridging loops. Most of the STP troubleshooting steps are simply using **show** commands to try to identify error conditions. Knowledge of the network helps focus your attention on the critical ports on the key devices.

Relevance

Because smooth functioning of networks is so critical to business today, being able to rapidly detect and solve problems is a primary consideration for network administrators.

Objectives

Upon completing this lesson, you will be able to:

- Select the most common problems that can occur in STP
- Identify a scenario that best represents duplex mismatching
- Identify the scenario that best represents unidirectional link failure
- Identify the scenario that best represents a PortFast configuration error
- Select an approach for troubleshooting STP so that the problem can be resolved effectively
- Identify the commands that display the debug messages

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Identifying STP Problems
- Identifying Duplex Mismatching
- Identifying Unidirectional Link Failure
- Identifying PortFast Configuration Error
- Troubleshooting the Spanning Tree Protocol
- Identifying Spanning Debug Commands
- Summary
- Quiz

Identifying STP Problems

This topic identifies the most common STP problems.

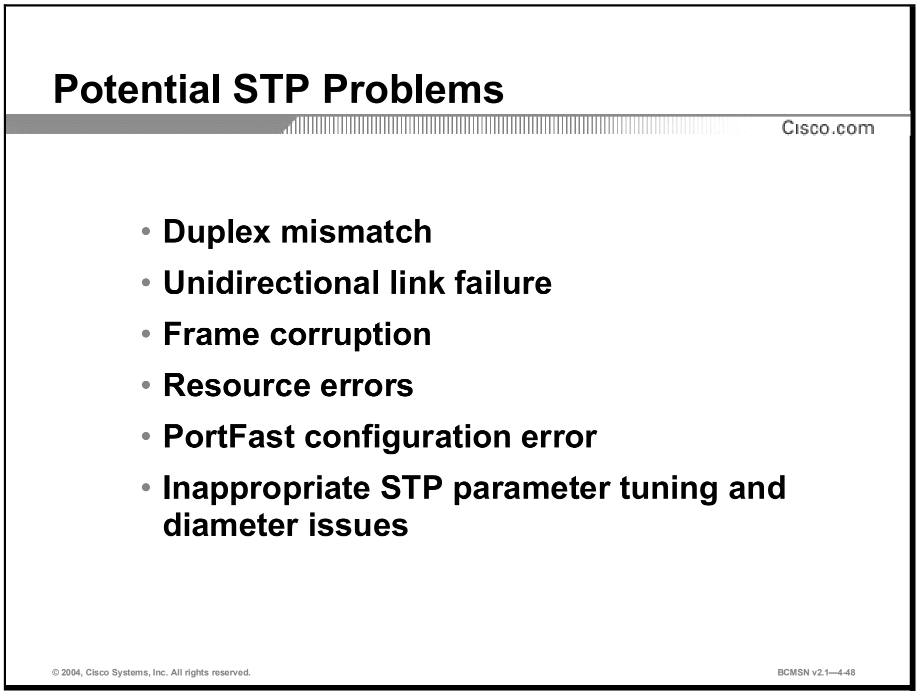

Potential STP Problems

Cisco.com

- Duplex mismatch
- Unidirectional link failure
- Frame corruption
- Resource errors
- PortFast configuration error
- Inappropriate STP parameter tuning and diameter issues

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-48

STP problems are generally evidenced by a bridging loop. Troubleshooting STP involves identifying and preventing such loops.

The primary function of the STA is to cut loops created by redundant links in bridged networks. The STP operates at Layer 2 of the Open System Interconnection (OSI) model. BPDUs are exchanged between bridges, thus electing the ports that will eventually forward or block traffic. This protocol can fail in some specific cases. Troubleshooting the resulting situation can be very difficult, depending on the design of the network.

Identifying Duplex Mismatching

This topic identifies the scenario that best represents duplex mismatching.

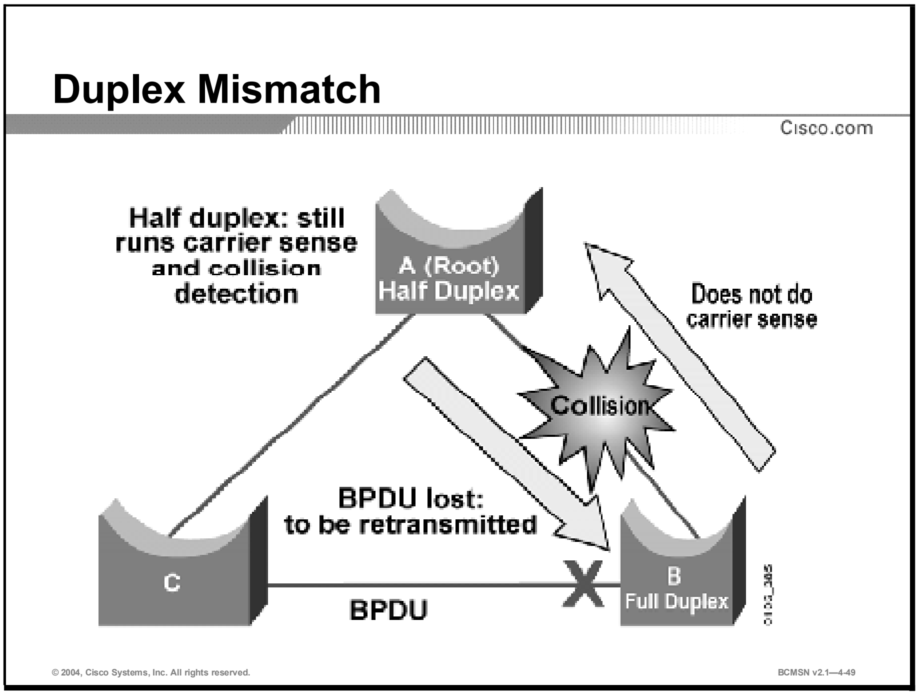

Duplex Mismatch

Cisco.com

Half duplex: still runs carrier sense and collision detection

A (Root) Half Duplex

Does not do carrier sense

Collision

BPDU lost: to be retransmitted

C

BPDU

X

B Full Duplex

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-49

Duplex mismatch on a point-to-point link is a common configuration error. When one side of the link is configured for full duplex and the other side is configured for autonegotiation, the port configured for autonegotiation will end up in half-duplex mode. Manually setting duplexity disables autonegotiation. The rule of autonegotiation is that upon failure, a port is required to assume half-duplex settings. The link will work but will have a very high error rate for frames.

The worst-case scenario is when a bridge sending BPDUs is configured for half duplex on a link, while the peer bridge is configured for full duplex. In the example, the duplex mismatch on the link between bridge A and B can easily lead to a bridging loop. Because B is configured for full duplex, it does not perform carrier sense when accessing the link. B will then start sending frames even if A is already using the link. This is a problem for A, which detects a collision and runs the backoff algorithm before attempting another transmission of its frame. The result is that, if there is enough traffic from B to A, every single packet (including the BPDUs) sent by A will be deferred or collisioned and eventually dropped. From an STP point of view, because it does not receive BPDUs from A anymore, bridge B lost its root. This leads B to unblock its port to C, thereby creating the loop.

Identifying Unidirectional Link Failure

This topic identifies the scenario that best represents unidirectional link failure.

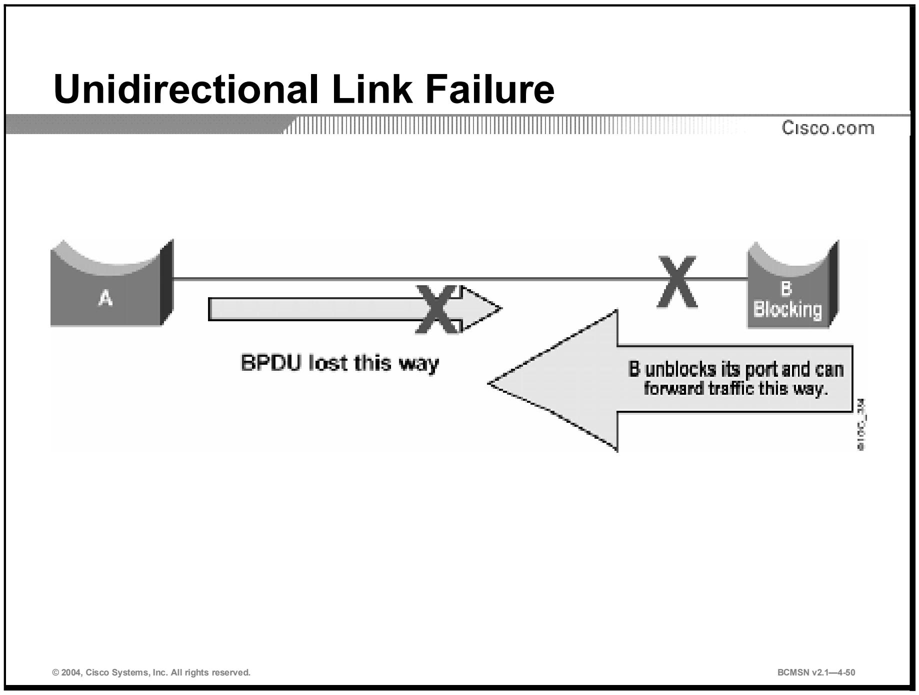

Unidirectional Link Failure

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-50

A unidirectional link is a very frequent cause for a bridging loop. Unidirectional links are often caused by an undetected failure on a fiber link or a problem with a transceiver. Anything that can cause a link to stay up while providing a one-way communication is very dangerous as far as STP is concerned.

Cisco introduced the UDLD protocol on high-end switches. This feature is able to detect wrong cabling or unidirectional links on Layer 2, and it will automatically break resulting loops by disabling some ports. It is a good idea to run UDLD in a bridged environment, wherever possible.

Example: Unidirectional Link Failure

In the example, the link between bridge A and bridge B is unidirectional and drops traffic from A to B while transmitting traffic from B to A. Suppose that the bridge B port should be blocking. A port can block only if it receives BPDUs from a bridge that has a higher priority. In this case, all the BPDUs coming from bridge A are lost and bridge B eventually forwards traffic, creating a loop. If the failure exists at startup, the STP will not converge correctly and rebooting the bridges will have absolutely no effect.

Frame Corruption

Frame corruption can also lead to an STP failure. If a link is experiencing a high rate of physical errors, a certain number of consecutive BPDUs could be lost, leading a blocking port to transition to the forwarding state. This is fairly rare because STP default parameters are very conservative. The blocking port would need to miss its BPDUs for 50 seconds before transitioning to the forwarding state, and a single BPDU successfully transmitted would break the loop. This situation usually occurs when STP parameters have been adjusted inappropriately (the max_age reduced, for example).

The causes of frame corruption are duplex mismatch, bad cables, or incorrect cable length.

Resource Errors

STP is implemented in software, even on high-end switches that perform most of their switching functions in hardware with specialized application-specific integrated circuits (ASICs). This means that if the CPU of the bridge is overutilized for any reason, it may lack the resources to send out BPDUs. The STA is generally not very processor-intensive and has priority over other processes. Therefore, a resource problem is very unlikely to arise.

Identifying PortFast Configuration Error

This topic identifies the scenario that best represents a PortFast configuration error.

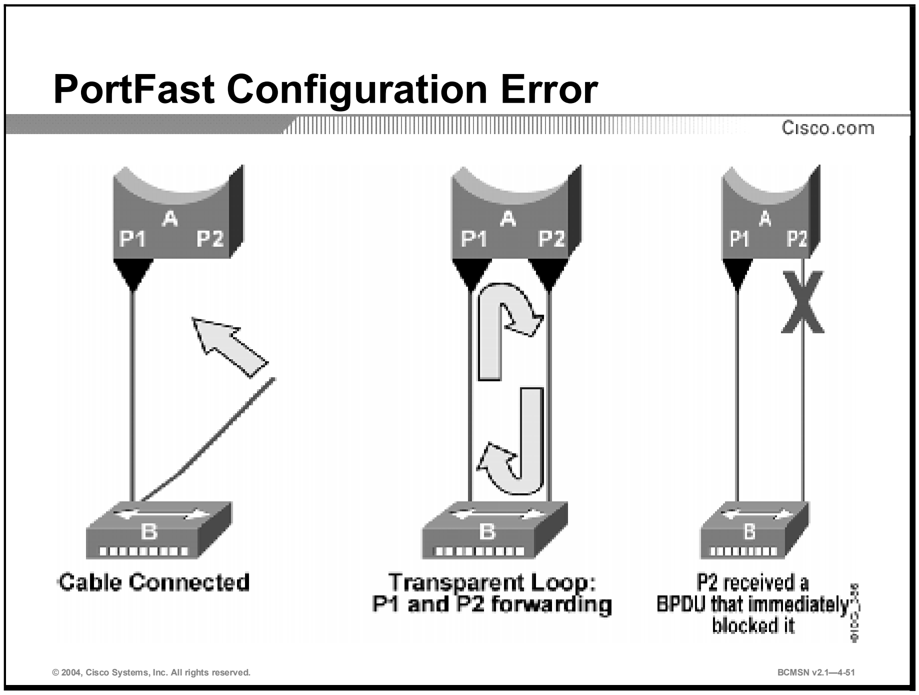

PortFast Configuration Error

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-51

PortFast is a feature that is typically enabled for a port connected to a host. When the link comes up on such a port, the first stages of the STA are skipped and the port directly transitions to the forwarding state. This can be dangerous when not used correctly.

Example: PortFast Configuration Error

A is a bridge with port P1 already forwarding and port P2 configured for PortFast. B is a hub. As soon as the second cable is plugged into A, port P2 goes to the forwarding state and creates a loop between P1 and P2. This will stop as soon as P1 or P2 receives a BPDU that will put one of these two ports in blocking state. The problem with this kind of transient loop is that if the looping traffic is very intensive, the bridge may have trouble successfully sending the BPDU that will stop the loop. This can delay the convergence considerably. Implementing BPDU guard will prevent this problem.

Inappropriate STP Parameter Tuning and Diameter Issues

An aggressive value for the max_age parameter and the forward_delay parameter can lead to a very unstable STP. The loss of some BPDUs can cause a loop. Another potential problem is related to the diameter of the bridged network. The conservative default values for STP impose a maximum network diameter of seven. This means that two distinct bridges in the network should not be more than seven hops away from one another. Part of this restriction comes from the age field carried by BPDUs. When a BPDU is propagated from the root bridge toward the leaves of the tree, the age field is incremented each time it goes though a bridge. Eventually, when the age field of a BPDU goes beyond max_age, it is discarded. Typically, this will occur if the root is too far away from some bridges of the network. This issue will have an impact on the convergence of spanning tree.

Troubleshooting the Spanning Tree Protocol

This topic identifies an approach for troubleshooting STP so that the problem can be resolved effectively.

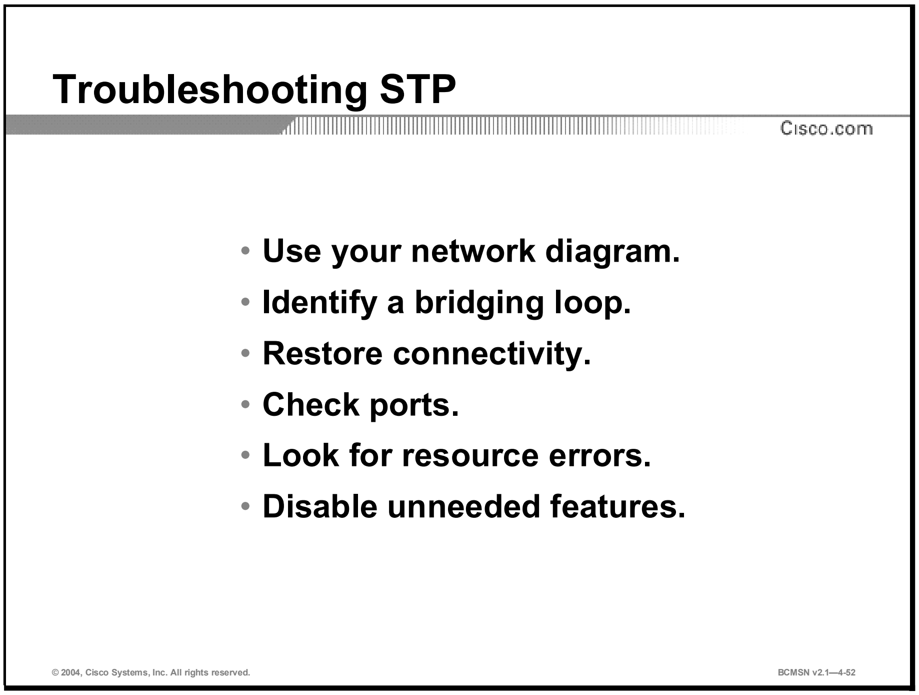

Troubleshooting STP

Cisco.com

- Use your network diagram.
- Identify a bridging loop.
- Restore connectivity.
- Check ports.
- Look for resource errors.
- Disable unneeded features.

BCMSN v2.1—4-52

You need to know some basic things about your network before troubleshooting a bridging loop. You need to know at least the following:

- The topology of the bridged network
- Where the root bridge is located
- Where the blocked ports (and therefore the redundant links) are located

This knowledge is essential for at least for the following two reasons:

- To identify a problem, you need to know how the network should look when it is operating correctly.
- Most of the STP troubleshooting steps are simply using **show** commands to try to identify error conditions. Knowledge of the network helps focus your attention on the critical ports on the key devices.

Identifying Spanning Debug Commands

This topic identifies the correct commands that display the debug messages.

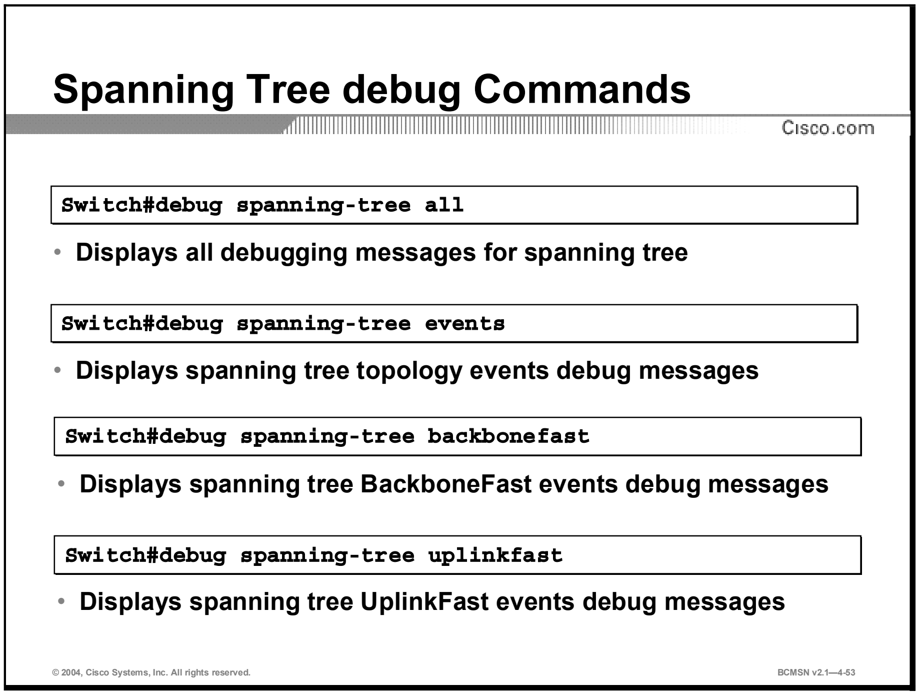

Spanning Tree debug Commands

Cisco.com

```
Switch#debug spanning-tree all
```

- Displays all debugging messages for spanning tree

```
Switch#debug spanning-tree events
```

- Displays spanning tree topology events debug messages

```
Switch#debug spanning-tree backbonefast
```

- Displays spanning tree BackboneFast events debug messages

```
Switch#debug spanning-tree uplinkfast
```

- Displays spanning tree UplinkFast events debug messages

BCMSN v2.1—4-53

The **debug spanning-tree** command is used to troubleshoot spanning-tree activities.

Identify a Bridging Loop

The best way to identify a bridging loop is to capture the traffic on a saturated link and to check whether similar packets are seen multiple times. If all users in a specific bridging domain have connectivity issues at the same time, you should suspect a bridging loop. Check the port utilization on your devices and look for abnormal values.

You can monitor BPDUs by using the **debug spanning-tree** command. This command is helpful in verifying correct bridging operation as well as troubleshooting bridging and spanning tree operation in bridged and switched networks.

Restore Connectivity Quickly

Bridging loops have severe consequences on a bridged network. Administrators generally do not have time to look for the reason for the loop. Instead, they prefer to restore connectivity as soon as possible and identify potential issues later.

Break the Loop Disabling Ports

A simple solution is to manually disable every port that is providing redundancy in the network. If you have been able to identify a part of the network that is affected more, start disabling ports in that area. If possible, start by disabling ports that should be blocking. Each time you disable a port, check to see if connectivity is restored in the network. If you know which port stopped the loop after being disabled, you can be sure that the failure was located on a redundant path where this port was located. If this port should have been blocking, you have probably found the link on which the failure occurred.

Log STP Events on Devices Hosting Blocked Ports

If you cannot identify precisely the source of the problem, or if the problem is only transient, enable logging of STP events on the bridges and switches of the network experiencing the failure. At a minimum, enable logging on devices hosting blocked ports, because it is always the transition of a blocked port that creates the loop.

Use the command **debug spanning-tree events** to enable STP debugging. Use the command **logging buffered** from global configuration mode to capture this debug information into the buffers of the device.

Check Ports

The ports to be investigated first are the blocking ports. What follows is a list of what you can look for on the different ports, with a brief description of the commands to enter.

Check That Blocked Ports Receive BPDUs

Check that BPDUs are being received periodically, especially on blocked ports and root ports.

If you are running Cisco IOS Release 12.0 or later release, the command **show spanning-tree** <*bridge-group* #> displays a field named BPDU, which will show you the number of BPDUs received on each interface. Issuing the command several times will quickly tell you if the device is receiving BPDUs.

Check for Duplex Mismatch

To look for a duplex mismatch, check on each side of a point-to-point link. Simply use the **show interface** command to check the speed and duplex status of the specified ports.

Check Port Utilization

An overloaded interface can fail to transmit vital BPDUs. An overloaded link is also an indication of a possible bridging loop.

Use the command **show interface** to determine interface utilization. Check the output for load and packet input and output.

Check Frame Corruption

Look for increases in the input errors field of the **show interface** command.

Look for Resource Errors

A high CPU utilization can be dangerous for a system running the STA. Use the **show processes cpu** command to check whether the CPU utilization is approaching 100 percent.

Disable Unneeded Features

Disabling as many features as possible helps simplify the network structure and eases the troubleshooting process. EtherChannel, for example, is an advanced feature that needs STP to logically bundle several different links into a single logical port. It can be helpful to disable this feature during troubleshooting. In general, simplifying the network configuration reduces the troubleshooting effort.

Example: STP debug Command

The command **debug spanning-tree** takes a variety of arguments to limit output to events that are specific to a certain STP feature. This example shows output regarding all events while interface GigabitEthernet 0/1 went down.

Caution As with all **debug** commands, be very careful with **debug spanning-tree**. This command is extremely resource-intensive and will interfere with normal network traffic processing.

```
Switch#debug spanning-tree events
Spanning Tree event debugging is on
Switch#
*Mar  5 21:23:14.994: STP: VLAN0013 sent Topology Change
Notice on Gi0/3
*Mar  5 21:23:14.994: STP: VLAN0014 sent Topology Change
Notice on Gi0/4
*Mar  5 21:23:14.994: STP: VLAN0051 sent Topology Change
Notice on Po3
*Mar  5 21:23:14.994: STP: VLAN0052 sent Topology Change
Notice on Po4
*Mar  5 21:23:15.982: %LINEPROTO-5-UPDOWN: Line protocol on
Interface GigabitEthernet0/1, changed state to down
*Mar  5 21:23:16.958: STP: VLAN0001 Topology Change rcvd on
Po1
*Mar  5 21:23:16.958: STP: VLAN0011 Topology Change rcvd on
Po1
*Mar  5 21:23:16.962: STP: VLAN0012 Topology Change rcvd on
Po1
*Mar  5 21:23:16.962: STP: VLAN0015 Topology Change rcvd on
Po1
*Mar  5 21:23:16.962: STP: VLAN0016 Topology Change rcvd on
Po1
*Mar  5 21:23:16.966: STP: VLAN0017 Topology Change rcvd on
Po1
*Mar  5 21:23:16.966: STP: VLAN0018 Topology Change rcvd on
Po1
*Mar  5 21:23:16.998: %LINK-3-UPDOWN: Interface
GigabitEthernet0/1, changed state to down
```

Summary

This topic summarizes the key points discussed in this lesson.

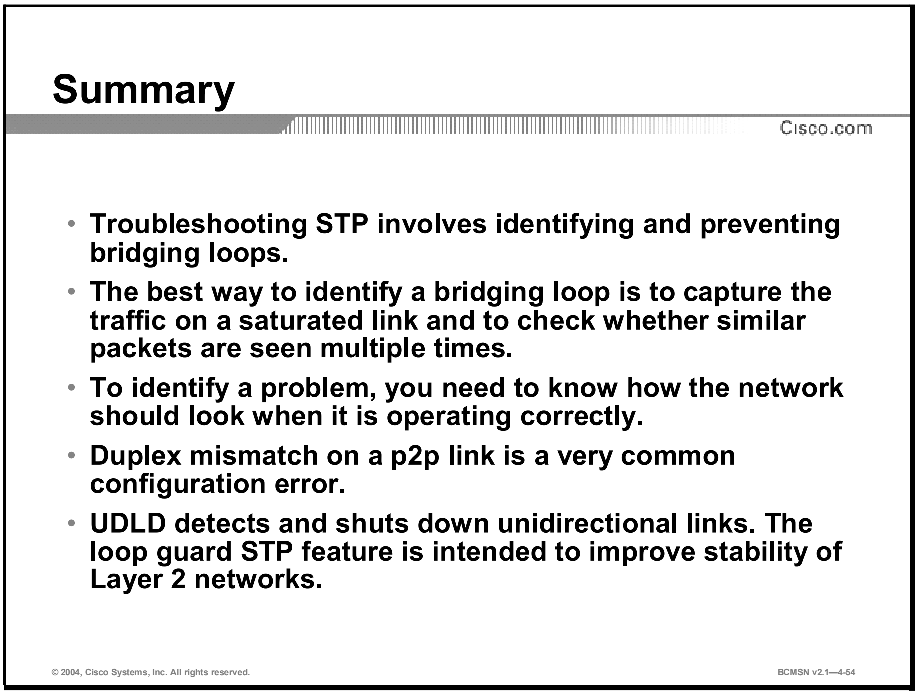

Summary

Cisco.com

- Troubleshooting STP involves identifying and preventing bridging loops.
- The best way to identify a bridging loop is to capture the traffic on a saturated link and to check whether similar packets are seen multiple times.
- To identify a problem, you need to know how the network should look when it is operating correctly.
- Duplex mismatch on a p2p link is a very common configuration error.
- UDLD detects and shuts down unidirectional links. The loop guard STP feature is intended to improve stability of Layer 2 networks.

BCMSN v2.1—4-54

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 4-1: Enhancing Spanning Tree Protocol

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are in the Quiz Answer Key.

Q1) Select the most common STP problems that can occur. (Choose two.)

A) frame corruption

B) redundant loops

C) duplex mismatch

D) spanning tree backup

Q2) Select the scenario that represents duplex mismatching.

A) One port is in full-duplex mode, and one port is in half-duplex mode.

B) One port is blocked and one port is full duplex.

C) Both links are configured for full duplex.

D) One port is blocked and one port is forwarding.

Q3) Select the scenario that represents unidirectional link failure.

A) when there is an undetected failure on a fiber link

B) when two switches are both blocking

C) when a link is up while providing one-way communication

D) when there are disabled ports on the interface

Q4) Select the scenario that represents a PortFast configuration error.

A) enabling PortFast on a link to a PC

B) enabling PortFast on a link to a Windows server

C) enabling PortFast on a link to a switch

D) enabling PortFast on a link to a UNIX server

Q5) Select the most appropriate approach to solving an STP problem. (Choose two.)

A) verify the number of switches in your network

B) verify where the root bridge is located

C) verify the last BPDUs that were sent

D) verify the location of the blocked ports

Q6) Select the correct commands that display the debug messages. (Choose two.)

A) Switch#**debug spanning-tree events uplinkfast**

B) Switch#**debug spanning-tree all backbonefast**

C) Switch#**debug spanning-tree uplinkfast**

D) Switch#**debug spanning-tree all**

Quiz Answer Key

Q1) A, C

Relates to: Identifying STP Problems

Q2) A

Relates to: Identifying Duplex Mismatching

Q3) A

Relates to: Identifying Unidirectional Link Failure

Q4) C

Relates to: Identifying PortFast Configuration Error

Q5) B, D

Relates to: Troubleshooting the Spanning Tree Protocol

Q6) C, D

Relates to: Identifying Spanning Debug Commands

Lesson Assessments

Overview

Use the lesson assessments here to test what you learned in this module. The correct answers and solutions are found in the Lesson Assessment Answer Key.

Outline

This section includes these assessments:

- Quiz 4-1: Optimizing Spanning Tree Protocol
- Quiz 4-2: Accelerating Spanning Tree Convergence
- Quiz 4-3: Configuring Spanning Tree
- Quiz 4-4: Tuning the Spanning Tree Protocol
- Quiz 4-5: Selecting a Troubleshooting Approach

Quiz 4-1: Optimizing Spanning Tree Protocol

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Identify the correct features that apply to PortFast
- Identify the correct features that apply to UplinkFast
- Identify the correct features that apply to BackboneFast
- Select the appropriate characteristic that applies to EtherChannel
- Match the correct features with the appropriate PAgP and LACP protocols
- Identify how STP enhances the Enterprise Composite Network model

Quiz

Answer these questions, based on the diagram:

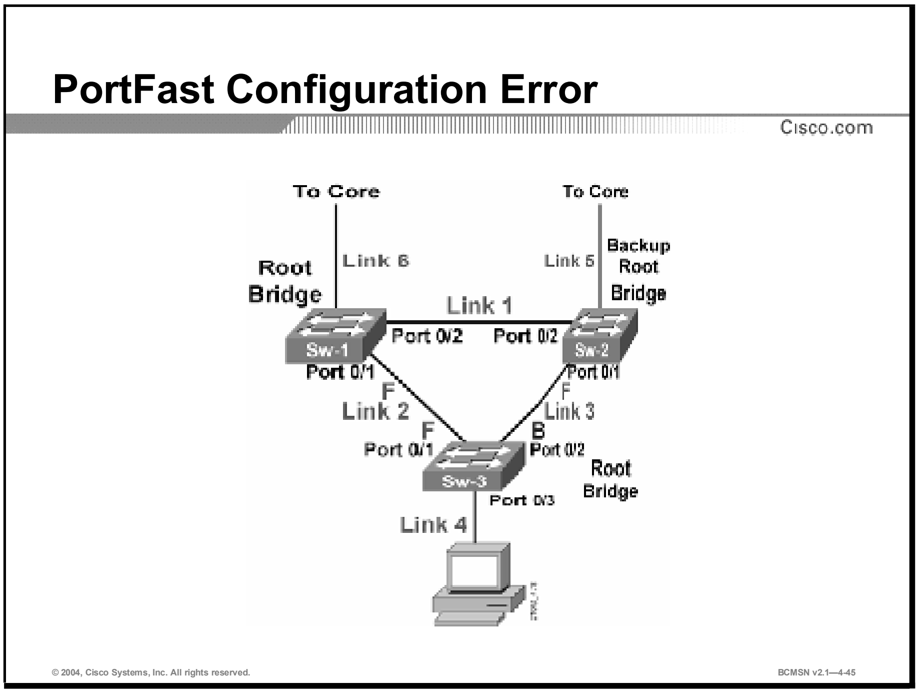

PortFast Configuration Error

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—4-45

Q1) Each switch is running IEEE 802.1D spanning tree. What is the correct location where the PortFast feature should be applied?

A) switch Sw-1 Port 0/1

B) switch Sw-2 Port 0/1

C) switch Sw-3 Port 0/1 and switch Sw-3 Port 0/2

D) switch Sw-3 Port 0/3

E) A, B, and C

F) PortFast would not be configured in this topology.

Q2) Each switch is running IEEE 802.1D spanning tree. Which is correct about where the UplinkFast feature is applied?

A) switch Sw-3

B) switches Sw-1 and Sw-2

C) all switches in the topology

D) all ports on switch Sw-3

E) UplinkFast would not be configured in this topology.

Q3) Each switch is running IEEE 802.1D spanning tree. Which is correct about where the BackboneFast feature is applied?

A) switch Sw-3

B) switches Sw-1 and Sw-2

C) all switches in the topology

D) all ports on switch Sw-3

E) BackboneFast would not be configured in this topology.

Q4) Each switch is running IEEE 802.1D spanning tree, and Sw-1 and Sw-2 are multilayer switches. Which are the optimal features for Link 5 and Link 6? (Choose two.)

A) Layer 2

B) Layer 3

C) EtherChannel

D) BackboneFast

E) PortFast

F) UplinkFast

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 67 percent or better.

Quiz 4-2: Accelerating Spanning Tree Convergence

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Select the correct features of RSTP
- Match the correct RSTP port states with their appropriate function
- Match the correct RSTP port role with the appropriate function
- Match the correct RSTP link type with the appropriate function
- Select the features that are attributed to RSTP BPDUs
- Select the correct features of rapid transition to the forwarding state
- Match the actions that apply to notifying topology changes with RSTP in their correct order
- Select the features that apply to multiple spanning tree
- Select the characteristics that apply to the MST region
- Select the appropriate benefits of applying RSTP in the Enterprise Composite Network model

Quiz

Answer these questions:

Q1) Select a benefit that RSTP provides over STP.

A) With RSTP, enhancements are very visible.

B) RSTP increases the speed of recalculating spanning tree when a Layer 2 network topology changes.

C) RSTP uses 802.1D delay timers.

D) RSTP requires additional configuration.

Q2) Select the port state feature that has changed as a result of RSTP.

A) The port state is identical to the state of the port.

B) RSTP can have many port states.

C) The port state is separated from the role of the port.

D) The port role and port state are determined by processes that are dependent on each other.

Q3) Which statement best describes a characteristic of RSTP BPDUs?

A) RSTP uses 2 bits in the type field.

B) RSTP BPDUs are relayed.

C) RSTP uses 6 bits of the flag byte.

D) RSTP BPDUs are generated when a BPDU is received on the root port.

Q4) Select the correct port in RSTP that causes a topology change.

A) nonedge ports

B) root port

C) point-to-point links

Q5) Select the feature of an IST instance.

A) An IST instance can have many blocked ports.

B) Loops can be broken only by a blocked port in the middle of the MST region.

C) Entire regions appear as one virtual bridge and run a CST.

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Quiz 4-3: Configuring Spanning Tree

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Identify the command to configure port-level tuning with PortFast
- Identify the command to configure UplinkFast
- Identify the command to configure BackboneFast
- Identify the command to enable MST
- Identify the command to verify MST
- Select the appropriate guidelines that apply to configuring EtherChannel
- Identify the command to configure EtherChannel
- Identify the command to verify EtherChannel
- Identify the command to configure PAgP and LACP
- Identify the command to verify PAgP and LACP
- Identify the command to configure EtherChannel load balancing

Quiz

Answer these questions:

Q1) You are setting up a network and want to be sure that your Layer 2 access port bypasses the listening and learning states and enters the forwarding state immediately. Select the correct command to ensure this will happen.

A) switch #**show running-config interface {{fastethernet | gigabitethernet}** *slot/port*} **| {port-channel** *port_channel_number*}

B) switch(config)#**interface fastethernet 5/8**

C) switch(config-if)#**[no] spanning-tree portfast**

D) switch#**show running-config interface fastethernet 5/8**

Q2) What action is performed by these commands?

```
switch(config)#[no] spanning tree uplinkfast [max-update-rate
max_update_rate]
```

```
switch# show spanning tree totals
```

A) verifies the UplinkFast configuration

B) enables or disables UplinkFast

C) identifies which VLANs have UplinkFast enabled

Q3) When configuring MST, what action does this command perform?

```
switch(config-mst)#instance instance_number vlan vlan_range
```

A) sets the MST region name

B) displays the current MST configuration

C) maps the VLANs to an MST instance

D) displays the new MST configuration to be applied

Q4) When configuring EtherChannel, match the command with the correct order in which it should be performed.

_____ 1. Step 1

_____ 2. Step 2

_____ 3. Step 3

_____ 4. Step 4

A) switch(config-if)#**channel-group number mode {auto | desirable | on}**

B) switch(config-if)#**interface port-channel** *port-channel-number*

C) switch(config)#**interface** *interface slot/port*

D) switch(config-if)#**ip address** *address mask*

Q5) When configuring LACP or PagP, match the command with the correct order in which it should be performed.

_____ 1. Step 1

_____ 2. Step 2

_____ 3. Step 3

_____ 4. Step 4

_____ 5. Step 5

A) switch(config-if)**#channel-protocol {lacp | pagp}**

B) switch(config-if)#**end**

C) switch(config)#**interface type1** *slot/port*

D) switch(config)#**lacp system-priority** *priority_value*

E) switch(config-if)#**lacp port-priority** *priority_value*

Q6) Match the EtherChannel load balancing keywords with the correct function.

_____ 1. src-mac

_____ 2. dst-mac

_____ 3. src-dst-ip

_____ 4. dst-ip

A) destination MAC address

B) destination IP address

C) source MAC address

D) source and destination IP addresses

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 75 percent or better.

Quiz 4-4: Tuning the Spanning Tree Protocol

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Identify the command to configure BPDU guard
- Identify the command to enable BPDU filtering
- Identify the command to enable BPDU skewing detection
- Identify the function of root guard
- Identify the command to configure root guard
- Identify the correct features for UDLD protocol
- Identify the command to configure UDLD
- Identify the features of loop guard
- Identify the command to configure loop guard
- Match the correct loop guard or UDLD feature with the appropriate name
- Identify the protection features used in the Enterprise Composite Network model

Quiz

Answer these questions:

Q1) What action is performed by this command?

```
switch(config)# [no] spanning-tree portfast bpduguard
```

A) enables BPDU guard

B) verifies BPDU configuration

C) enables BPDU guard to shut down PortFast-configured interfaces that receive BPDUs

Q2) Select the command used to prevent the system from sending or receiving BPDUs on all ports.

A) switch #**show spanning-tree summary totals**

B) switch(config)#**spanning-tree portfast bpdufilter default**

C) switch(config)#**spanning-tree portfast bpduguard**

D) switch(config)#**[no] spanning-tree portfast bpduguard**

Q3) Select the command that determines whether any ports are in a root-inconsistent state.

A) switch #**show spanning-tree inconsistentports**

B) switch (config)#**show spanning tree inconsistentports**

C) switch (config-if)#**show spanning tree inconsistentports**

D) switch #**show spanning-tree portsinconsistent**

Q4) Select the command that disables UDLD on individual nonfiber-optic interfaces.

A) switch(config-if)#**no udld disable**

B) switch(config-if)#**udld enable**

C) switch(config)#**no udld enable**

D) switch(config-if)#**no udld enable**

Q5) Select the command that disables loop guard.

A) switch(config)#**spantree guard none mod/port**

B) switch(config)#**spantree guard mod/port**

C) switch(config-if)#**spantree guard none mod/port**

D) switch(config-if)#**spantree guard mod/port**

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 75 percent or better.

Quiz 4-5: Selecting a Troubleshooting Approach

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Troubleshoot a duplex mismatching
- Troubleshoot a unidirectional link failure

Quiz

Answer these questions:

Q1) Every packet (including the BPDUs) that is being sent by a particular bridge is deferred and eventually dropped, thus creating a bridging loop. Select the best solution to this problem.

A) ensure that one link is configured for half duplex and the others are configured for full duplex

B) ensure that both links are configured for half duplex

C) ensure that there is not an undetected failure on a fiber link

D) ensure that a port that is connected to a host is not directly transitioning to the forwarding state

Q2) You have a bridged environment and are having a problem with a loop; rebooting your machine does not help. Select the best solution to the problem.

A) ensure that both links are configured for half duplex

B) ensure that one link is configured for half duplex and the others are configured for full duplex

C) run UDLD

D) configure a root guard

Q3) You discover a lot of frame corruption. Select the most appropriate actions to solve the problem.

A) verify that cable lengths are correct

B) verify that there is not an undetected failure on a fiber link

C) verify that there is not a duplex mismatch

D) verify that there is not a unidirectional link problem

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 75 percent or better.

Lesson Assessment Answer Key

Quiz 4-1: Optimizing Spanning Tree Protocol

Q1) D

Q2) A

Q3) E

Q4) B, C

Quiz 4-2: Accelerating Spanning Tree Convergence

Q1) B

Q2) C

Q3) C

Q4) A

Q5) C

Quiz 4-3: Configuring Spanning Tree

Q1) C

Q2) B

Q3) C

Q4) 1=B, 2=D, 3=C, 4=A

Q5) 1=C, 2=A, 3=E, 4=B, 5=D

Q6) 1=C, 2=A, 3=D, 4=B

Quiz 4-4: Tuning the Spanning Tree Protocol

Q1) C

Q2) B

Q3) A

Q4) D

Q5) A

Quiz 4-5: Selecting a Troubleshooting Approach

Q1) B

Q2) C

Q3) A, C

Module 5

Implementing Multilayer Switching in the Network

Overview

Multilayer switching provides high-performance hardware-based Layer 2 and Layer 3 switching for networking. With multilayer switching, unicast IP data packets are switched between IP subnets using advanced application-specific integrated circuit (ASIC) switching hardware. This approach offloads processor-intensive packet routing from network routers.

The packet forwarding function is moved onto Layer 3 switches whenever a complete switched path exists between two hosts. Standard routing protocols, such as Open Shortest Path First Protocol (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), and Routing Information Protocol (RIP) are used for route determination.

Upon completing this module, you will be able to:

- Select multilayer switching architectures, given specific multilayer switching needs
- Configure, monitor, and troubleshoot Multilayer Switching and Cisco Express Forwarding
- Configure inter-VLAN routing

Outline

The module contains these components:

- Examining Multilayer Switching
- Configuring Multilayer Switching
- Routing Between VLANS
- Lesson Assessments

Examining Multilayer Switching

Overview

Multilayer switching refers to a class of high-performance devices optimized for the campus LAN or intranet. You can improve the performance of your networked applications by using the high-speed packet switching, routing, and enhanced network services that a multilayer switching network offers.

Relevance

Layer 3 switching offers performance enhancements not possible on Layer 2 switches.

Objectives

Upon completing this lesson, you will be able to:

- Identify the operation of the key components required to implement Layer 3 switching
- Compare Layer 2 and multilayer forwarding, and explain the packet flow with each type of forwarding
- Explain the multilayer switching table architectures
- Describe centralized and distributed forwarding, demand-based switching, and topology-based switching

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Introducing Layer 3 Switching
- Multilayer Switch Packet Forwarding
- Switching Table Architectures
- Switch Forwarding Architectures
- Summary
- Quiz

Introducing Layer 3 Switching

Layer 3 switching refers to a class of high-performance switches optimized for the campus LAN or intranet. The same technology and intelligence that exists in traditional routing platforms has been incorporated into Layer 3 switches. This topic identifies the operation of the key components required to implement Layer 3 switching.

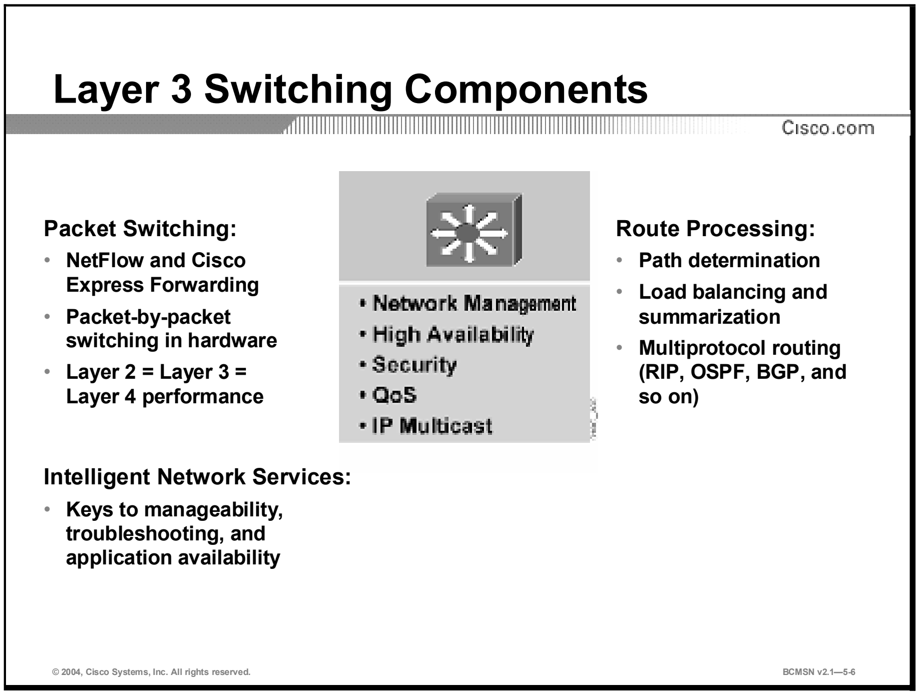

Layer 3 Switching Components

Cisco.com

Packet Switching:

- **NetFlow and Cisco Express Forwarding**
- **Packet-by-packet switching in hardware**
- **Layer 2 = Layer 3 = Layer 4 performance**

- Network Management
- High Availability
- Security
- QoS
- IP Multicast

Route Processing:

- **Path determination**
- **Load balancing and summarization**
- **Multiprotocol routing (RIP, OSPF, BGP, and so on)**

Intelligent Network Services:

- **Keys to manageability, troubleshooting, and application availability**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—5-6

With Layer 3 switches, the prime area of focus is often raw performance, which refers to the aggregate number of packets that a device can switch in and out over a fixed period of time. Layer 3 switches tend to have packet-switching throughputs in the millions of packets per second (pps), while traditional general-purpose routers have evolved from the 100,000 pps range to over one million pps. However, route processing and intelligent network services must be considered in determining the significant benefit that Layer 3 switching will have on applications running on your network.

The two core elements to consider for switching implementations are where and how the switching decision is to be made.

- **Centralized switching:** In this method, the switching decision is made centrally. A central forwarding table, typically controlled by an ASIC, performs Layer 2 or Layer 3 lookups. The ASIC implementation must be able to search in the table extremely quickly, so that it does not become a bottleneck.
- **Distributed switching:** In this method, a switching decision can either be made locally by a port or, in the case of some modular chassis, on a line card level. In distributed switching models, address tables that have been distributed must be synchronized to account for adds, moves, and changes. Without synchronization of the local tables with the master table, tables may contain out-of-date data and try to forward packets to a port of exit that does not exist.

The next item to examine is how Layer 3 switching takes place using one of these methods:

- **Route caching:** Also known as flow-based or demand-based switching, route caching is a model in which a table is built based on traffic flow into the switch.
- **Topology-based switching:** This alternative method of Layer 3 switching is based on a Forwarding Information Base (FIB). In a topology-based switching implementation, the route cache is prepopulated based on the information in the routing table. Because the cache is prepopulated without traffic having to flow through the switch, lookups can be done very quickly and need not be based on traffic. This type of implementation is a more complicated architecture to build, because the search algorithm used to search the information base needs to be highly efficient to avoid becoming a bottleneck. This Cisco topology-based switching forwarding information base is called Cisco Express Forwarding (CEF).

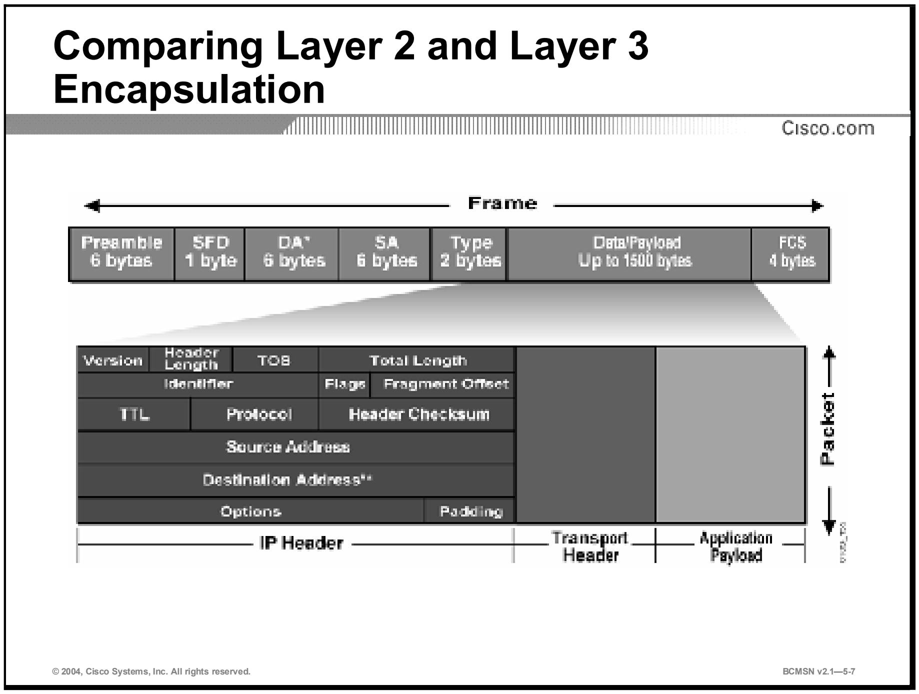

Comparing Layer 2 and Layer 3 Encapsulation

Cisco.com

BCMSN v2.1—5-7

Layer 2 and Layer 3 switches each make decisions about how to handle units of data, known as frames (Layer 2) and packets (Layer 3). A frame is a Layer 2 encapsulation, while a packet is a Layer 3 encapsulation. For example, Ethernet is used for a Layer 2 frame encapsulation, while IP is used for a Layer 3 packet encapsulation.

Layer 2 and Layer 3 switches both read the destination address (DA) field. Layer 2 uses the frame DA field to make switching decisions, while Layer 3 makes switching decisions based on the packet DA.

Multilayer Switch Packet Forwarding

Multilayer switches handle Layer 2, Layer 3, and Layer 4 forwarding in hardware. This topic compares Layer 2 and multilayer forwarding, and the packet flow with each type of forwarding is shown.

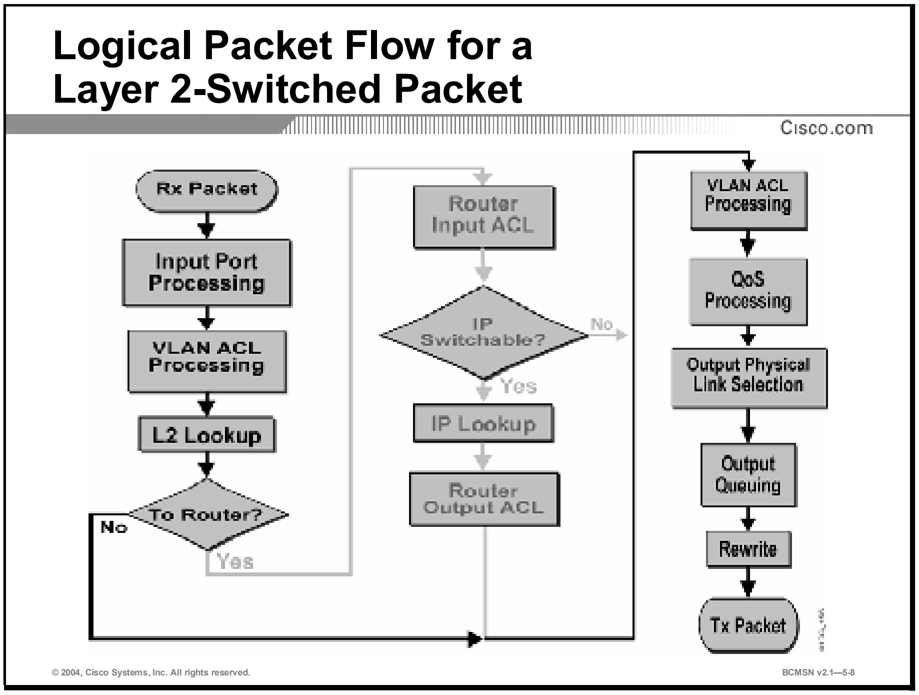

Logical Packet Flow for a Layer 2-Switched Packet

BCMSN v2.1—5-8

Layer 2 forwarding in hardware is based on the destination MAC address. The Layer 2 switch learns the address based on the source MAC address. The MAC address table lists MAC and VLAN pairs with associated interfaces.

The figure describes how a Layer 2 switch forwards packets.

Step 1 The Layer 2 engine receives a frame or the frame header over the data bus.

Step 2 The switch performs these operations:

- The Layer 2 lookup engine looks up the destination MAC address.
- The Layer 2 engine performs the input security access control list (ACL) lookup.

Step 3 The switch then performs these operations:

- The Layer 2 forwarding engine performs the outbound security ACL lookup.
- The Layer 2 forwarding engine performs the outbound quality of service (QoS) ACL lookup.

Step 4 The Layer 2 forwarding engine forwards the packet.

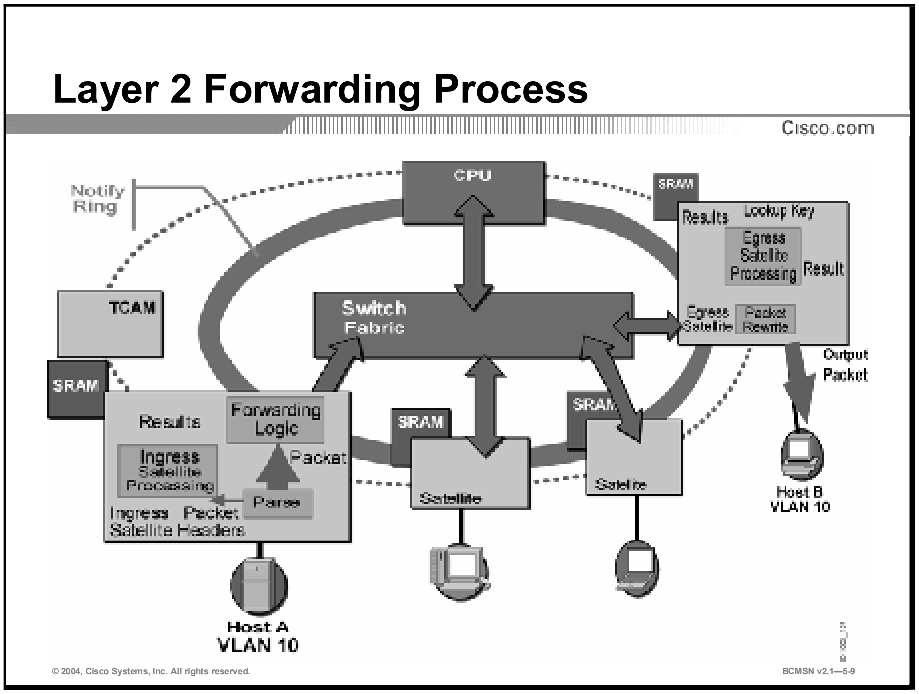

Layer 2 Forwarding Process

Cisco.com

BCMSN v2.1—5-9

The Layer 2 switch receives packets on satellite ASICs. Packets are buffered in the shared memory fabric.

Key headers are generated for lookups. Lookups are done in shared ternary content addressable memory (TCAM) or a forwarding engine. Lookup results may point to additional information stored in satellite static RAM. The results are then transmitted to the egress satellite. Egress lookups and packet rewrites occur on egress satellite ASICs.

In the figure, assume that host A is sending data to host B. Both hosts are located on VLAN10. To make the forwarding decision, the ingress satellite ASIC for host A performs a Layer 2 lookup in the switch TCAM table and inputs ACL and QoS requirements. If the switch did not already know the recipient address, learning would be required. The CPU would then learn the source MAC address. In this case, because both hosts are on the same VLAN, no learning is required.

The egress satellite within the switch receives the packet and performs a lookup in the TCAM table to determine any ACL and QoS requirements. The egress satellite then forwards the packet to host B.

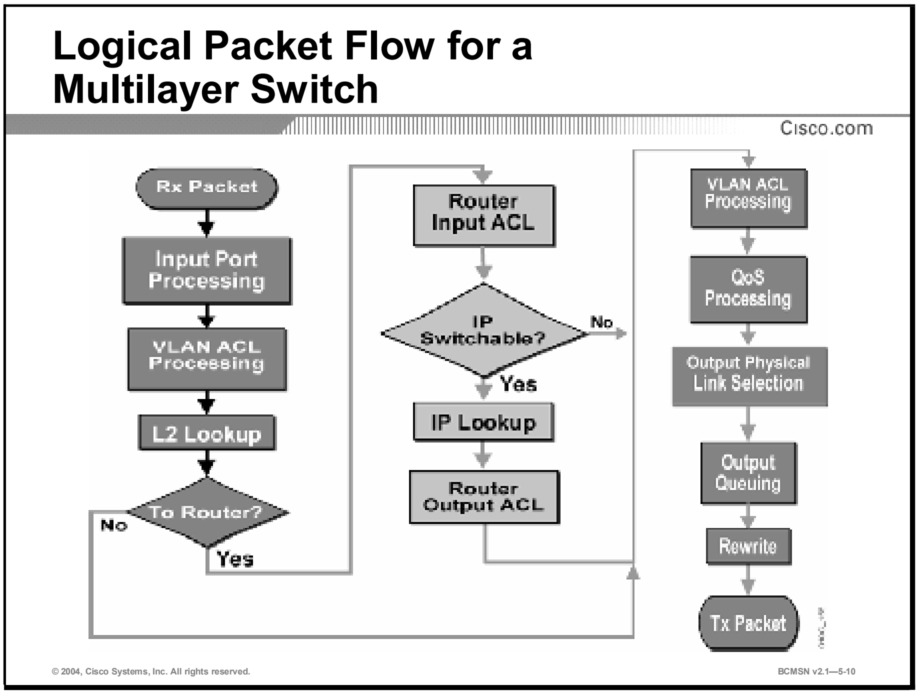

Logical Packet Flow for a Multilayer Switch

Cisco.com

BCMSN v2.1—5-10

Multilayer switches forward Layer 3 unicast packets in hardware using CEF-based Multilayer Switching (MLS), which is implemented on Cisco multilayer switches.

The figure describes how a multilayer switch forwards packets.

Step 1 The Layer 2 and Layer 3 engines receive a frame or the frame header over the data bus.

Step 2 The switch performs these operations:

- The Layer 2 lookup engine looks up the destination MAC address.
- The Layer 2 engine performs the input security ACL lookup.
- The Layer 2 engine performs the input QoS ACL lookup.
- The Layer 3 engine performs an FIB table lookup.
- The Layer 3 engine performs a NetFlow (for devices with NetFlow capability) table lookup.

Step 3 The switch then performs these operations:

- The Layer 2 forwarding engine forwards the input ACL and QoS results to the Layer 3 forwarding engine.
- The Layer 3 forwarding engine sends the destination VLAN information for the packet to the Layer 2 forwarding engine.

Step 4 The switch then performs these operations:

- The Layer 3 forwarding engine performs the adjacency lookup.
- The Layer 2 forwarding engine performs the outbound security ACL lookup.
- The Layer 2 forwarding engine performs the outbound QoS ACL lookup.

Step 5 The Layer 2 forwarding engine sends the results of the security and QoS ACL lookups to the Layer 3 forwarding engine.

Step 6 The switch then performs these operations:

- The Layer 3 forwarding engine generates the rewrite result and sends it over to the Layer 2 forwarding engine.
- The Layer 3 forwarding engine updates the table statistics.

Step 7 The Layer 2 forwarding engine looks up the destination MAC address from the Layer 3 forwarding engine. The Layer 2 forwarding engine then chooses between a Layer 2 and Layer 3 result and sends the result onto the result bus.

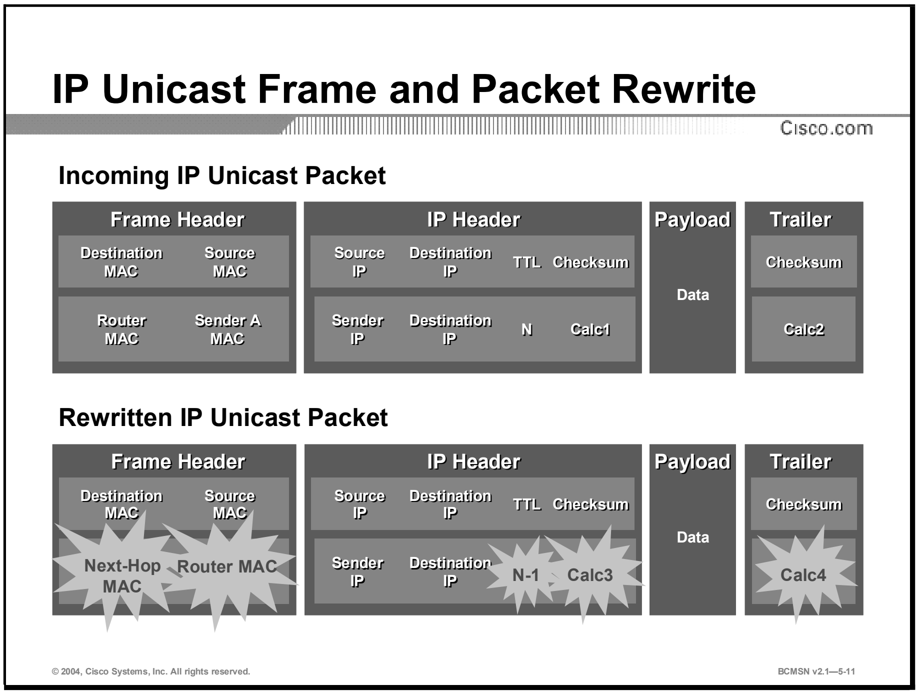

IP Unicast Frame and Packet Rewrite

Incoming IP Unicast Packet

Frame Header		IP Header				Payload	Trailer
Destination MAC	Source MAC	Source IP	Destination IP	TTL	Checksum	Data	Checksum
Router MAC	Sender A MAC	Sender IP	Destination IP	N	Calc1		Calc2

Rewritten IP Unicast Packet

Frame Header		IP Header				Payload	Trailer
Destination MAC	Source MAC	Source IP	Destination IP	TTL	Checksum	Data	Checksum
Next-Hop MAC	Router MAC	Sender IP	Destination IP	N-1	Calc3		Calc4

BCMSN v2.1—5-11

IP unicast packets are rewritten on the output interface as follows:

- The source MAC address changes from the sender MAC address to the router MAC address.
- The destination MAC address changes from the router MAC to the next-hop MAC address.
- The Time to Live (TTL) is decremented by one.
- The IP header and frame checksums are recalculated.

Switching Table Architectures

When routing, bridging, or ACL tables are built, the information is stored in a high-speed table memory. Lookups are done with efficient search algorithms. The table architectures used in switching are the content addressable memory (CAM) table and the TCAM table. This topic describes the multilayer switching table architectures.

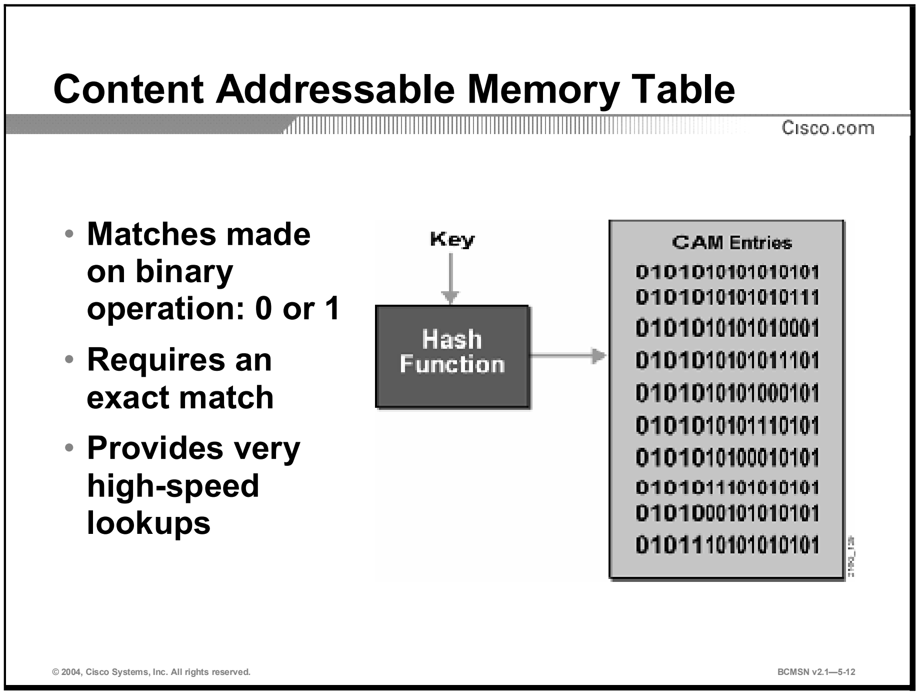

Content Addressable Memory Table

Cisco.com

- Matches made on binary operation: 0 or 1
- Requires an exact match
- Provides very high-speed lookups

BCMSN v2.1—5-12

In a CAM table, matches are made on a binary operation. The switch must find an exact match. CAM provides very high-speed lookup in large tables. CAM tables are used as follows:

- **Catalyst 6500 series:** Layer 2 and NetFlow tables
- **Catalyst 4000 series:** Layer 2 tables

The information used to do a lookup in a CAM table is called a “key.” For example, a Layer 2 lookup would use a destination MAC address and a VLAN ID (VID) for a key. The key is fed into a hashing algorithm. The hashing algorithm then produces a pointer into the table. The system knows exactly where the pointer is and finds the result without searching the entire table.

In a Layer 2 table, all information, including VLAN, destination MAC address, route destination, destination ports or virtual circuits, and protocol type is important. However, in more complicated tables, there can be information that you care about and information that you do not care about. For example, you might be concerned with the first 24 bits of an IP address. In that case, you do not want to do an exact match; instead, you want a match only on the values that you care about.

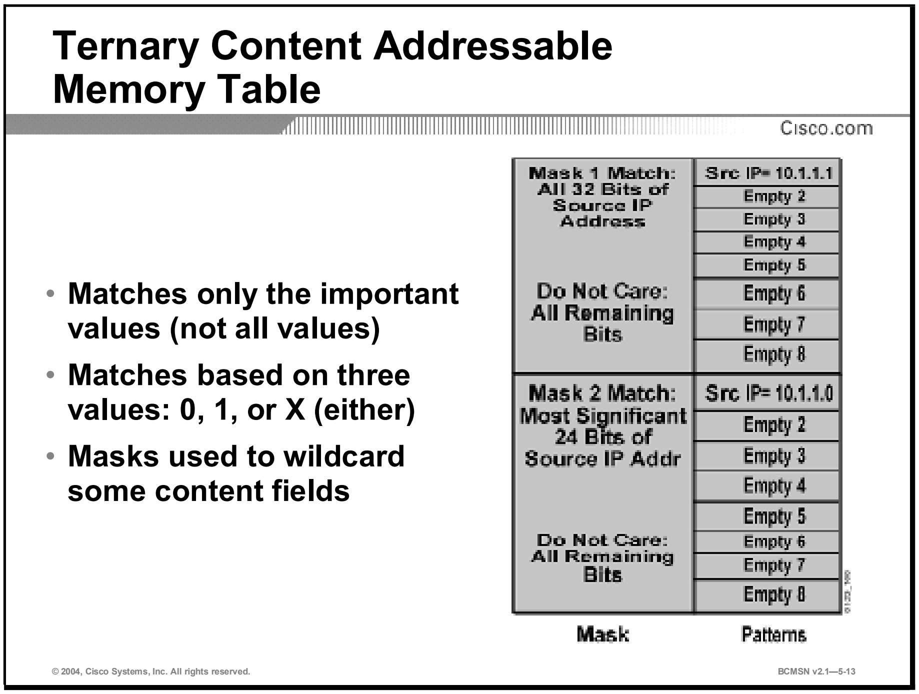

Ternary Content Addressable Memory Table

- Matches only the important values (not all values)
- Matches based on three values: 0, 1, or X (either)
- Masks used to wildcard some content fields

Mask	Patterns
Mask 1 Match: All 32 Bits of Source IP Address	Src IP= 10.1.1.1
	Empty 2
	Empty 3
	Empty 4
	Empty 5
Do Not Care: All Remaining Bits	Empty 6
	Empty 7
	Empty 8
Mask 2 Match: Most Significant 24 Bits of Source IP Addr	Src IP= 10.1.1.0
	Empty 2
	Empty 3
	Empty 4
	Empty 5
Do Not Care: All Remaining Bits	Empty 6
	Empty 7
	Empty 8

BCMSN v2.1—5-13

TCAM is a specialized piece of memory designed for rapid table lookups—based on packets passing through the switch—performed by the ACL engine. The result of the ACL engine lookup into the TCAM table determines how the switch handles a packet. For example, the packet might be permitted or denied.

The TCAM table has a limited number of entries that are populated with pattern values and mask values, each with an associated result. These are known as value, mask, result (VMR) entries. The term VMR simply refers to the format in which access control entries (ACEs) are represented in the TCAM table.

The "value" in VMR refers to the pattern that is to be matched (such as IP addresses, protocol ports, and so on). The "mask" refers to the mask bits associated with the pattern. The "result" refers to the result or action that occurs in the case of a lookup returning a hit for the pattern and mask. This result might be a simple "permit" or "deny," or it may be a pointer to other more complex information. For example, in policy-based routing (PBR), the result is a pointer to an entry in the hardware adjacency table that contains the next-hop information.

With TCAM, matching is based on three values: 0, 1, or *x* (where *x* is either number), hence the term ternary. The memory structure is broken into a series of patterns and masks. Masks are shared among a specific number of patterns and are used to wildcard some content fields.

To perform a lookup in a TCAM table, all entries are checked in parallel. The performance is independent of the number of entries. This allows a switch to use the longest match lookup when needed and provides fixed latency to ignore fields.

These platforms use TCAMs for Layer 3 switching:

- Catalyst 6500
- Catalyst 4000
- Catalyst 3550

An example of using ACLs stored in TCAMs is shown in the figure. Assume these two entries:

```
access-list 101 permit ip host 10.1.1.1 any
access-list 101 deny ip 10.1.1.0 0.0.0.255 any
```

The two mask value entries and two pattern value entries are consumed. You can mask the values that you do not care about. The remaining mask bits are "do not care bits," meaning that the destination IP address, port numbers, and so on, are of no concern.

The TCAM table consists of these types of regions:

- **Exact-match region:** The exact-match region consists of Layer 3 entries for multiple protocol regions, such as IP adjacencies and Internetwork Packet Exchange (IPX) node.
- **Longest-match region:** Each longest-match region consists of groups of Layer 3 address entries ("buckets") organized in decreasing order by mask length. All entries within a bucket share the same mask value and key size. The buckets can change their size dynamically by borrowing address entries from neighboring buckets. Although the size of the whole protocol region is fixed, you can reconfigure it. The reconfigured size of the protocol region is effective only after the next system reboot.
- **First-match region:** The first-match region consists of ACL entries. Lookup stops after first match of the entry.

The table lists default partitioning for each protocol region in TCAM.

Protocol Region	Lookup Type	Key Size	Default Size	Number of TCAM Entries
ipx-bvi-network	Exact-match	32 bits	32	32
ip-adjacency	Exact-match	32 bits	2048	2048
ipx-node	Exact-match	64 bits	2048	4096
ip-prefi	Longest-match	32 bit	8192	8192
ipx-network	Exact-match	32 bits	6144	6144
ip-mcast	Longest-match	64 bits	3072	6144
l2-switching	Exact-match	64 bits	1024	2048
udp-flooding	Exact-match	64 bits	256	512
access-list	First-match	128 bits	512	8192

The enhanced Gigabit Ethernet interface module supports TCAM sizes of 32 KB, 64 KB, or 256 KB. Each entry in TCAM is 32 bits wide. You can configure the various protocol regions in TCAM based on your requirements and on the size of TCAM on your Gigabit Ethernet interface module.

Switch Forwarding Architectures

Multilayer switches may offer centralized forwarding or distributed forwarding and may offer NetFlow switching or topology-based switching. This topic explains the switch forwarding architectures supported on Cisco multilayer switches.

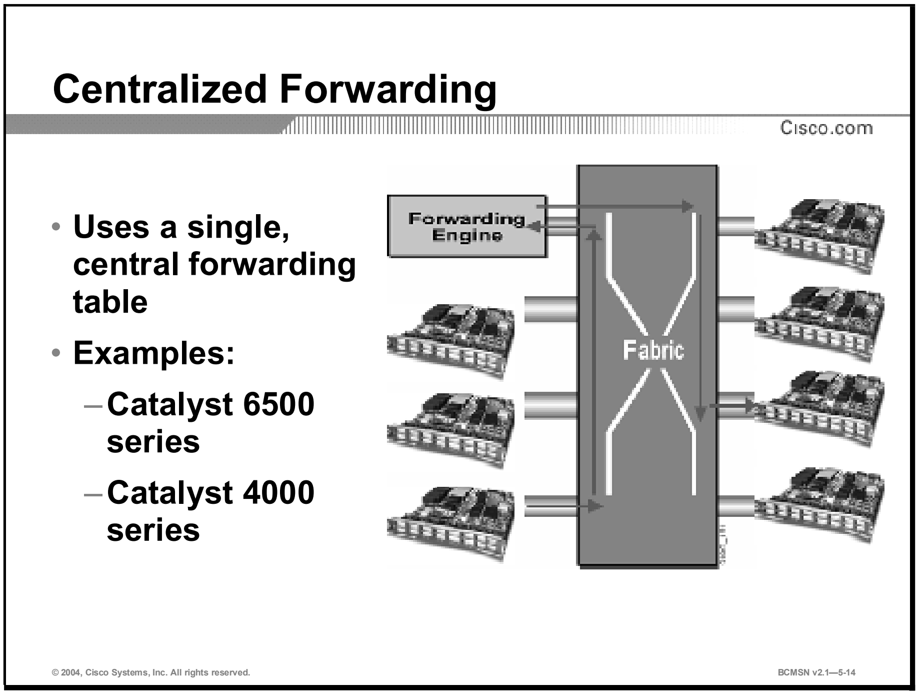

Centralized Forwarding

Cisco.com

- Uses a single, central forwarding table
- Examples:
  - Catalyst 6500 series
  - Catalyst 4000 series

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—5-14

With centralized forwarding, a single, central forwarding table is used. All forwarding operations are completed centrally, including Layer 2, Layer 3, QoS, ACLs, and so on. System performance capacity is based on the central forwarding engine.

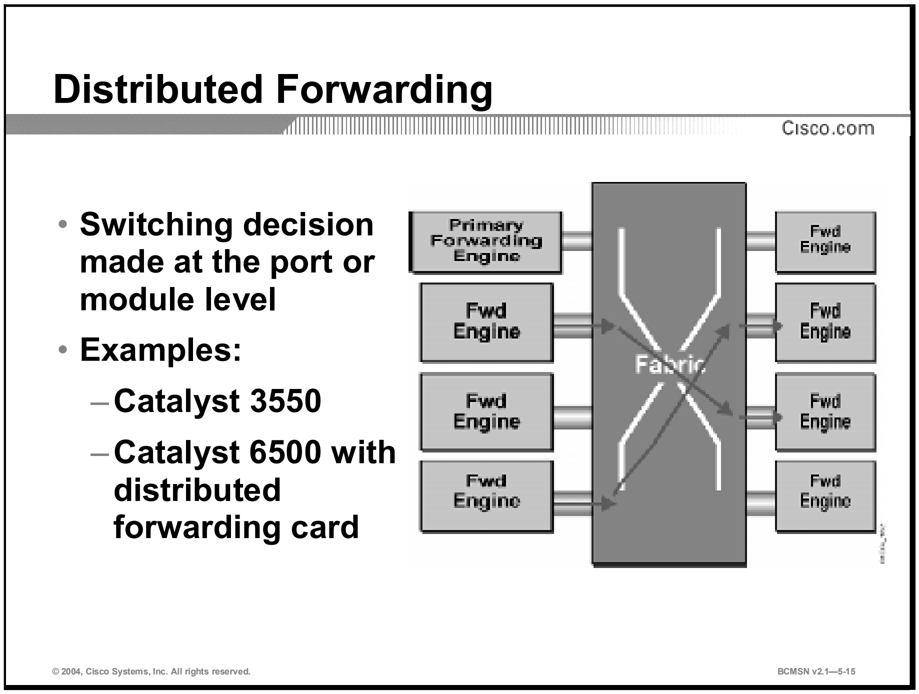

Distributed Forwarding

Cisco.com

- Switching decision made at the port or module level
- Examples:
  - Catalyst 3550
  - Catalyst 6500 with distributed forwarding card

BCMSN v2.1—5-15

With distributed forwarding, the switching decision is made at the port or module level. Forwarding tables must be synchronized to account for topology changes. The primary forwarding engine manages the distributed tables.

System performance is equal to the aggregate of all forwarding engines. Distributed forwarding enables switches to achieve rates over 100 Mpps.

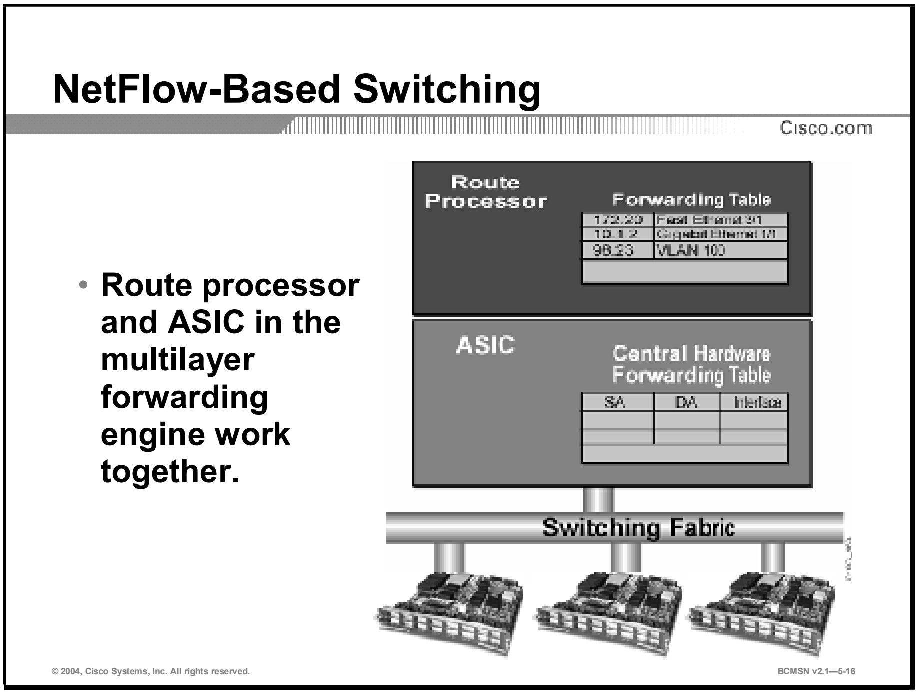

NetFlow-Based Switching

Cisco.com

- **Route processor and ASIC in the multilayer forwarding engine work together.**

BCMSN v2.1—5-16

NetFlow-based switching uses multilayer forwarding engines. The route processor and ASICs work together. An entry contains source, source and destination, or full flow (up to Layer 4) information. All traffic forwarding is unidirectional.

The first packet in a stream is switched in software. The destination MAC address must be for the default gateway. The forwarding decision is programmed in the hardware-forwarding subtable for subsequent packets. Subsequent packets in that stream are forwarded via the hardware-forwarding table.

All packets are forwarded at Layer 3 in the ASIC hardware.

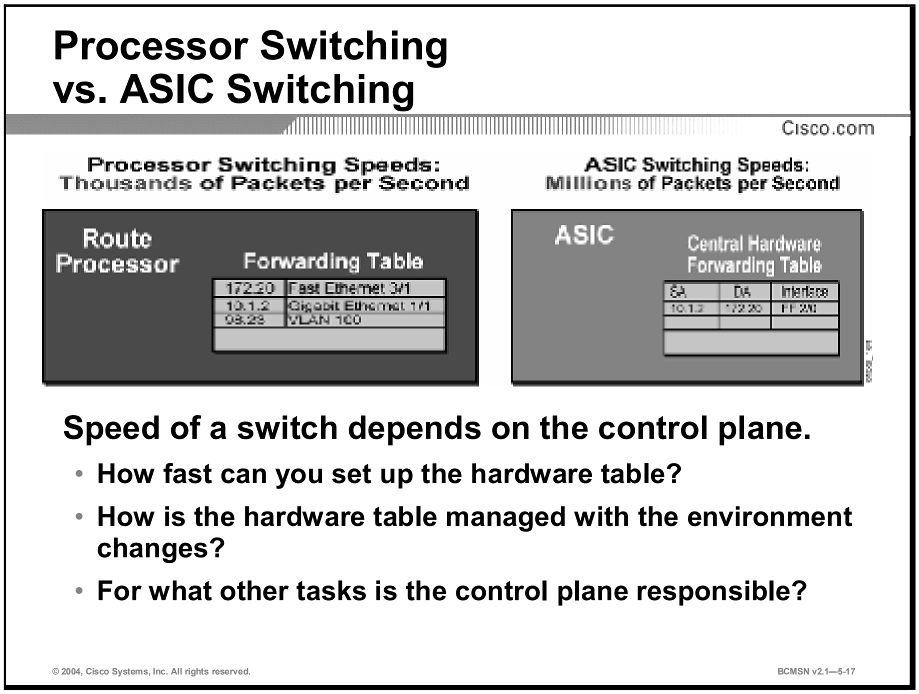

Processor Switching vs. ASIC Switching

Cisco.com

Processor Switching Speeds:
Thousands of Packets per Second

Route Processor

Forwarding Table

172.20	Fast Ethernet 3/1
10.1.2	Gigabit Ethernet 1/1
08.23	VLAN 100

ASIC Switching Speeds:
Millions of Packets per Second

ASIC

Central Hardware Forwarding Table

SA	DA	Interface
10.1.2	172.20	FE 2/0

Speed of a switch depends on the control plane.

- **How fast can you set up the hardware table?**
- **How is the hardware table managed with the environment changes?**
- **For what other tasks is the control plane responsible?**

BCMSN v2.1—5-17

Layer 3 switching software employs a distributed architecture in which the control path and data path are relatively independent. The control path code, such as routing protocol, runs on the processor, whereas the data packets are switched by ASICs, either centralized or distributed. The interfaces provide access to the media, and the switching fabric provides the transport between interfaces.

A microcoded ASIC handles packet switching. The main functions of the control layer between the routing protocol and the firmware data path microcode are as follows:

- Managing the internal data and control circuits for the packet forwarding and control functions
- Extracting the other routing and packet forwarding-related control information from the Layer 2 and Layer 3 bridging and routing protocols and the configuration data, and then conveying the information to the ASICs to control the data path
- Collecting the data path information, such as traffic statistics, from the interfaces to the processor
- Handling certain data packets sent from the interfaces to the processor

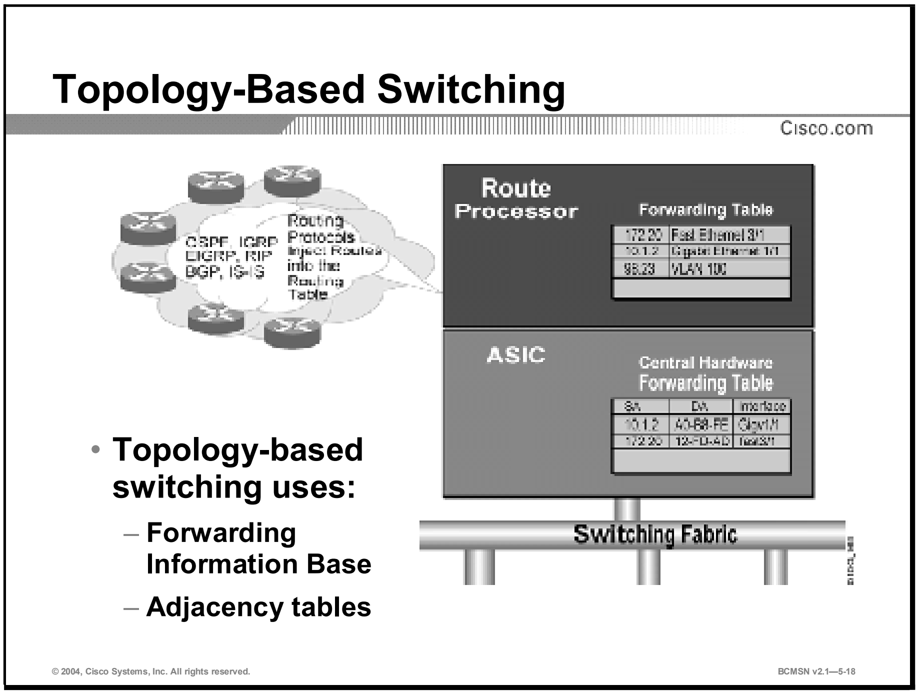

Topology-Based Switching

Cisco.com

OSPF, IGRP
EIGRP, RIP
BGP, IS-IS
Routing Protocols Inject Routes into the Routing Table
Route Processor
Forwarding Table
ASIC
Central Hardware Forwarding Table
Switching Fabric

172.20	Fast Ethernet 3/1
10.1.2	Gigabit Ethernet 1/1
98.23	VLAN 100

SA	DA	Interface
10.1.2	A0-B8-FE	Gig1/1
172.20	12-FD-AD	fast3/1

- **Topology-based switching uses:**
  - **Forwarding Information Base**
  - **Adjacency tables**

BCMSN v2.1—5-18

In CEF, the control plane data is installed in hardware. CEF scales to large networks and is not based on traffic flow.

The two main components of CEF operation are the following:

- **FIB:** CEF uses a FIB to make IP destination prefix-based switching decisions. The FIB is conceptually similar to a routing table or information base. The FIB maintains a mirror image of the forwarding information contained in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and those changes are reflected in the FIB. The FIB maintains next-hop address information based on the information in the IP routing table.
- **Adjacency tables:** Network nodes in the network are said to be adjacent if they can reach each other with a single hop across a link layer. In addition to the FIB, CEF uses adjacency tables to prepend Layer 2 addressing information. The adjacency table maintains Layer 2 next-hop addresses for all FIB entries.

You can maintain the FIB centrally or you can distribute it.

Summary

This topic summarizes the key points discussed in this lesson.

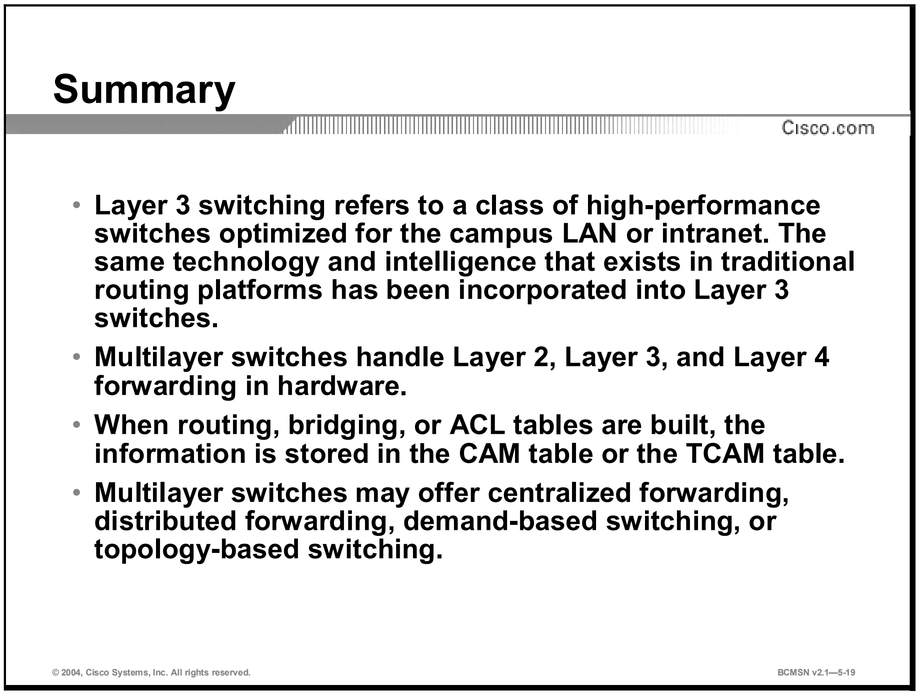

Summary

Cisco.com

- Layer 3 switching refers to a class of high-performance switches optimized for the campus LAN or intranet. The same technology and intelligence that exists in traditional routing platforms has been incorporated into Layer 3 switches.
- Multilayer switches handle Layer 2, Layer 3, and Layer 4 forwarding in hardware.
- When routing, bridging, or ACL tables are built, the information is stored in the CAM table or the TCAM table.
- Multilayer switches may offer centralized forwarding, distributed forwarding, demand-based switching, or topology-based switching.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—5-19

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which type of Layer 3 switching uses a FIB?

A) route caching

B) flow-based switching

C) demand-based switching

D) topology-based switching

Q2) Which two operations are performed when a multilayer switch receives a packet? (Choose two.)

A) The Layer 3 engine performs a FIB table lookup.

B) The Layer 3 engine performs the input QoS ACL lookup.

C) The Layer 3 lookup engine looks up the destination MAC address.

D) The Layer 2 lookup engine looks up the destination MAC address.

E) The Layer 2 forwarding engine performs the outbound security ACL lookup.

Q3) Which choice describes a TCAM mask associated with the entry **access-list 101 deny ip 10.1.0.0 0.0.255.255 any**?

A) all 32 bits of source address

B) most significant 8 bits of source address

C) most significant 24 bits of source address

D) most significant 16 bits of source address

Q4) Which element manages the forwarding tables in a distributed forwarding architecture?

A) the distributed ASICs

B) each individual module

C) the central forwarding table

D) the primary forwarding engine

Quiz Answer Key

Q1) D

Relates to: Introducing Layer 3 Switching

Q2) A, D

Relates to: Multilayer Switch Packet Forwarding

Q3) D

Relates to: Switching Table Architectures

Q4) D

Relates to: Switch Forwarding Architectures

Configuring Multilayer Switching

Overview

CEF is enabled by default on the latest Cisco products to optimize switching and forwarding in a multilayer switched network. This lesson describes how to configure, verify, and troubleshoot CEF.

Relevance

Using the fastest switching mechanism available helps optimize performance in your multilayer switched network. Configuring CEF-based MLS will help you achieve optimal performance in your network.

Objectives

Upon completing this lesson, you will be able to:

- Explain the features and operation of CEF-based MLS
- Configure and verify CEF
- Troubleshoot CEF

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- CEF-Based MLS
- Configuring and Verifying CEF
- Troubleshooting CEF
- Summary
- Quiz

CEF-Based MLS

CEF-based MLS is a forwarding model implemented on the latest generation of Cisco multilayer switches. CEF-based MLS is topology-based. The control plane information is forwarded to the data plane at the port or line card in hardware. All packets are then switched in hardware with ASICs. This topic explains the features and operation of CEF-based MLS.

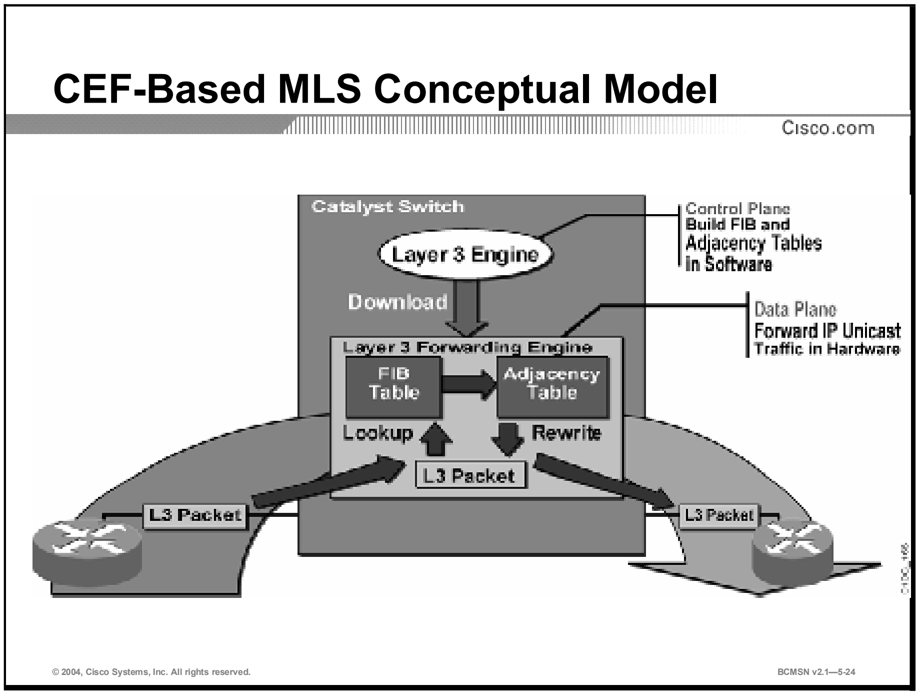

CEF-Based MLS Conceptual Model

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—5-24

CEF separates the control plane hardware from the data plane hardware and switching. ASICs in switches are used to separate the control plane and data plane, thereby achieving higher data throughput. The control plane is responsible for building the FIB table and adjacency tables in software. The data plane is responsible for forwarding IP unicast traffic using hardware.

Software switching occurs when traffic cannot be processed in hardware. The following types of exception packets are processed in software at a much slower rate:

- Packets that use IP header options (Packets that use TCP header options are switched in hardware because they do not affect the forwarding decision.)
- Packets that have an expiring IP TTL counter
- Packets that are forwarded to a tunnel interface
- Packets that arrive with nonsupported encapsulation types
- Packets that are routed to an interface with nonsupported encapsulation types
- Packets that exceed the maximum transmission unit (MTU) of an output interface and must be fragmented
- Packets that require an Internet Group Management Protocol (IGMP) redirect to be routed
- 802.3 Ethernet packets

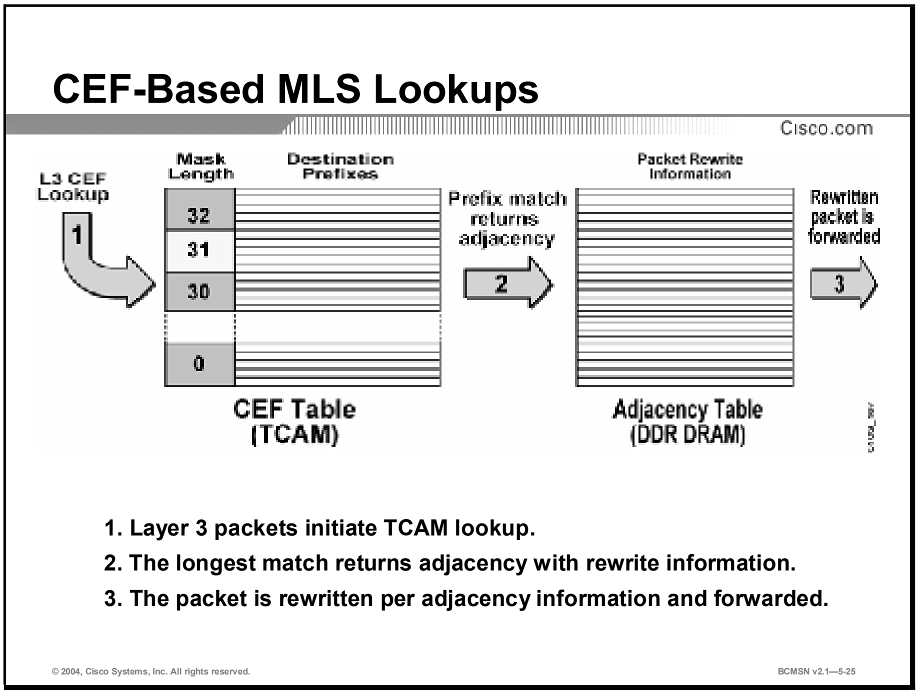

CEF-Based MLS Lookups

Cisco.com

L3 CEF Lookup
1
Mask Length
Destination Prefixes
32
31
30
0
CEF Table (TCAM)
Prefix match returns adjacency
2
Packet Rewrite Information
Adjacency Table (DDR DRAM)
Rewritten packet is forwarded
3

1. **Layer 3 packets initiate TCAM lookup.**
2. **The longest match returns adjacency with rewrite information.**
3. **The packet is rewritten per adjacency information and forwarded.**

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—5-25

CEF-based MLS consists of these features:

- The FIB lookup is based on the Layer 3 destination address prefix (longest match).
- The FIB is derived from the IP routing table and is arranged for maximum lookup throughput.
- The adjacency table is derived from the Address Resolution Protocol (ARP) table, and it contains Layer 2 rewrite (MAC) information for the next hop.
- CEF IP destination prefixes are stored in the TCAM table from the most specific to the least specific entry.
- Adjacency (rewrite) information and statistics are maintained in double-data-rate onboard DRAM.
- CEF maintains a one-to-one CEF-to-adjacency mapping for accurate statistics tracking.
- When the CEF TCAM table is full, a wildcard entry redirects to the Layer 3 engine. The nonstatic host routes discarded packets first. More specific mask lengths are honored.
- When the adjacency table is full, a CEF TCAM table entry points to the Layer 3 engine to redirect the adjacency.

FIB Table Updates

The FIB table is updated when the following occurs:

- An ARP entry for the destination next hop changes, ages out, or is removed.
- The routing table entry for a prefix changes.
- The routing table entry for the next hop changes.

CEF-Based MLS Operation

Cisco.com

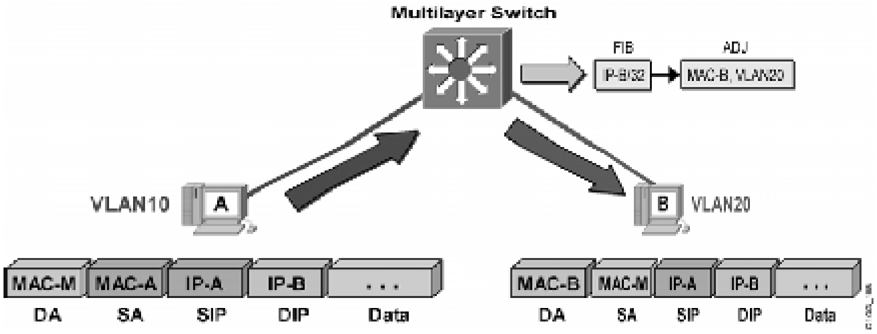

BCMSN v2.1—5-26

The figure provides an example of CEF-based MLS operation. These actions occur:

Step 1 The Layer 3 engine queries the switch for a physical MAC address.

Step 2 The switch selects a MAC address from the chassis MAC range and assigns it to the Layer 3 engine. This MAC address is assigned by the Layer 3 engine as its burned-in address for all VLANs and is used by the switch to initiate Layer 3 packet lookups.

Step 3 The switch installs wildcard CEF entries, which point to drop adjacencies (for handling CEF table lookup misses).

Step 4 The Layer 3 engine informs the switch of its interfaces participating in MLS (MAC address and associated VLAN). The switch creates the (MAC, VLAN) Layer 2 CAM entry for the Layer 3 engine.

Step 5 The Layer 3 engine informs the switch about features for interfaces participating in MLS.

Step 6 The Layer 3 engine informs the switch about all CEF entries related to its interfaces and connected networks. The switch populates the CEF entries and points them to Layer 3 engine redirect adjacencies.

Step 7 Host A sends a packet to host B. The switch recognizes the frame as a Layer 3 packet because the destination MAC (MAC-M) matches the Layer 3 engine MAC.

Step 8 The switch performs a CEF lookup based on the destination IP address (IP-B). The packet hits the CEF entry for the connected (VLAN20) network and is redirected to the Layer 3 engine using a “glean” adjacency.

Step 9 The Layer 3 engine installs an ARP throttling adjacency in the switch for the host B IP address.

Step 10 The Layer 3 engine sends ARP requests for host B on VLAN20.

Step 11 Host B sends an ARP response to the Layer 3 engine.

Step 12 The Layer 3 engine installs the resolved adjacency in the switch (removing ARP throttling adjacency).

Step 13 The switch forwards the packet to host B.

Step 14 The switch receives a subsequent packet for host B (IP-B).

Step 15 The switch performs a Layer 3 lookup and finds a CEF entry for host B. The entry points to the adjacency with rewrite information for host B.

Step 16 The switch rewrites packets per the adjacency information and forwards the packet to host B on VLAN20.

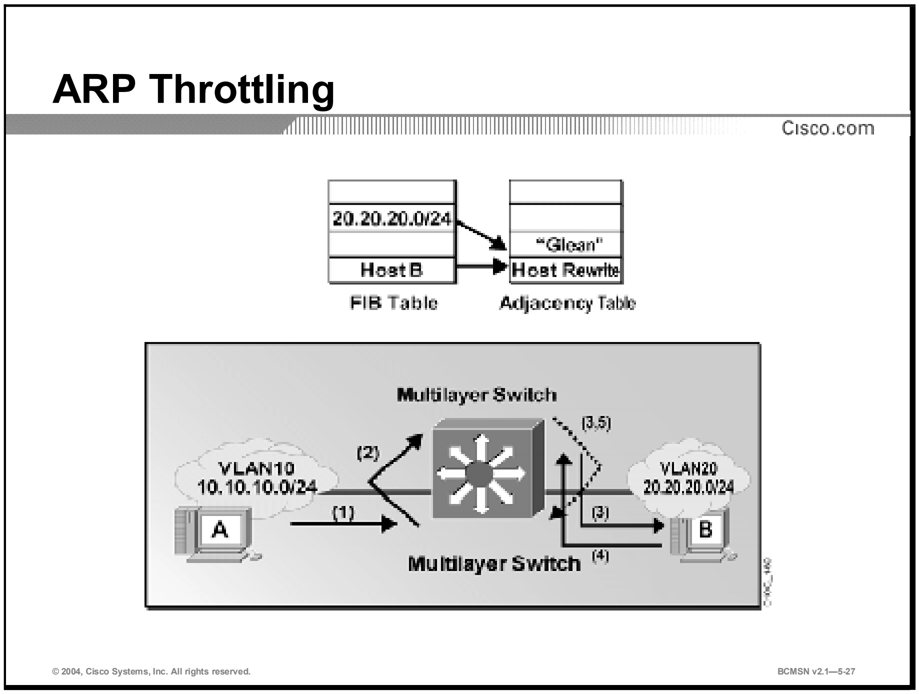

ARP Throttling

Cisco.com

BCMSN v2.1—5-27

Only the first few packets for a connected destination reach the Layer 3 engine so that the Layer 3 engine can use ARP to locate the host. Throttling adjacency is installed so that subsequent packets to that host are dropped in hardware until an ARP response is received. The throttling adjacency is removed when an ARP reply is received (and a complete rewrite adjacency is installed for the host). The switch removes throttling adjacency if no ARP reply is seen within 2 seconds to allow more packets through to reinitiate ARP. This relieves the Layer 3 engine from excessive ARP processing or from ARP-based denial of service attacks.

The figure provides an example of ARP throttling, which consists of these steps:

Step 1 Host A sends packet to host B.

Step 2 The switch forwards the packet to the Layer 3 engine based on the "glean" entry in the FIB.

Step 3 The Layer 3 engine sends an ARP request for host B and installs the drop adjacency for host B.

Step 4 Host B responds to the ARP request.

Step 5 The Layer 3 engine installs adjacency for host B and removes the drop adjacency.

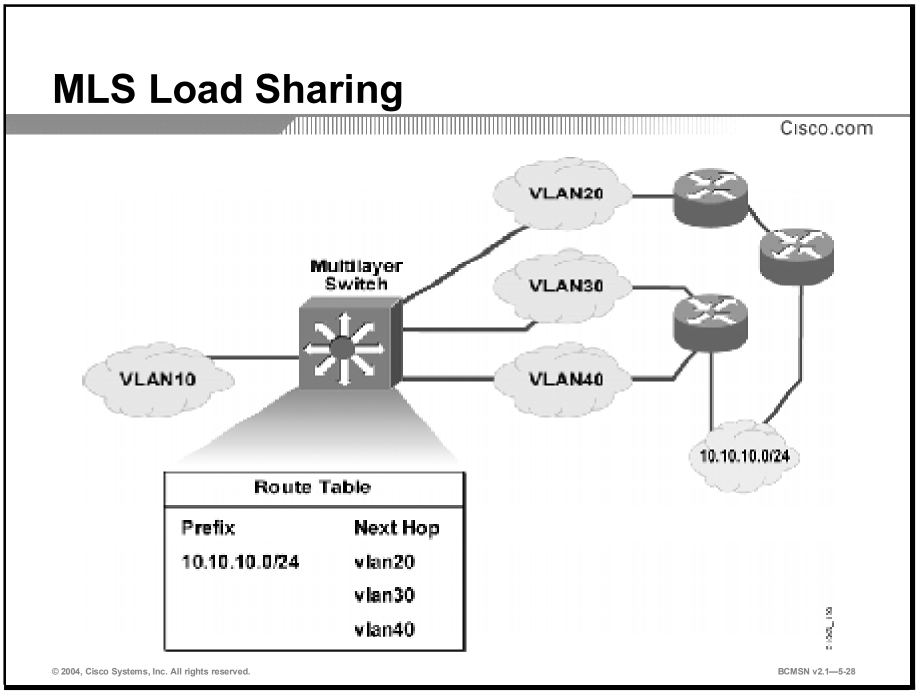

MLS Load Sharing

Cisco.com

Route Table

Prefix	Next Hop
10.10.10.0/24	vlan20
	vlan30
	vlan40

BCMSN v2.1—5-28

Per-flow load sharing (equal-cost or nonequal-cost) for MLS is supported in Catalyst multilayer switch hardware. A single FIB entry can point to up to six adjacencies for load sharing.

The selection of an adjacency is calculated as a hash function based on the following:

- Source IP address
- Destination IP address
- Source and destination IP Layer 4 ports

No special configuration is required. Simply configure IP routing on the Layer 3 engine.

Configuring and Verifying CEF

Cisco Catalyst multilayer switches have CEF enabled by default. You can view statistics about interfaces to determine how CEF is operating. This topic explains the default CEF states and how to view information about CEF.

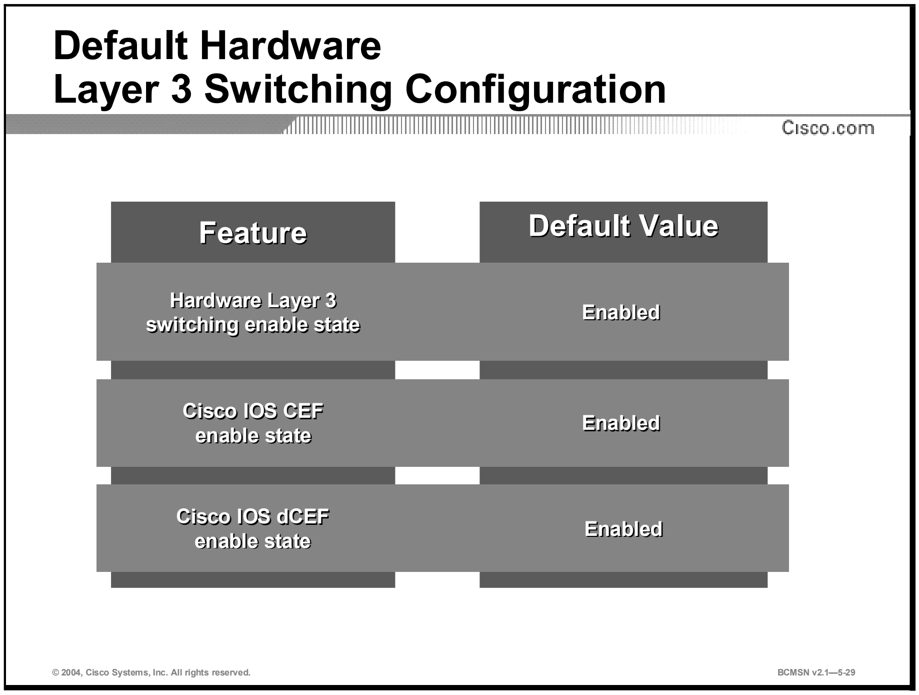

Default Hardware Layer 3 Switching Configuration

Cisco.com

Feature	Default Value
Hardware Layer 3 switching enable state	Enabled
Cisco IOS CEF enable state	Enabled
Cisco IOS dCEF enable state	Enabled

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—5-29

Hardware Layer 3 switching is permanently enabled on Catalyst 6500 series Supervisor Engine 2 with Policy Feature Card 2 (PFC2), Multilayer Switch Feature Card 2 (MSFC2), and Distributed Forwarding Card (DFC). No configuration is required.

You can use the **no ip cef** command to disable CEF on the Catalyst 4000 or the **no ip route-cache cef** command on a Catalyst 3550 interface.

The default configuration, which Cisco recommends, is CEF enabled on all Layer 3 interfaces. If you disable CEF on an interface, you can enable CEF as follows:

- On the Catalyst 3550 switch, use the **ip route-cache cef** interface configuration command to enable CEF on an interface.
- On the Catalyst 4000 switch, use the **ip cef** interface configuration command to enable CEF on an interface after it has been disabled.
- On the Catalyst 6500 with PFC2, DFCs, and MSFC2, you cannot disable CEF.

Per-destination load balancing is enabled by default when you enable CEF. To use per-destination load balancing, you do not perform any additional tasks after you enable CEF.

Per-destination load balancing allows the router to use multiple paths to achieve load sharing. Packets for a given source-destination host pair are guaranteed to take the same path, even if multiple paths are available. Per-destination load balancing ensures packets for a given host pair arrive in order. Traffic destined for different pairs tend to take different paths. Per-destination load balancing is enabled by default when you enable CEF, and it is the load-balancing method of choice for most situations.

Because per-destination load balancing depends on the statistical distribution of traffic, load sharing becomes more effective as the number of source-destination pairs increase.

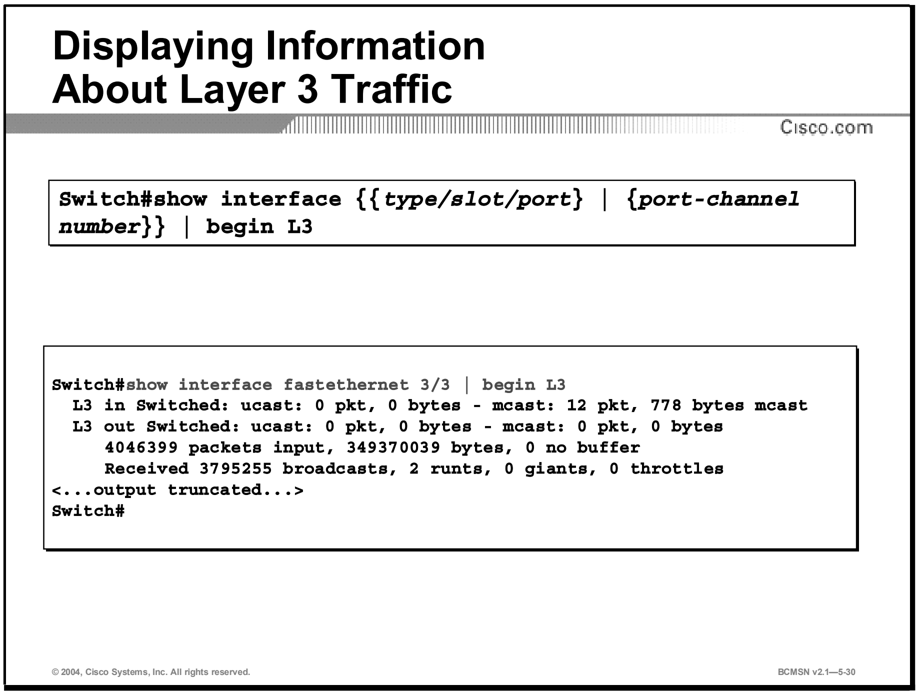

Displaying Information About Layer 3 Traffic

Cisco.com

```
Switch#show interface {{type/slot/port} | {port-channel
number}} | begin L3
```

```
Switch#show interface fastethernet 3/3 | begin L3
  L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 12 pkt, 778 bytes mcast
  L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
     4046399 packets input, 349370039 bytes, 0 no buffer
     Received 3795255 broadcasts, 2 runts, 0 giants, 0 throttles
<...output truncated...>
Switch#
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—5-30

Use the **show interface** command to display information about Layer 3 switched traffic.

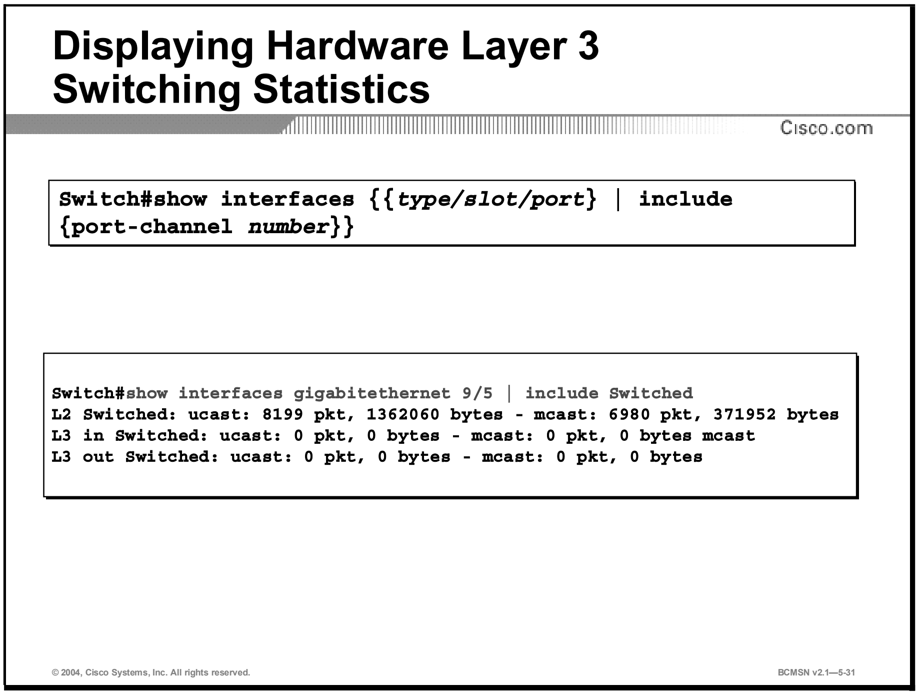

Displaying Hardware Layer 3 Switching Statistics

Cisco.com

```
Switch#show interfaces {{type/slot/port} | include
{port-channel number}}
```

```
Switch#show interfaces gigabitethernet 9/5 | include Switched
L2 Switched: ucast: 8199 pkt, 1362060 bytes - mcast: 6980 pkt, 371952 bytes
L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes mcast
L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—5-31

Hardware Layer 3 switching statistics are obtained on a per-VLAN basis. To display hardware Layer 3 switching statistics, use the **show interfaces** command.

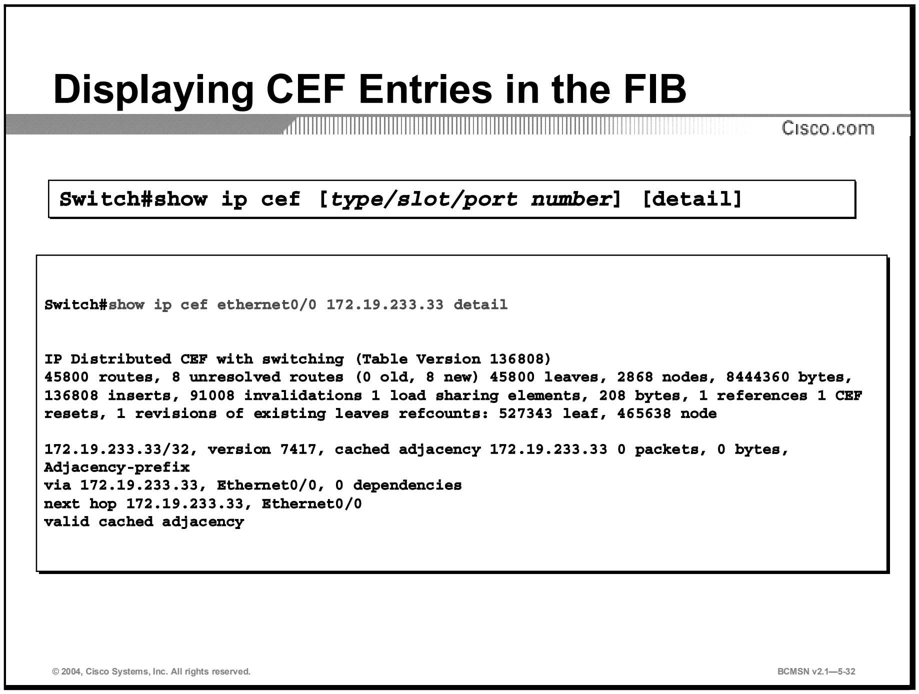

Displaying CEF Entries in the FIB

Cisco.com

```
Switch#show ip cef [type/slot/port number] [detail]
```

```
Switch#show ip cef ethernet0/0 172.19.233.33 detail

IP Distributed CEF with switching (Table Version 136808)
45800 routes, 8 unresolved routes (0 old, 8 new) 45800 leaves, 2868 nodes, 8444360 bytes,
136808 inserts, 91008 invalidations 1 load sharing elements, 208 bytes, 1 references 1 CEF
resets, 1 revisions of existing leaves refcounts: 527343 leaf, 465638 node

172.19.233.33/32, version 7417, cached adjacency 172.19.233.33 0 packets, 0 bytes,
Adjacency-prefix
via 172.19.233.33, Ethernet0/0, 0 dependencies
next hop 172.19.233.33, Ethernet0/0
valid cached adjacency
```

BCMSN v2.1—5-32

The **show ip cef** command shows a display of all FIB entries.

The **show ip cef detail** command shows detailed FIB entry information for all FIB entries.

Note The **show ip cef** command is the only CEF **show** command available on the Catalyst 3550.

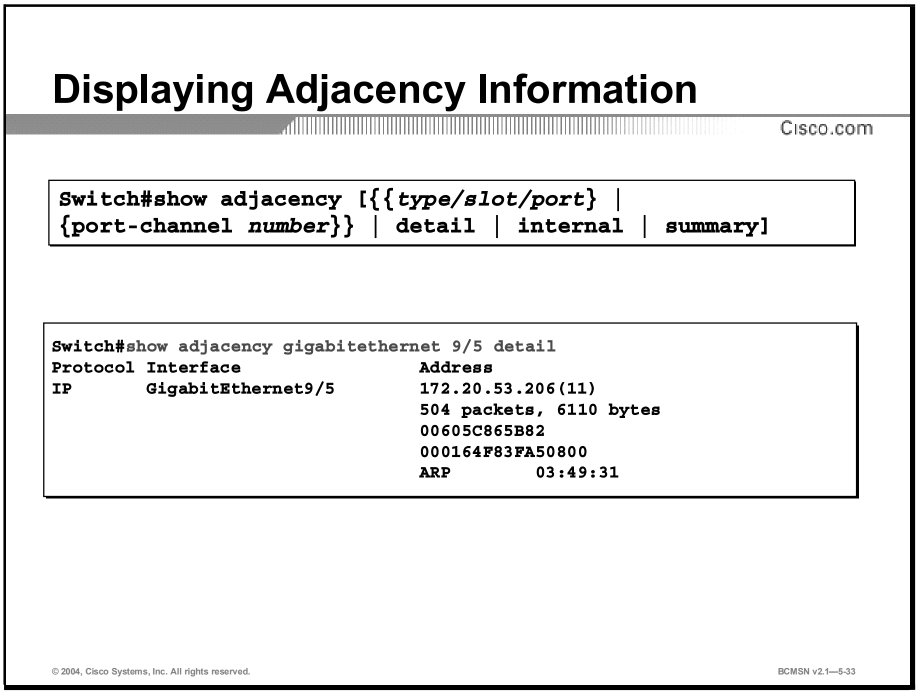

Displaying Adjacency Information

Cisco.com

```
Switch#show adjacency [{{type/slot/port} |
{port-channel number}} | detail | internal | summary]
```

```
Switch#show adjacency gigabitethernet 9/5 detail
Protocol Interface                Address
IP       GigabitEthernet9/5       172.20.53.206(11)
                                  504 packets, 6110 bytes
                                  00605C865B82
                                  000164F83FA50800
                                  ARP        03:49:31
```

BCMSN v2.1—5-33

The adjacency table is populated as adjacencies are discovered. Each time an adjacency entry is created (such as through the ARP protocol), a link-layer header for that adjacent node is precomputed and stored in the adjacency table. After a route is determined, it points to a next hop and corresponding adjacency entry. The route is subsequently used for encapsulation during CEF switching of packets.

A route might have several paths to a destination prefix, such as when a router is configured for simultaneous load balancing and redundancy. For each resolved path, a pointer is added for the adjacency corresponding to the next-hop interface for that path. This mechanism is used for load balancing across several paths.

In addition to adjacencies associated with next-hop interfaces (host-route adjacencies), other types of adjacencies are used to expedite switching when certain exception conditions exist. When the prefix is defined, prefixes requiring exception processing are cached with one of the following special adjacencies:

- **Null adjacency:** Packets destined for a “Null0” interface are dropped. This can be used as an effective form of access filtering.
- **Glean adjacency:** When a router is connected directly to several hosts, the FIB table on the router maintains a prefix for the subnet rather than for the individual host prefixes. The subnet prefix points to a glean adjacency. When packets need to be forwarded to a specific host, the adjacency database is gleaned for the specific prefix.
- **Punt adjacency:** Features that require special handling, or features that are not yet supported in conjunction with CEF switching paths, are forwarded to the next switching layer for handling; for example, the packet may require CPU processing. Features that are not supported are forwarded to the next higher switching level.
- **Discard adjacency:** Packets are discarded.
- **Drop adjacency:** Packets are dropped, but the prefix is checked.

When a link-layer header is prepended to packets, FIB requires the prepend to point to an adjacency corresponding to the next hop. If an adjacency was created by FIB and not discovered through a mechanism such as ARP, the Layer 2 addressing information is not known and the adjacency is considered incomplete. After the Layer 2 information is known, the packet is forwarded to the route processor, and the adjacency is determined through ARP.

To display adjacency table information, use the **show adjacency** command. The optional **detail** keyword displays detailed adjacency information, including Layer 2 information. Adjacency statistics are updated approximately every 60 seconds.

Troubleshooting CEF

You can use **debug** commands to display debugging information about CEF. This topic discusses troubleshooting CEF.

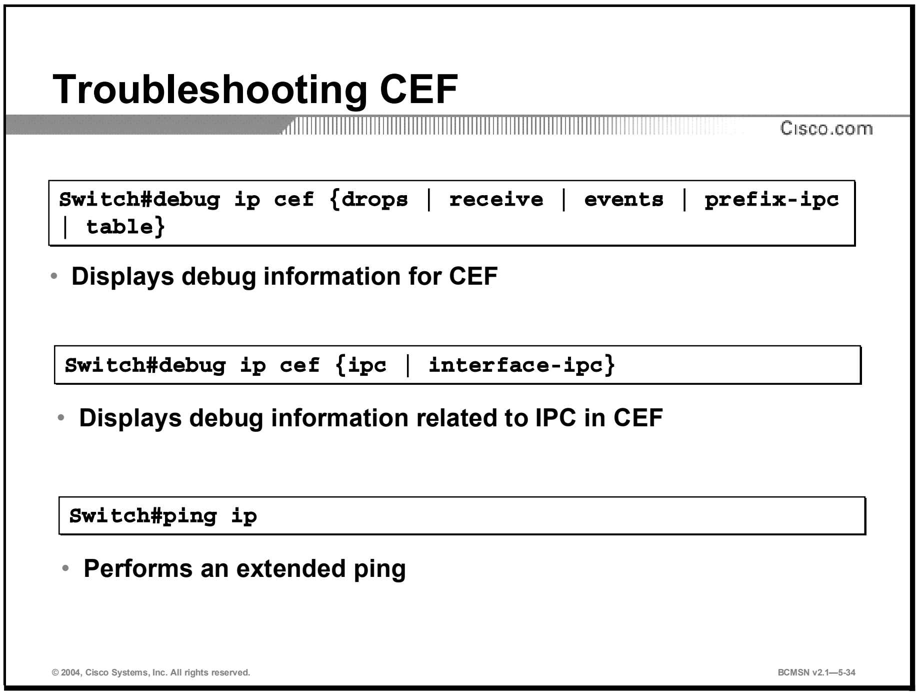

Troubleshooting CEF

Cisco.com

```
Switch#debug ip cef {drops | receive | events | prefix-ipc
| table}
```

- **Displays debug information for CEF**

```
Switch#debug ip cef {ipc | interface-ipc}
```

- **Displays debug information related to IPC in CEF**

```
Switch#ping ip
```

- **Performs an extended ping**

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—5-34

Use the **debug ip cef** EXEC commands for troubleshooting CEF. The syntax is as follows:

```
debug ip cef {drops [access-list] | receive [access-list] |
events [access-list] | prefix-ipc [access-list] | table
[access-list]}

debug ip cef {ipc | interface-ipc}
```

The arguments to the **debug** command include the following:

- **drops:** Records dropped packets
- **access-list (Optional):** Controls collection of debugging information from specified lists
- **receive:** Records packets that are not switched using information from the FIB table, but that are received and sent to the next switching layer
- **events:** Records general CEF events
- **prefix-ipc:** Records updates related to IP prefix information, including the following:
  - Debugging of IP routing updates in a line card
  - Reloading of a line card with a new table
  - Adding a route update from the route processor to the line card exceeds the maximum number of routes
  - Control messages related to FIB table prefixes

- **table:** Produces a table showing events related to the FIB table. Possible types of events include the following:
  - Routing updates that populate the FIB table
  - Flushing of the FIB table
  - Adding or removing of entries to the FIB table
  - Table reloading process
- **ipc:** Records information related to interprocess communications (IPC) in CEF. Possible types of events include the following:
  - Transmission status of IPC messages
  - Status of buffer space for IPC messages
  - IPC messages received out of sequence
  - Status of resequenced messages
  - Throttle requests sent from a line card to the route processor
- **interface-ipc:** Records IPC updates related to interfaces. Possible reporting includes an interface coming up or going down and updates to fibhwidb, fibidb, and so forth.

When a normal **ping** command is sent, the source address of the ping is the IP address of the interface that the packet uses to exit the switch. If you use an extended **ping** command, you can change the source IP address to any IP address on the switch. The extended ping is used to perform a more advanced check of host reachability and network connectivity. The extended **ping** command works only at the privileged EXEC command line. The normal **ping** command works both in the user EXEC mode and the privileged EXEC mode. To use the extended ping feature, enter **ping ip** at the command line and press **Return**. You are prompted for the fields described in the table.

Field	Description
Target IP address	Prompts for the IP address or host name of the destination node you plan to ping. If you have specified a supported protocol other than IP, enter an appropriate address for that protocol here.
Repeat count 5	Number of ping packets that will be sent to the destination address. The default is five packets.
Datagram size	Size of the ping packet (in bytes).
Timeout in seconds	Timeout interval. Default: 2 (seconds). The ping is declared successful only if the echo reply packet is received before this time interval.
Extended commands	Specifies whether a series of additional commands appears.
Source address or interface	The interface or IP address of the router to use as a source address for the probes. The router normally picks the IP address of the outbound interface to use. The interface may also be mentioned, but with the correct syntax.
Type of service	Specifies the type of service (ToS). The requested ToS is placed in each probe, but there is no guarantee that all routers will process the ToS. It depends on the quality selection of the Internet service.
Set DF bit in IP header?	Specifies whether the do not fragment (DF) bit is to be set on the ping packet. If yes is specified, the DF option does not allow this packet to be fragmented when it has to go through a segment with a smaller MTU. You will receive an error message from the device that wanted to fragment the packet. This is useful for determining the smallest MTU in the path to a destination.
Validate reply data?	Specifies whether to validate the reply data.
Data pattern	Specifies the data pattern. Different data patterns are used to troubleshoot framing errors and clocking problems on serial lines.
Loose, Strict, Record, Timestamp, Verbose	IP header options. This prompt offers more than one option to be selected.
Sweep range of sizes	Allows you to vary the sizes of the echo packets being sent. This is used to determine the minimum sizes of the MTUs configured on the nodes along the path to the destination address. Performance problems caused by packet fragmentation are thus reduced.
!!!!!	Each exclamation point (!) denotes receipt of a reply. A period (.) denotes that the network server timed out while waiting for a reply.
Success rate is 100 percent	Percentage of packets successfully echoed back to the switch. Anything less than 80 percent is usually considered problematic.
round-trip min/avg/max = 1/2/4 ms	Round-trip travel time intervals for the protocol echo packets, including minimum/average/maximum (in milliseconds).

Summary

This topic summarizes the key points discussed in this lesson.

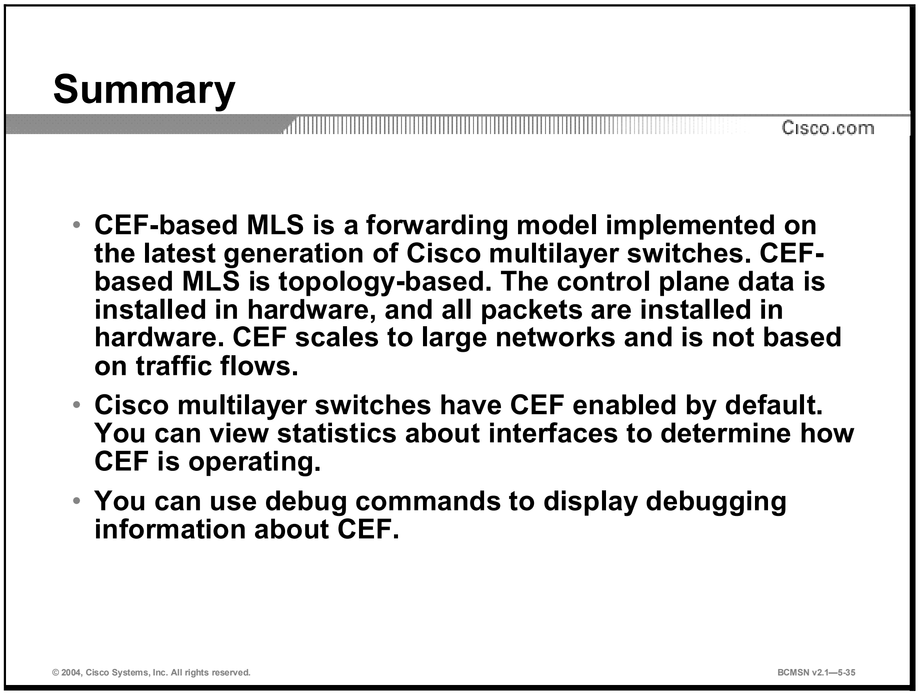

Summary

Cisco.com

- CEF-based MLS is a forwarding model implemented on the latest generation of Cisco multilayer switches. CEF-based MLS is topology-based. The control plane data is installed in hardware, and all packets are installed in hardware. CEF scales to large networks and is not based on traffic flows.
- Cisco multilayer switches have CEF enabled by default. You can view statistics about interfaces to determine how CEF is operating.
- You can use debug commands to display debugging information about CEF.

BCMSN v2.1—5-35

References

For additional information, refer to this resource:

- Software configuration guide for your Cisco Catalyst switch

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which statement accurately describes CEF-based MLS?

A) Each flow requires a cache entry.

B) All traffic switching is flow-based.

C) All packets are handled in hardware.

D) The first packet in each flow is handled in software.

Q2) In CEF-based multilayer switching, which type of adjacency entry is used for features that require special handling or for features that are not yet supported in conjunction with CEF switching paths?

A) null adjacency

B) punt adjacency

C) glean adjacency

D) next-hop adjacency

Q3) Which CEF **debug** command displays information about control messages related to FIB table prefixes?

A) **debug ip cef ipc**

B) **debug ip cef table**

C) **debug ip cef prefix-ipc**

D) **debug ip cef interface-ipc**

Quiz Answer Key

Q1) C

Relates to: CEF-Based MLS

Q2) B

Relates to: Configuring and Verifying CEF

Q3) C

Relates to: Troubleshooting CEF

Routing Between VLANs

Overview

Network devices in different VLANs cannot communicate with one another without a Layer 3 switch or router to forward traffic between the VLANs. In most network environments, VLANs are associated with individual networks or subnetworks.

Relevance

When an end station in one VLAN needs to communicate with an end station in another VLAN, inter-VLAN communication is required. You configure one or more devices to route traffic to the appropriate destination VLAN.

Objectives

Upon completing this lesson, you will be able to:

- Describe and configure the different interface types on a multilayer switch for routing between VLANs
- Explain the operation and configuration of inter-VLAN routing
- Explain the operation and configuration of router on a stick
- Verify the inter-VLAN routing configuration

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Connecting VLANs with Multilayer Switches
- Inter-VLAN Routing
- Router on a Stick
- Verifying the Inter-VLAN Routing Configuration
- Summary
- Quiz

Connecting VLANs with Multilayer Switches

Multilayer switches support different types of interfaces that you can use to route between VLANs. This topic covers the different port types on multilayer switches and explains how to configure those ports to support routing between VLANs.

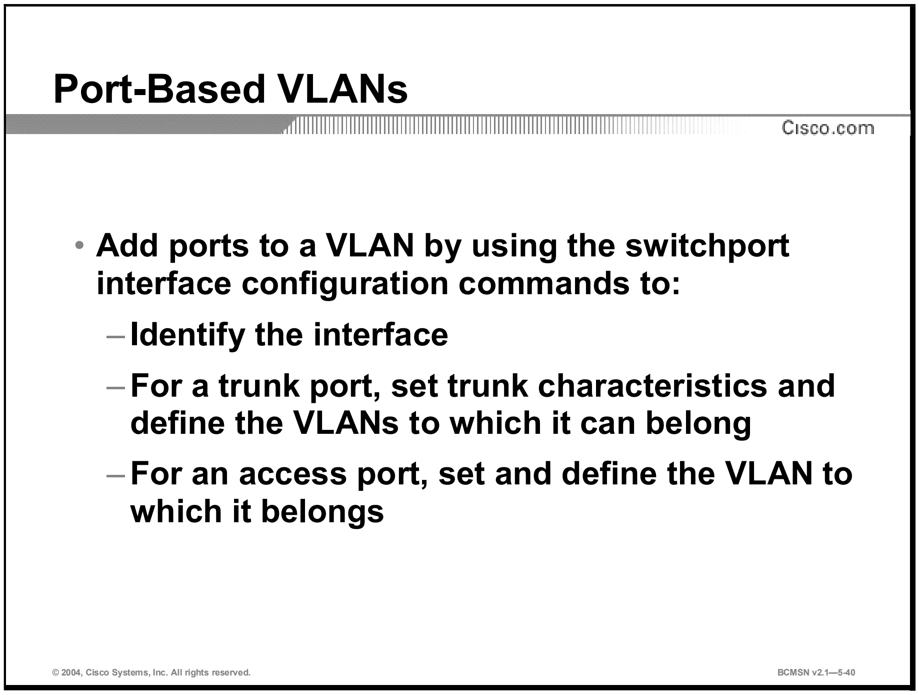

Port-Based VLANs

- **Add ports to a VLAN by using the switchport interface configuration commands to:**
  - **Identify the interface**
  - **For a trunk port, set trunk characteristics and define the VLANs to which it can belong**
  - **For an access port, set and define the VLAN to which it belongs**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—5-40

A VLAN is a switched network that is logically segmented by function, team, or application, without regard to the physical location of the users. Frames received on a port are forwarded only to ports that belong to the same VLAN as the receiving port. Network devices in different VLANs cannot communicate with one another without a Layer 3 device to route traffic between the VLANs.

VLAN partitions provide hard firewalls for traffic in the VLAN, and each VLAN has its own MAC address table. A VLAN comes into existence when a local port is configured to be associated with the VLAN, when the VLAN Trunking Protocol (VTP) learns of its existence from a neighbor on a trunk, or when a user adds a VLAN to the local VTP database.

Add ports to a VLAN by using the **switchport** interface configuration commands to do the following:

- Identify the interface.
- For a trunk port, set trunk characteristics, and if desired, define the VLANs to which it can belong.
- For an access port, set and define the VLAN to which it belongs.

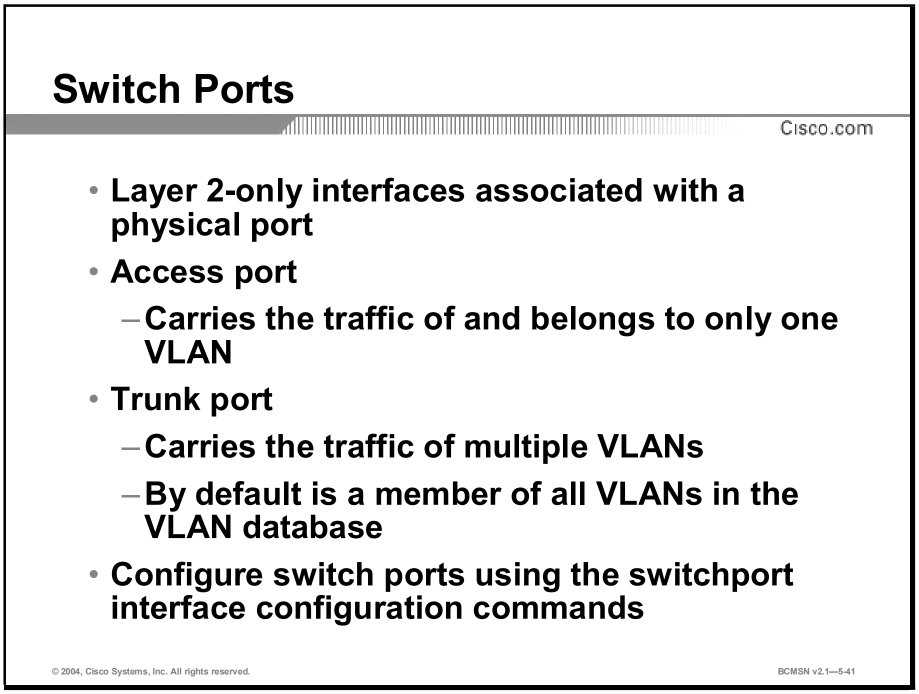

Switch Ports

- Layer 2-only interfaces associated with a physical port
- Access port
  - Carries the traffic of and belongs to only one VLAN
- Trunk port
  - Carries the traffic of multiple VLANs
  - By default is a member of all VLANs in the VLAN database
- Configure switch ports using the switchport interface configuration commands

BCMSN v2.1—5-41

Switch ports are Layer 2-only interfaces associated with a physical port. A switch port can be either an access port or a trunk port. You can configure a port as an access port or trunk port; or, you can let the Dynamic Trunking Protocol (DTP) operate on a per-port basis to determine if a switch port should be an access port or a trunk port. Switch ports are used for managing the physical interface and associated Layer 2 protocols and do not handle routing or bridging.

Configure switch ports (access ports and trunk ports) by using the **switchport** interface configuration command.

Access Ports

An access port carries the traffic of, and belongs to, only one VLAN. Traffic is received and sent in native formats with no VLAN tagging. Traffic arriving on an access port is assumed to belong to the VLAN assigned to the port. If an access port receives a tagged frame (Inter-Switch Link [ISL] or 802.1Q tagged), the frame is dropped, the source address is not learned, and the frame is counted in the “no destination” statistic.

These two types of access ports are supported:

- **Static access ports:** These access ports are manually assigned to a VLAN.
- **Dynamic access ports**: These access ports learn of their VLAN membership through incoming frames. By default, a dynamic access port is a member of no VLAN, and forwarding to and from the port is enabled only when a VLAN membership of the port is discovered.

Trunk Ports

A trunk port carries the traffic of multiple VLANs. By default, a trunk port is a member of all VLANs in the VLAN database. These two types of trunk ports are supported:

- **ISL trunk port:** In an ISL trunk port, all received frames are expected to be encapsulated with an ISL header, and all transmitted frames are sent with an ISL header. Native (untagged) frames received from an ISL trunk port are dropped.
- **IEEE 802.1Q trunk port:** An IEEE 802.1Q trunk port supports simultaneous tagged and untagged traffic. An 802.1Q trunk port is assigned a default port VLAN ID (PVID), and all untagged traffic travels on the port default PVID. All untagged traffic and tagged traffic with a “NULL VLAN ID” are assumed to belong to the port default PVID. A frame with a VID equal to the outgoing port default PVID is sent untagged. All other traffic is sent with a VLAN tag.

Although by default, a trunk port is a member of every VLAN known to the VLAN Trunk Protocol (VTP), you can limit VLAN membership by configuring an allowed list of VLANs for each trunk port. The list of allowed VLANs does not affect any other port but the associated trunk port. By default, all possible VLANs (VID 1 to 1005) are in the allowed list. A trunk port can become a member of a VLAN only if VTP knows of the VLAN and the VLAN is in the enabled state. If VTP learns of a new, enabled VLAN and the VLAN is in the allowed list for a trunk port, the trunk port automatically becomes a member of that VLAN, and traffic is forwarded to and from the trunk port for that VLAN. If VTP learns of a new, enabled VLAN that is not in the allowed list for a trunk port, the port does not become a member of the VLAN, and no traffic for the VLAN is forwarded to or from the port.

Note VLAN1 cannot be excluded from the allowed list.

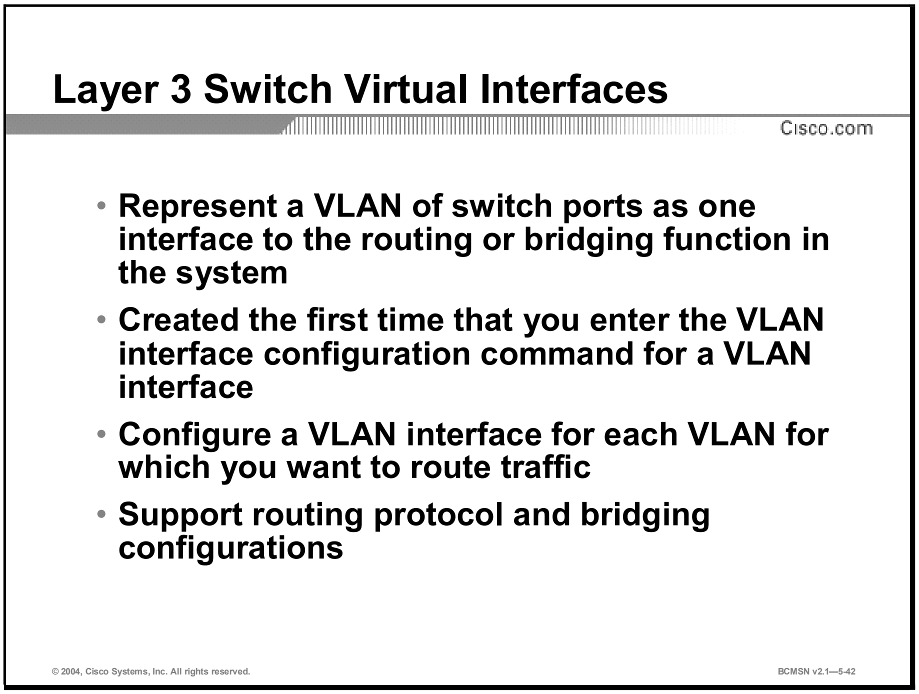

Layer 3 Switch Virtual Interfaces

- Represent a VLAN of switch ports as one interface to the routing or bridging function in the system
- Created the first time that you enter the VLAN interface configuration command for a VLAN interface
- Configure a VLAN interface for each VLAN for which you want to route traffic
- Support routing protocol and bridging configurations

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—5-42

A switched virtual interface (SVI) represents a VLAN of switch ports as one interface to the routing or bridging function in the system. Only one SVI can be associated with a VLAN. You configure an SVI for a VLAN for these reasons:

- When you route between VLANs
- When you have fallback-bridge nonroutable protocols between VLANs
- To provide IP host connectivity to the switch

By default, an SVI is created for the default VLAN (VLAN1) to permit remote switch administration. You must explicitly configure additional SVIs. In Layer 2 mode, SVIs provide IP host connectivity only to the system; in Layer 3 mode, you can configure routing across SVIs.

SVIs are created the first time that you enter the **vlan** interface configuration command for a VLAN interface. The VLAN corresponds to the VLAN tag associated with data frames on an ISL or 802.1Q encapsulated trunk or the VLAN ID configured for an access port. Configure a VLAN interface for each VLAN for which you want to route traffic and assign an IP address to the interface.

SVIs support routing protocol and bridging configurations.

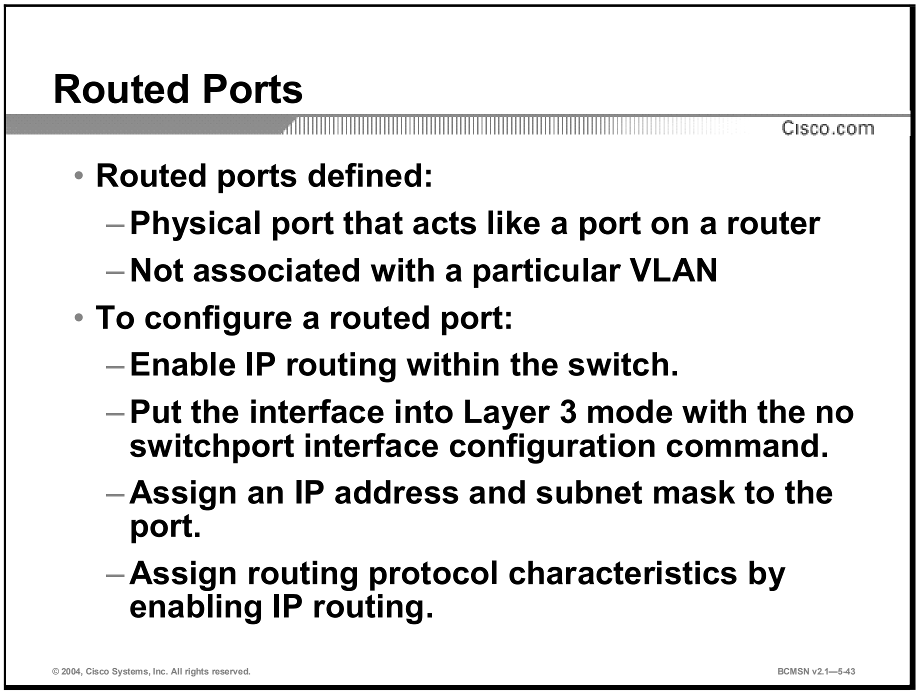

Routed Ports

Cisco.com

- Routed ports defined:
  - Physical port that acts like a port on a router
  - Not associated with a particular VLAN
- To configure a routed port:
  - Enable IP routing within the switch.
  - Put the interface into Layer 3 mode with the no switchport interface configuration command.
  - Assign an IP address and subnet mask to the port.
  - Assign routing protocol characteristics by enabling IP routing.

BCMSN v2.1—5-43

A routed port is a physical port that acts like a port on a router; a routed port does not have to be connected to a router. A routed port is not associated with a particular VLAN, as is an access port. A routed port behaves like a regular router interface, except that it does not support VLAN subinterfaces. Routed ports can be configured with a Layer 3 routing protocol.

Configure routed ports by putting the interface into Layer 3 mode with the **no switchport** interface configuration command. Then assign an IP address to the port, enable routing, and assign routing protocol characteristics by using the **ip routing** and **router protocol** global configuration commands.

Entering a **no switchport** interface configuration command shuts the interface down and then re-enables it. This might generate messages on the device to which the interface is connected. Furthermore, when you use this command to put the interface into Layer 3 mode, you are deleting any Layer 2 characteristics configured on the interface.

The number of routed ports and SVIs that you can configure is not limited by software. However, the interrelationship between this number and the number of other features being configured might have an impact on CPU utilization because of hardware limitations.

Inter-VLAN Routing

This topic describes the operation and configuration of inter-VLAN routing.

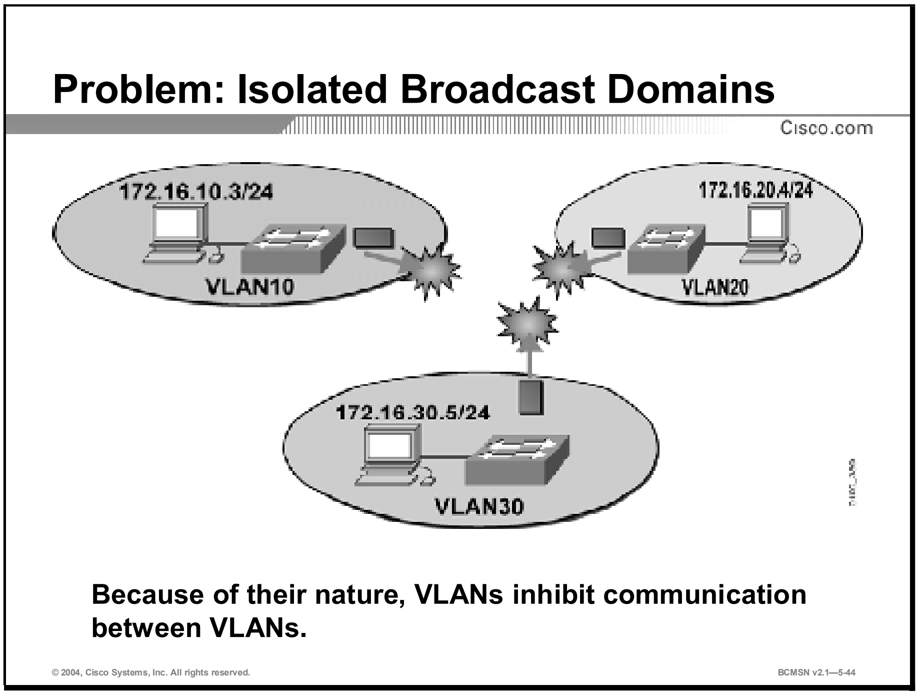

Problem: Isolated Broadcast Domains

Because of their nature, VLANs inhibit communication between VLANs.

BCMSN v2.1—5-44

VLANs are designed to control the size of the broadcast domain and to keep local traffic local. Because VLANs isolate traffic to a defined collision domain or subnet, network devices in different VLANs cannot communicate with one another without some intervening device to forward frames between subnets.

In Layer 2-switched networks, a routing device is used to provide communication between VLANs. The router provides VLAN access to shared resources and connects to other parts of the network, or it provides access to remote sites across wide-area links.

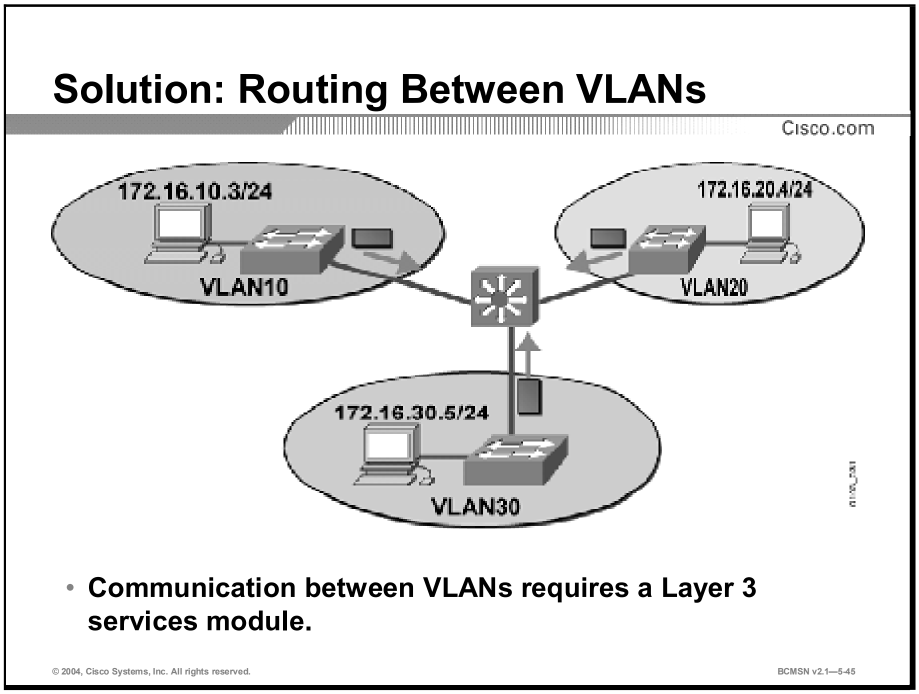

Solution: Routing Between VLANs

Cisco.com

- Communication between VLANs requires a Layer 3 services module.

BCMSN v2.1—5-45

Before you can configure routing between VLANs, you must first define the VLANs on the switches in your network. Issues related to network design and VLAN definition should be addressed during your network design. You need to consider these issues:

- Sharing resources between VLANs
- Load balancing
- Redundant links
- Addressing
- Segmenting networks with VLANs

So that each end device does not have to manage its own routing tables, most devices are configured with the IP address of a designated Layer 3 switch. This designated Layer 3 switch performs as the default router through which all nonlocal network packets are sent. The routing engine then forwards the packets toward the proper destination.

If the Layer 3 switch has no knowledge of that network segment, it discards the packet.

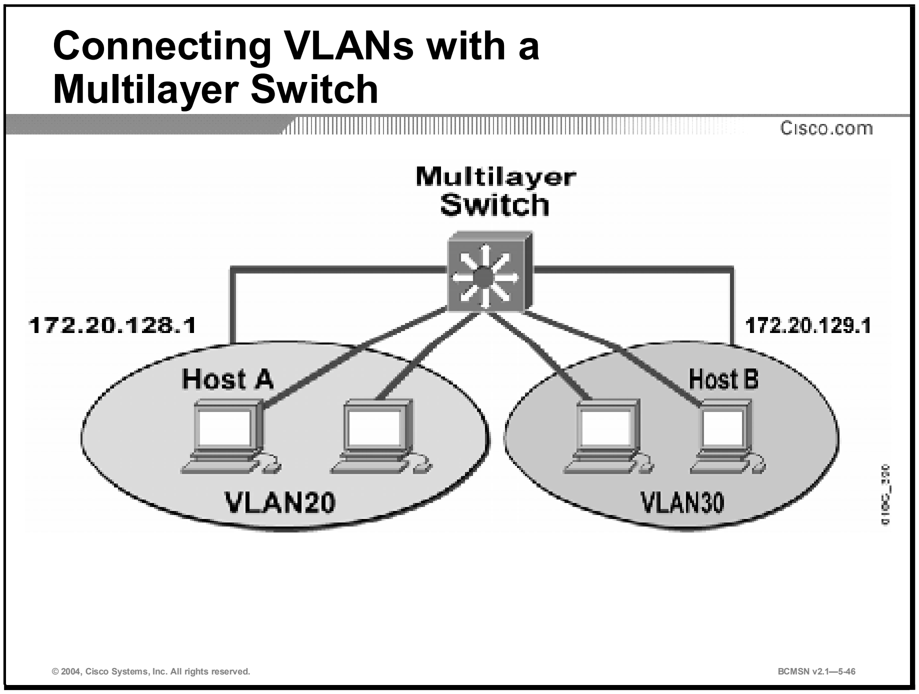

Connecting VLANs with a Multilayer Switch

Cisco.com

BCMSN v2.1—5-46

Devices within a single VLAN can communicate directly through any switch. Ports in different VLANs cannot exchange data without going through a routing device or interface. By using a multilayer switch, when you configure two VLANs (such as VLAN20 and VLAN30 in the figure), each with an SVI to which an IP address is assigned, the switch can send frames from host A to host B directly through the switch with no need for an external router.

Catalyst multilayer switches support two methods of forwarding traffic between interfaces: routing and fallback bridging. Whenever possible, to maintain high performance, forwarding is done by switch hardware. IP version 4 packets with Ethernet II encapsulation can be routed in hardware. All other types of IP traffic are routed in software using fast switching.

The routing function can be enabled on all SVIs and routed ports. Catalyst switches route only IP traffic. When IP routing protocol parameters and address configurations are added to an SVI or routed port, any IP traffic received from these ports is routed.

Fallback bridging forwards traffic not routed by the switch with the enhanced multilayer software image, which does not route traffic belonging to a nonroutable protocol. Fallback bridging connects multiple VLANs into one bridge domain by bridging between two or more SVIs or routed ports. When configuring fallback bridging, you assign SVIs or routed ports to bridge groups, with each SVI or routed port assigned to only one bridge group. All interfaces in the same group belong to the same bridge domain.

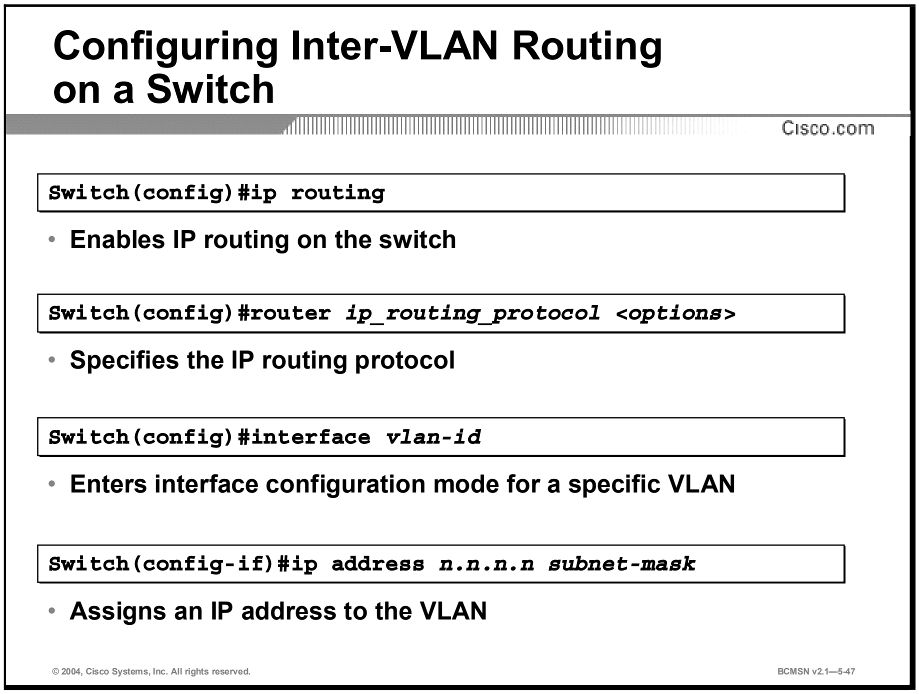

Configuring Inter-VLAN Routing on a Switch

Cisco.com

```
Switch(config)#ip routing
```

- Enables IP routing on the switch

```
Switch(config)#router ip_routing_protocol <options>
```

- Specifies the IP routing protocol

```
Switch(config)#interface vlan-id
```

- Enters interface configuration mode for a specific VLAN

```
Switch(config-if)#ip address n.n.n.n subnet-mask
```

- Assigns an IP address to the VLAN

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—5-47

To configure inter-VLAN routing for a routed interface on a Catalyst switch (such as the Catalyst 6000 series), perform these steps:

Step	Description	Notes and Comments
1.	(Optional) Enable IP routing on the router. `Switch(config)#ip routing`	This step is necessary if you have multiple routers in the network.
2.	(Optional) Specify an IP routing protocol. `Switch(config)#router ip_routing_protocol <options>`	This step is necessary if you enabled IP routing in Step 1. The routing protocol you specify may require additional options. This step might include other commands, such as using the network router configuration command to specify the networks to route. Refer to the documentation for your router platform for detailed information on configuring routing protocols.
3.	Specify a VLAN interface. `Switch(config)#interface vlan-id`	
4.	Assign an IP address to the VLAN. `Switch(config-if)#ip address n.n.n.n subnet-mask`	
5.	Exit configuration mode. `Switch(config-if)#Ctrl-Z`	

Example: Inter-VLAN Routing

The example shows how to enable IP routing on a Catalyst 6500, create a VLAN interface, and assign an IP address to the interface:

```
Switch#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)#ip routing
Switch(config)#router rip
Switch(config-router)#network 10.0.0.0
Switch(config)#interface vlan 100
Switch(config-if)#ip address 10.1.1.1 255.0.0.0
Switch(config-if)#^Z
Switch#
```

Router on a Stick

An alternative method of implementing inter-VLAN routing uses the router on a stick. A router is used to provide routing between VLANs. This topic describes the operation and configuration of router on a stick.

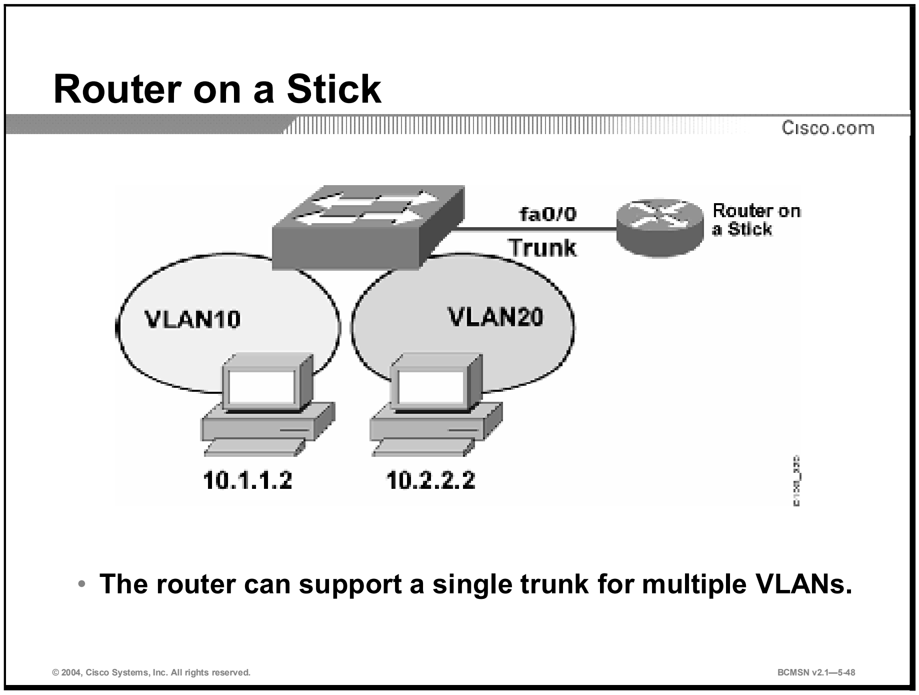

Router on a Stick

Cisco.com

- **The router can support a single trunk for multiple VLANs.**

BCMSN v2.1—5-48

To maintain integrity between VLAN traffic, the router on a stick is required to identify each VLAN frame. A single trunk can carry traffic from multiple VLANs. When implementing an 802.1Q trunk, it is important to ensure that the native VLAN assigned on each side of the link matches.

In the figure, the clients on VLAN10 and VLAN20 need to establish sessions with a server that is attached to a port designated to be in a separate VLAN. Because the file server resides in a different VLAN than any of the requestors, you need to configure inter-VLAN routing. The router would perform this function in this way:

1. The router accepts the packets from each VLAN. This is because the route processor is configured to route VLAN10 and VLAN20 traffic.

2. The router then classifies the packet based on the destination network address. The router applies the appropriate VLAN identification to the packet.

3. The router then routes the packets to the appropriate interface.

In the figure, the router can receive packets on one VLAN and forward them to another VLAN. To perform inter-VLAN routing functions, the router must know how to reach all VLANs being interconnected. The router must have a separate logical connection for each VLAN. You must enable ISL or 802.1Q trunking on a single physical connection. The router already knows about directly connected networks. The router must learn routes to networks not connected directly to it.

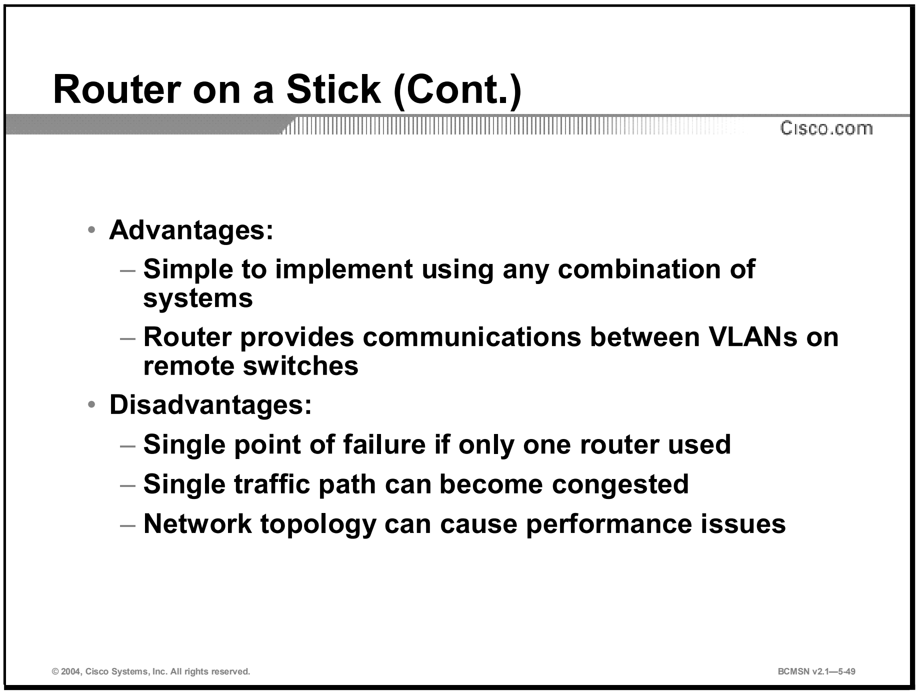

Router on a Stick (Cont.)

Cisco.com

- Advantages:
  - Simple to implement using any combination of systems
  - Router provides communications between VLANs on remote switches
- Disadvantages:
  - Single point of failure if only one router used
  - Single traffic path can become congested
  - Network topology can cause performance issues

BCMSN v2.1—5-49

The figure describes the advantages and disadvantages of using router on a stick for inter-VLAN routing.

Router on a Stick Between VLANs with ISL Trunks

Cisco.com

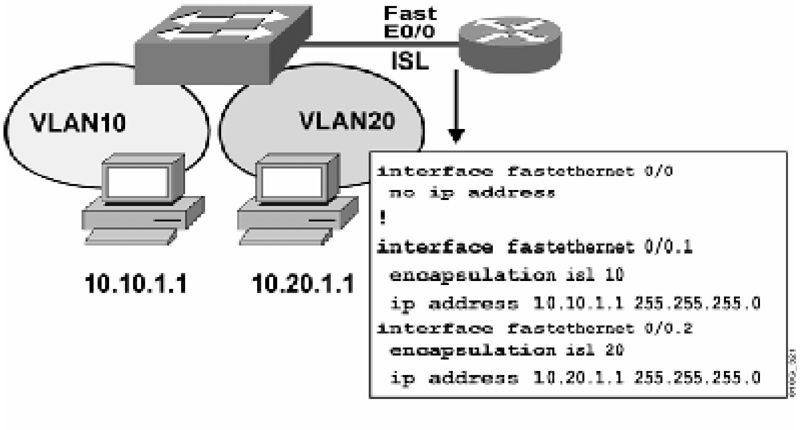

```
interface fastethernet 0/0
 no ip address
!
interface fastethernet 0/0.1
 encapsulation isl 10
 ip address 10.10.1.1 255.255.255.0
interface fastethernet 0/0.2
 encapsulation isl 20
 ip address 10.20.1.1 255.255.255.0
```

BCMSN v2.1—5-50

Use the **encapsulation isl** *vlan_id* subinterface configuration command to enable ISL trunking on a router subinterface (where *vlan identifier* is the VLAN number).

To configure the "router on a stick" for inter-VLAN routing, complete these tasks:

1. Enable ISL trunking on the switch port connecting to the router.

2. Enable ISL encapsulation on the router Fast Ethernet subinterface.

3. Assign a network layer address to each subinterface.

Note	In this example, the VLANs are directly connected. Routing between networks not directly connected requires that the router learn the routes either statically or dynamically (such as via a routing protocol).

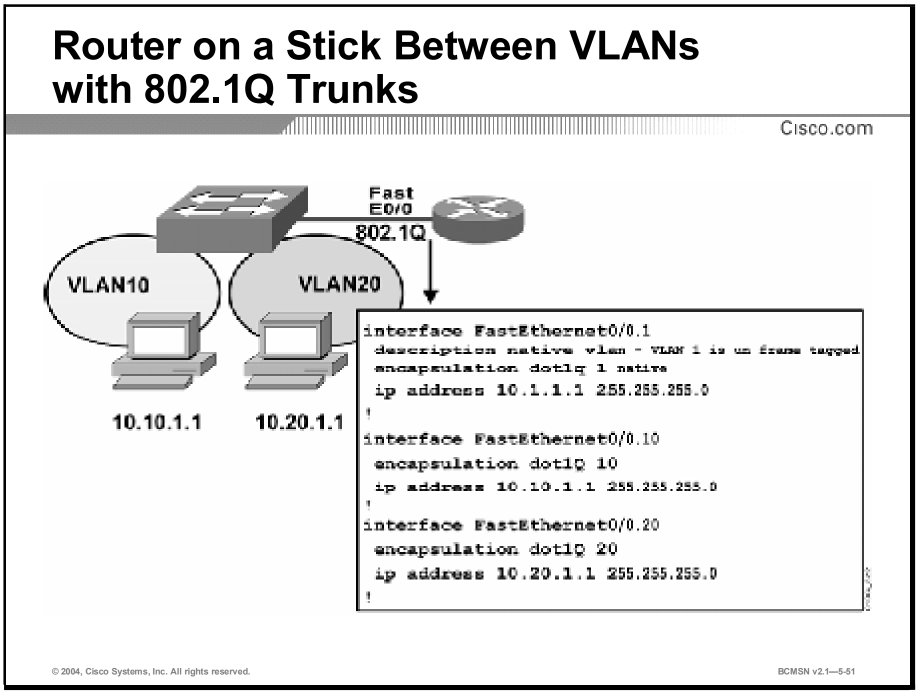

Router on a Stick Between VLANs with 802.1Q Trunks

Cisco.com

```
interface FastEthernet0/0.1
 description native vlan - VLAN 1 is un frame tagged
 encapsulation dot1q 1 native
 ip address 10.1.1.1 255.255.255.0
!
interface FastEthernet0/0.10
 encapsulation dot1Q 10
 ip address 10.10.1.1 255.255.255.0
!
interface FastEthernet0/0.20
 encapsulation dot1Q 20
 ip address 10.20.1.1 255.255.255.0
!
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—5-51

Use the **encapsulation dot1q** subinterface configuration command to enable 802.1Q encapsulation trunking on a router subinterface.

802.1Q is slightly different from ISL. The native VLAN frames in 802.1Q do not carry a tag. Therefore, the major interface of a trunk has an address. Any other configuration information for the native VLAN subinterfaces is configured with the dot1Q encapsulation**,** IP address, and so on. The subinterface number need not equal the dot1Q VLAN number. However, management is easier when the two numbers are the same.

The Catalyst 2950 switches support only 802.1Q encapsulation, which is configured automatically when trunking is enabled on the interface by using the **switchport mode trunk** command.

Verifying the Inter-VLAN Routing Configuration

This topic discusses how to verify the inter-VLAN routing configuration.

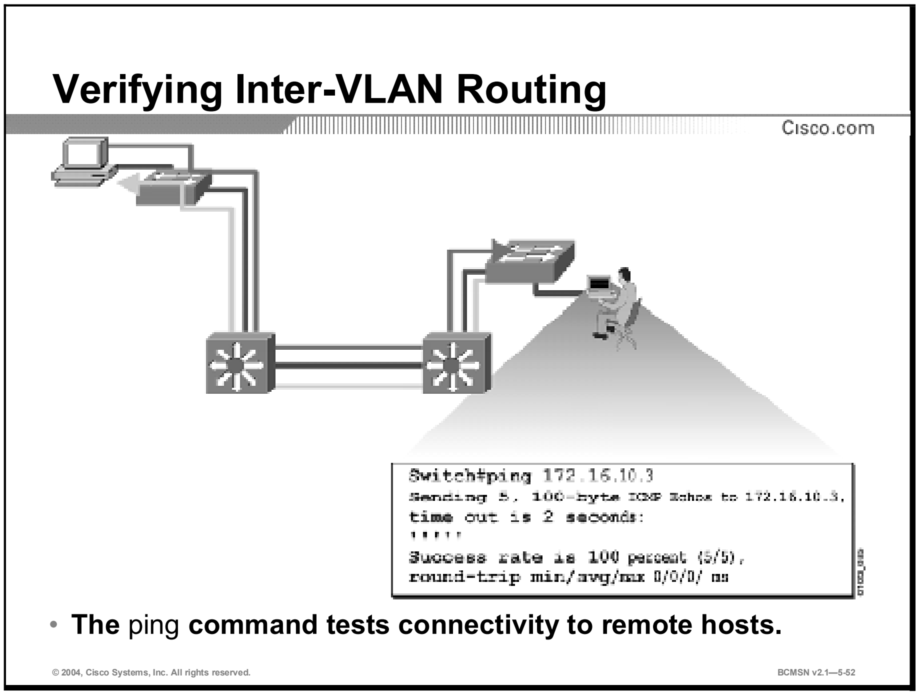

Verifying Inter-VLAN Routing

```
Switch#ping 172.16.10.3
Sending 5, 100-byte ICMP Echos to 172.16.10.3,
time out is 2 seconds:
!!!!!
Success rate is 100 percent (5/5),
round-trip min/avg/max 0/0/0/ ms
```

- **The** ping **command tests connectivity to remote hosts.**

After the router is properly configured and connected to the network, the router can communicate with other nodes on the network.

To test connectivity to remote hosts, enter the **ping** command in privileged mode:

```
Switch#ping destination-ip-address
```

The **ping** command will return one of these responses:

- **Success rate is 100 percent or *ip-address* is alive:** This response occurs in one to ten milliseconds, depending on network traffic and the number of Internet Control Message Protocol (ICMP) packets sent.
- **Destination does not respond:** No answer message is returned if the host does not respond.
- **Unknown host:** This response occurs if the targeted host cannot be resolved.
- **Destination unreachable:** This response occurs if the default gateway cannot reach the specified network or is being blocked.
- **Network or host unreachable:** This response occurs if the TTL times out. The TTL default is 2 seconds.

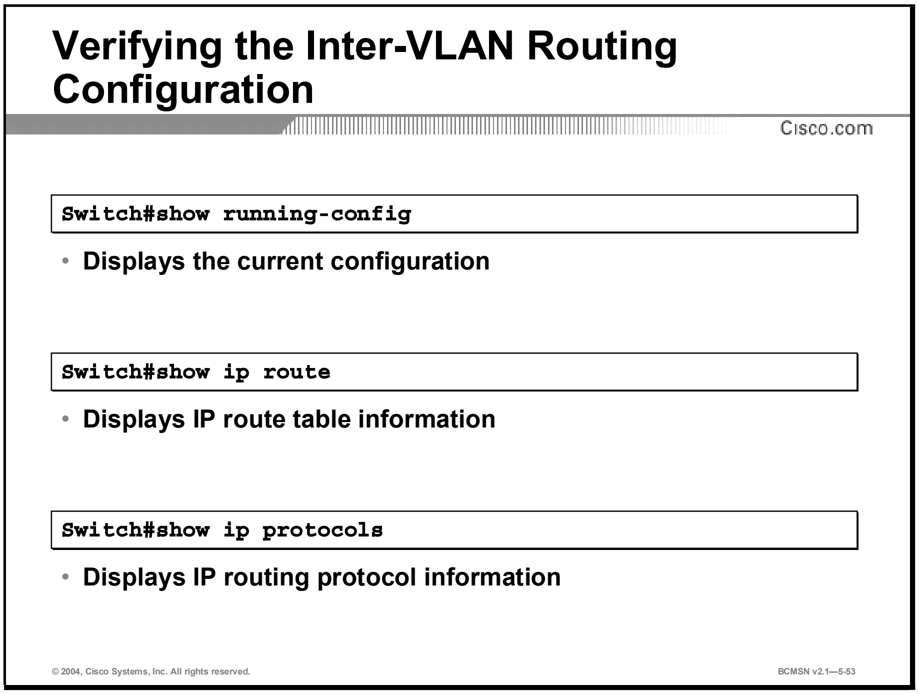

Verifying the Inter-VLAN Routing Configuration

Cisco.com

```
Switch#show running-config
```

- Displays the current configuration

```
Switch#show ip route
```

- Displays IP route table information

```
Switch#show ip protocols
```

- Displays IP routing protocol information

BCMSN v2.1—5-53

Use the **show** commands to display the current (running) configuration, IP routing information, and IP protocol information.

Example: Displaying Inter-VLAN Configuration Information

This example displays inter-VLAN configuration information:

```
Switch#show running-config
(text deleted)
!
interface VLAN1
 ip address 172.16.1.111 255.255.255.0
 no ip route-cache
!
interface VLAN11
 ip address 172.16.11.111 255.255.255.0
 no ip route-cache
!
[output omitted]

Switch#show ip route
Codes: C -connected,S -static,I -IGRP,R -RIP,M -mobile,B -BGP
D -EIGRP,EX_-EIGRP external,O -OSPF,IA -OSPF inter area
N1 -OSPF NSSA external type 1,N2 -OSPF NSSA external type 2
E1 -OSPF external type 1, E2 -OSPF external type 2, E -EGP
```

```
I -IS-IS,L1 -IS-IS level-1,L2 -IS-IS level-2,ia -IS-IS inter
area

* -candidate default,U -pre-user static route,o -ODR

P -periodic downloaded static route

Gateway of last resort is not set

172.16.0.0/24 is subnetted, 5 subnets

C     172.16.11.0 is directly connected, Vlan11

C     172.16.12.0 is directly connected, Vlan12

C     172.16.13.0 is directly connected, Vlan13

C     172.16.14.0 is directly connected, Vlan14

C     172.16.1.0 is directly connected, Vlan1
```

Example: Displaying IP Routing Protocol Information

This example displays IP routing protocol information:

```
Switch#show ip protocols
Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 1
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    172.16.0.0
  Passive Interface(s):
    Vlan1
    Vlan11
    Vlan12
    Vlan13
    Vlan14
  Routing Information Sources:
    Gateway         Distance      Last Update
    172.16.117.202        90      20:25:10
    172.16.113.201        90      20:25:10
    Gateway         Distance      Last Update
    172.16.115.202        90      20:25:12
    172.16.111.201        90      20:25:12
  Distance: internal 90 external 170
Switch#
```

Summary

This topic summarizes the key points discussed in this lesson.

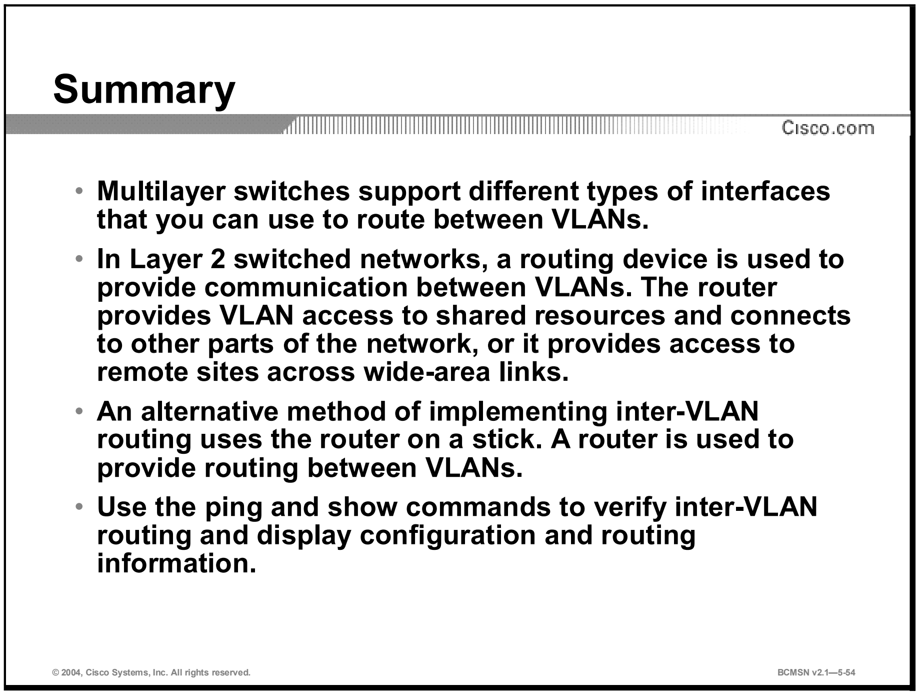

Summary

Cisco.com

- Multilayer switches support different types of interfaces that you can use to route between VLANs.
- In Layer 2 switched networks, a routing device is used to provide communication between VLANs. The router provides VLAN access to shared resources and connects to other parts of the network, or it provides access to remote sites across wide-area links.
- An alternative method of implementing inter-VLAN routing uses the router on a stick. A router is used to provide routing between VLANs.
- Use the ping and show commands to verify inter-VLAN routing and display configuration and routing information.

BCMSN v2.1—5-54

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Next Steps

For the associated lab exercises, refer to the following section of the course Lab Guide:

- Lab Exercise 5-1: Implementing Layer 3 Services
- Lab Exercise 5-2: Implementing Multilayer Switching in the Network

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which type of port or interface on a multilayer switch represents a VLAN of switch ports as one interface to the routing or bridging function in the system?

A) trunk port

B) routed port

C) switched port

D) switch virtual interface

Q2) What is one advantage of using a trunk link for inter-VLAN routing?

A) scalability

B) availability

C) performance

D) bandwidth utilization

Q3) Which Cisco IOS command enables IP routing on a routed interface on a Catalyst switch?

A) **`ip routing`**

B) **`interface`** ***`vlan-id`***

C) **`ip address`** *`n.n.n.n mask`*

D) **`router`** *`ip_routing_protocol`*

Q4) Which command displays information about the specific routing protocol(s) in use?

A) **`show eigrp`**

B) **`show ip route`**

C) **`show protocols`**

D) **`show ip protocols`**

Quiz Answer Key

Q1) D

Relates to: Connecting VLANs with Multilayer Switches

Q2) A

Relates to: Inter-VLAN Routing

Q3) A

Relates to: Router on a Stick

Q4) D

Relates to: Verifying the Inter-VLAN Routing Configuration

Lesson Assessments

Overview

Use the lesson assessments here to test what you learned in this module. The correct answers and solutions are found in the Lesson Assessment Answer Key.

Outline

This section includes these assessments:

- Quiz 5-1: Examining Multilayer Switching
- Quiz 5-2: Configuring Multilayer Switching
- Quiz 5-3: Routing Between VLANs

Quiz 5-1: Examining Multilayer Switching

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Identify the operation of the key components required to implement Layer 3 switching
- Compare Layer 2 and multilayer forwarding, and explain the packet flow with each type of forwarding
- Explain the multilayer switching table architectures
- Describe centralized and distributed forwarding, demand-based switching, and topology-based switching

Quiz

Answer these questions:

Q1) What does a Layer 3 switch examine to make switching decisions?

A) the frame source address

B) the packet source address

C) the frame destination address

D) the packet destination address

Q2) How are entries in a forwarding table sorted?

A) by arrival time

B) numerically

C) by frequency of use

D) by frequency of advertisement

Q3) Which ACL definition would create a TCAM mask for the most significant 24 bits of the source address?

A) access-list 101 permit ip 10.1.123.0 0.0.0.255 any

B) access-list 101 permit ip 10.1.0.0 0.0.255.255 any

C) access-list 101 permit ip 10.1.0.0 0.255.255.255 any

D) access-list 101 permit ip 10.1.123.0 0.0.255.255 any

Q4) With which architecture is the first packet in a stream switched in software, with the remaining packets switched in hardware?

A) distributed forwarding

B) centralized forwarding

C) demand-based switching

D) topology-based switching

Q5) Which operation is performed when a multilayer switch rewrites an IP unicast packet on the outbound interface?

A) The source MAC address changes from the router MAC to the next-hop MAC address.

B) The destination MAC address changes from the router MAC to the next-hop MAC address.

C) The source MAC address changes from the next-hop MAC address to the router MAC address.

D) The destination MAC address changes from the sender MAC address to the router MAC address.

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Quiz 5-2: Configuring Multilayer Switching

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Explain the features and operation of CEF-based MLS
- Configure and verify CEF
- Troubleshoot CEF

Quiz

Answer these questions:

Q1) From which table is the FIB table derived in CEF-based multilayer switching?

A) ARP table

B) adjacency table

C) IP routing table

D) MAC address table

Q2) On which switch model does the **no ip cef** command disable CEF?

A) Catalyst 3550

B) Catalyst 4000

C) Catalyst 6000

D) Catalyst 6500

Q3) What is displayed in the output of the **debug ip cef table** command?

A) general CEF events

B) events related to the FIB table

C) updates related to IP prefix information

D) control messages related to FIB table prefixes

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 67 percent or better.

Quiz 5-3: Routing Between VLANs

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Describe and configure the different interface types on a multilayer switch for routing between VLANs
- Explain the operation and configuration of inter-VLAN routing
- Explain the operation and configuration of router on a stick
- Verify the inter-VLAN routing configuration

Quiz

Answer these questions:

Q1) On a multilayer switch, which type of port carries the traffic of, and belongs to, only one VLAN?

A) trunk port

B) access port

C) switch port

D) routed port

Q2) Why is inter-VLAN routing necessary?

A) to allow for scalability

B) to provide high availability

C) to enable hosts within a VLAN to communicate with each other

D) to enable hosts on one VLAN to communicate with hosts on a different VLAN

Q3) Why is router on a stick a possible design choice?

A) because Layer 2 LAN switches can also operate at Layer 3

B) because routers cannot route packets between multiple VLANs

C) because Layer 2 LAN switches can only support a single VLAN

D) because Layer 2 LA N switches cannot switch frames between multiple VLANs

Q4) What information can you look for in the output of the **show run** command that helps confirm your inter-VLAN routing configuration?

A) port interface IP addresses

B) VLAN interface IP addresses

C) port interface routing protocol

D) VLAN interface encapsulation

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 75 percent or better.

Lesson Assessment Answer Key

Quiz 5-1: Examining Multilayer Switching

Q1) D

Q2) A

Q3) A

Q4) C

Q5) B

Quiz 5-2: Configuring Multilayer Switching

Q1) C

Q2) B

Q3) B

Quiz 5-3: Routing Between VLANs

Q1) B

Q2) D

Q3) A

Q4) B

Module 6

Improving Availability on Multilayer Switched Networks

Overview

High availability optimizes design and tools to ensure end-to-end availability for services and clients. Tools include reliable, fault-tolerant network devices to automatically identify and overcome failures, and resilient network technologies, such as Hot Standby Router Protocol (HSRP), to bring resilience to the critical junction between hosts and backbone links.

Upon completing this module, you will be able to:

- Identify considerations for installing and troubleshooting redundant modules to provide high availability in multilayer switched networks
- Configure and troubleshoot router redundancy to provide high availability on multilayer switched networks

Outline

The module contains these components:

- Implementing Module Redundancy in a Multilayer Switched Network
- Implementing Router Redundancy in a Multilayer Switched Network
- Lesson Assessments

Implementing Module Redundancy in a Multilayer Switched Network

Overview

As enterprises rely more heavily on their IP network for core business practices, a high degree of network availability becomes critical. System downtime translates into significant productivity and revenue losses.

Relevance

Maximizing network uptime requires the use of operational best practices and redundant network designs in conjunction with high-availability technologies within network elements. Several of these technologies are embedded in Cisco IOS software.

Objectives

Upon completing this lesson, you will be able to:

- Explain the types of redundancy in a multilayer switched network, including hardware and software redundancy
- Implement redundant supervisor modules in Catalyst switches
- Implement redundant supervisor uplink modules in Catalyst switches
- Implement redundant power supplies

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Introducing Redundancy
- Implementing Redundant Supervisor Engines in Catalyst Switches
- Implementing Redundant Supervisor Uplink Modules in Catalyst Switches
- Implementing Redundant Power Supplies
- Summary
- Quiz

Introducing Redundancy

Providing hardware redundancy in a switched network can be accomplished by implementing redundant modules within devices or by deploying redundant devices. This topic introduces hardware redundancy.

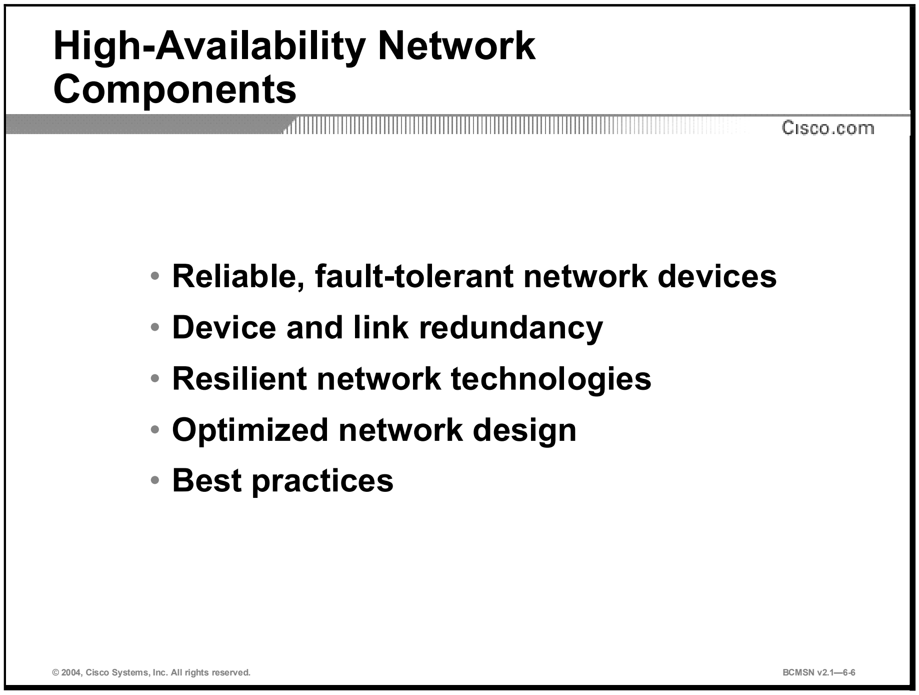

High-Availability Network Components

- Reliable, fault-tolerant network devices
- Device and link redundancy
- Resilient network technologies
- Optimized network design
- Best practices

BCMSN v2.1—6-6

To achieve high network availability, these network components are required:

- **Reliable, fault-tolerant network devices:** Hardware and software reliability to automatically identify and overcome failures.
- **Device and link redundancy:** Entire devices may be redundant or modules within devices can be redundant. Links may also be redundant.
- **Resilient network technologies:** Intelligence that ensures fast recovery around any device or link failure.
- **Optimized network design:** Well-defined network topologies and configurations designed to ensure there is no single point of failure.
- **Best practices:** Documented procedures for deploying and maintaining a robust e-commerce network infrastructure.

High availability implies that a device or network is ready for use as close to 100 percent of the time as possible. Fault tolerance indicates the ability of a device or network to recover from the failure of a component or device. Achieving high availability relies on eliminating any single point of failure and on distributing intelligence throughout the architecture. You can increase availability by adding redundant components, including redundant network devices and connections to redundant Internet services. With the proper design, no single point of failure will have an impact on the availability of the overall system.

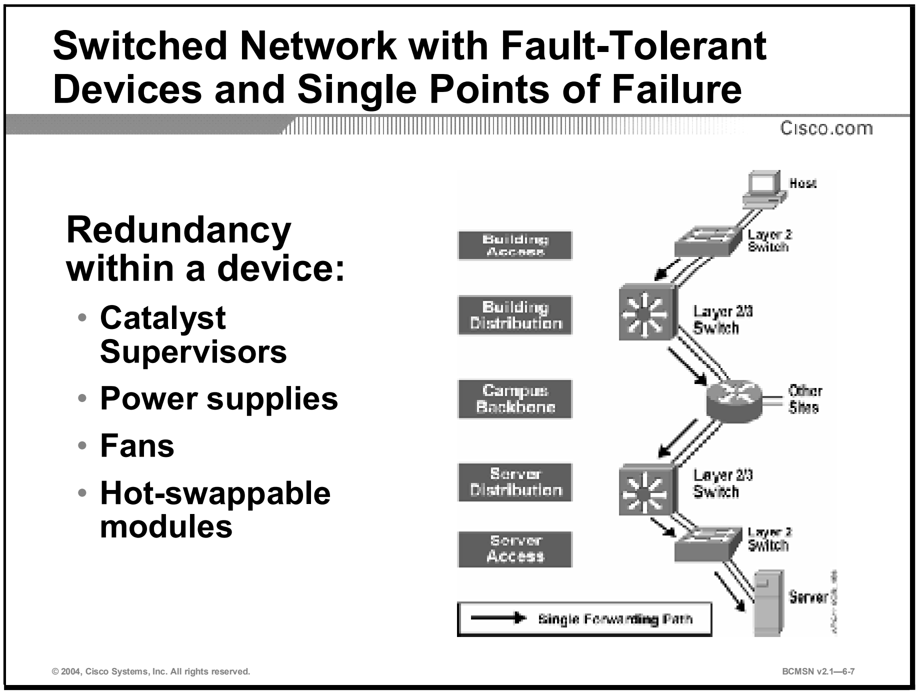

Switched Network with Fault-Tolerant Devices and Single Points of Failure

Cisco.com

Redundancy within a device:

- **Catalyst Supervisors**
- **Power supplies**
- **Fans**
- **Hot-swappable modules**

BCMSN v2.1—6-7

One approach to building highly available networks is to use extremely fault-tolerant network devices throughout the network. To achieve high availability end –to end, the fault tolerance of each device is optimized. This is achieved by providing redundant backup within the device for each of its key components. Fault tolerance offers these benefits:

- Minimizes time periods during which the system is nonresponsive to requests (for example, while the system is being reconfigured because of a component failure or recovery)
- Eliminates all single points of failure that would cause the system to stop
- Provides disaster protection by allowing the major system components to be separated geographically

Trying to achieve high network availability solely through device-level fault tolerance has a number of drawbacks.

- Massive redundancy within each device adds significantly to its cost. Massive redundancy also reduces physical capacity of each device by consuming slots that could otherwise house network interfaces or provide useful network services.
- Redundant subsystems within devices are often maintained in a hot-standby mode. In hot standby mode, such redundant subsystems cannot contribute additional performance because they are only fully activated when the primary component fails.
- Focusing on device-level hardware reliability may result in a number of other failure mechanisms being overlooked. Network elements are not standalone devices; they are components of a network system whose internal operations and system-level interactions are governed by software and configuration parameters.

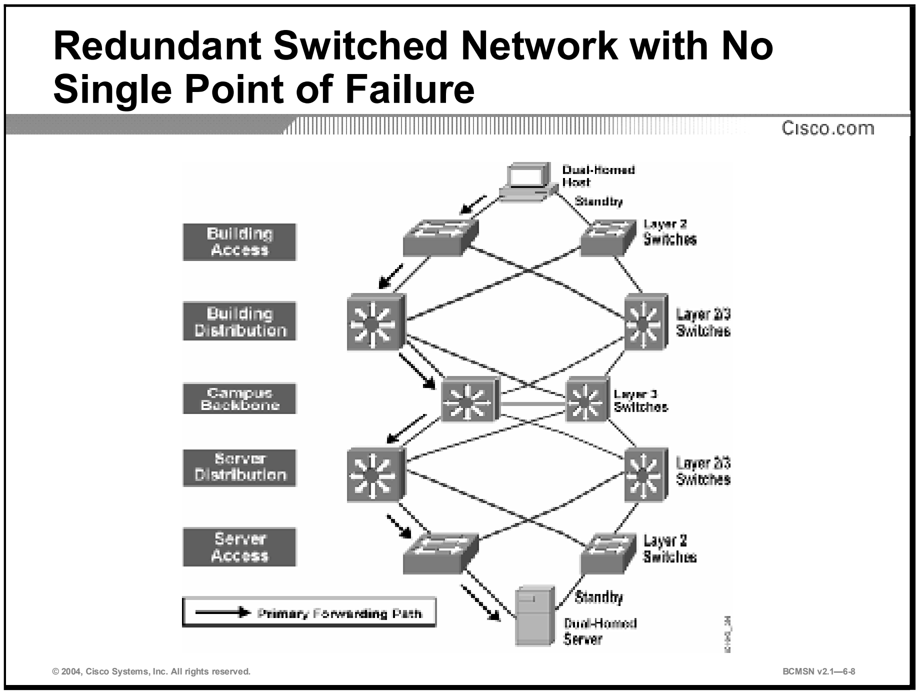

Redundant Switched Network with No Single Point of Failure

Cisco.com

BCMSN v2.1—6-8

A complementary way to build highly available networks is to provide reliability through redundancy in the network topology rather than primarily within the network devices themselves. In the campus network design shown in the figure, there is a backup for every link and for every network device in the path between the client and server. This approach to network reliability offers these advantages:

- The network elements providing redundancy need not be colocated with the primary network elements. This reduces the probability that problems with the physical environment will interrupt service.
- Software discrepancies, software upgrades, or configuration errors and changes can be dealt with separately in the primary and secondary forwarding paths without completely interrupting service. Therefore, network-level redundancy can also reduce the impact of nonhardware failure mechanisms.
- With the redundancy provided by the network, each network device no longer needs to be configured for optimal standalone fault tolerance. Device-level fault tolerance can be concentrated in the Building Core and Building Distribution layers of the network where a hardware failure would affect a larger number of users. By partially relaxing the requirement for device-level fault tolerance, the cost per network device is reduced. To some degree, this offsets the requirement for more devices.
- With appropriate resiliency features combined with careful design and configuration, the traffic load between the respective layers of the network topology (that is, the Building Access submodule to the Building Distribution submodule) can be shared between the primary and secondary forwarding paths. Therefore, network-level redundancy can also provide increased aggregate performance and capacity.
- Redundant networks can be configured to automatically fail over from primary to secondary facilities without operator intervention. The duration of service interruption is equal to the time it takes for failover to occur. Failover times as low as a few seconds are possible.

Fast EtherChannel (FEC) and Gigabit EtherChannel (GEC) are link-aggregation (and access-link) technologies based on grouping together multiple full-duplex Fast Ethernet or Gigabit Ethernet ports to provide fault-tolerant, high-speed links between switches, routers, and servers. EtherChannel uses a peer-to-peer control protocol that provides autoconfiguration and minimal convergence times for parallel links.

The drawbacks of link redundancy include the following:

- Increased media costs
- More difficulty managing and troubleshooting

Implementing Redundant Supervisor Engines in Catalyst Switches

Route processor redundancy enables a second Supervisor Engine to provide failover capabilities within a Catalyst switch. This topic explains how to implement redundant Supervisor Engines.

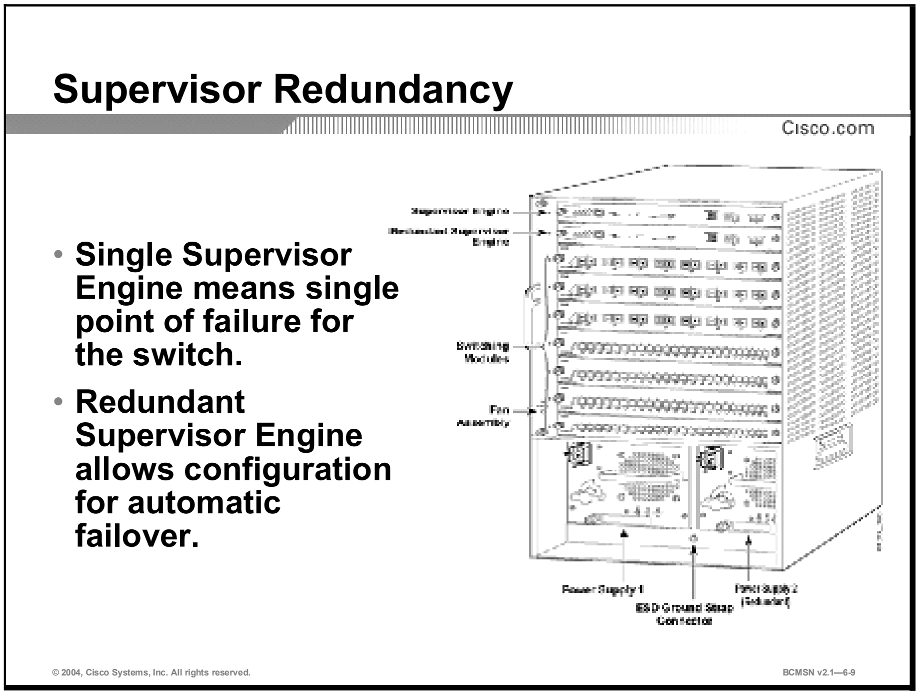

Supervisor Redundancy

Cisco.com

- Single Supervisor Engine means single point of failure for the switch.
- Redundant Supervisor Engine allows configuration for automatic failover.

Supervisor Engine
Redundant Supervisor Engine
Switching Modules
Fan Assembly
Power Supply 1
ESD Ground Strap Connector
Power Supply 2 (Redundant)

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-9

You can ensure availability by configuring dual Supervisor Engines within a Catalyst 6000 series switch. Dual Supervisor engines within a single device provide redundancy without having to add a separate switch. This solution can be a cost-effective alternative to deploying multiple devices. When multiple redundant switches are used, configuring redundant Supervisor Engines adds an extra level of availability assurance.

Cisco IOS Release 12.1(13)E and later support Supervisor Engine redundancy with Route Processor Redundancy (RPR) and Route Processor Redundancy Plus (RPR+).

Note The Cisco Catalyst switch features described in this topic apply to the Catalyst 6500 series. For the Catalyst 4507R, for example, the RPR has a subminute failover time.

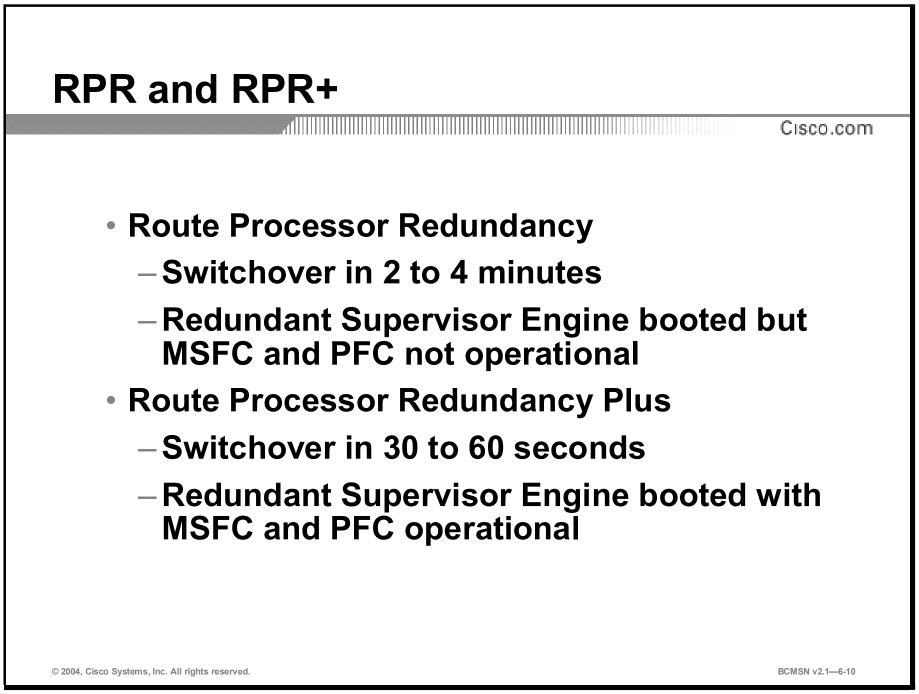

RPR and RPR+

Cisco.com

- Route Processor Redundancy
  - Switchover in 2 to 4 minutes
  - Redundant Supervisor Engine booted but MSFC and PFC not operational
- Route Processor Redundancy Plus
  - Switchover in 30 to 60 seconds
  - Redundant Supervisor Engine booted with MSFC and PFC operational

BCMSN v2.1—6-10

Catalyst 6500 series switches support fault resistance by allowing a redundant Supervisor Engine to take over if the primary Supervisor Engine. RPR supports a switchover time of two to four minutes and RPR+ supports a switchover time of 30 to 60 seconds.

When RPR+ mode is used, the redundant Supervisor Engine is fully initialized and configured, which shortens the switchover time. The active Supervisor Engine checks the image version of the redundant Supervisor Engine when the redundant Supervisor Engine comes online. If the image on the redundant Supervisor Engine does not match the image on the active Supervisor Engine, RPR redundancy mode is used.

RPR supports these features:

- Auto-startup and bootvar synchronization between active and redundant Supervisor Engines.
- Hardware signals that detect and decide the active or redundant status of Supervisor Engines.
- Clock synchronization every 60 seconds from the active to the redundant Supervisor Engine.
- A redundant Supervisor Engine that is booted without all subsystems needing to be up: If the active Supervisor Engine fails, the redundant Supervisor Engine becomes fully operational.
- An operational Supervisor Engine in place of the failed unit becomes the redundant Supervisor Engine.
- Support for fast software upgrade.

Note The two Gigabit Ethernet interfaces on the redundant Supervisor Engine are always active.

When the switch is powered on, RPR runs between the two Supervisor Engines. The Supervisor Engine that boots first, either in slot 1 or 2 becomes the RPR active Supervisor Engine. The Multilayer Switch Feature Card (MSFC) or Multilayer Switch Feature Card 2 (MSFC2) and the Policy Feature Card (PFC) or Policy Feature Card 2 (PFC2) become fully operational. The MSFC and PFC on the redundant Supervisor Engine come out of reset but are not operational.

These events cause an RPR switchover:

- Clock synchronization failure between Supervisor Engines
- MSFC or PFC failure on the active Supervisor Engine
- A manual switchover

In a switchover, the redundant Supervisor Engine becomes fully operational and these events occur:

- All switching modules recycle power.
- The remaining subsystems on the MSFC (including Layer 2 and Layer 3 protocols) are brought up.
- Access control lists (ACLs) are reprogrammed into Supervisor Engine hardware.

Note In a switchover, there is a disruption of traffic. This is because some address states are lost and then restored after they are dynamically redetermined.

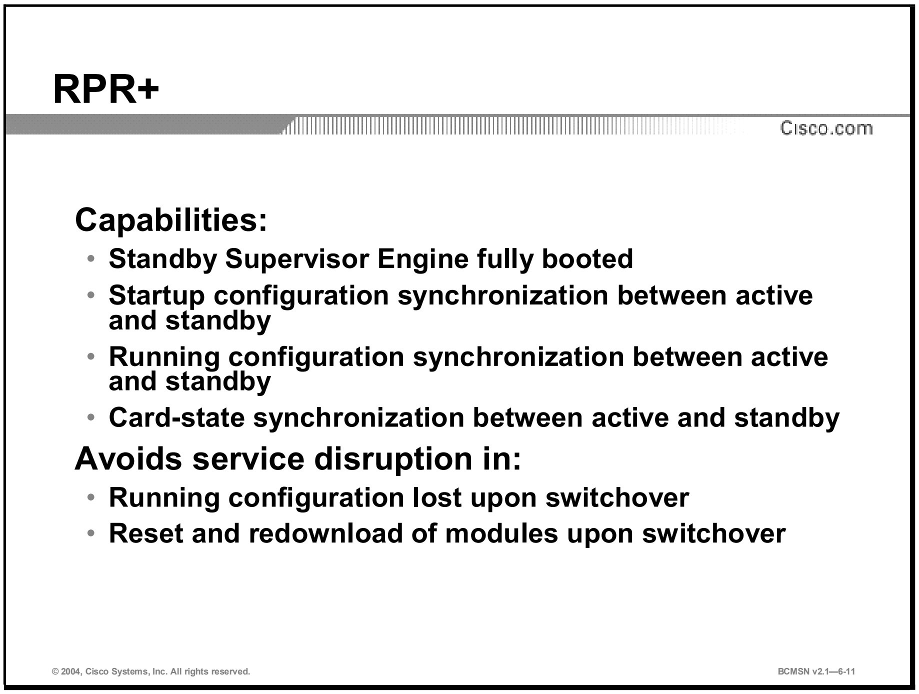

RPR+

Capabilities:

- **Standby Supervisor Engine fully booted**
- **Startup configuration synchronization between active and standby**
- **Running configuration synchronization between active and standby**
- **Card-state synchronization between active and standby**

Avoids service disruption in:

- **Running configuration lost upon switchover**
- **Reset and redownload of modules upon switchover**

BCMSN v2.1—6-11

With RPR+, the redundant Supervisor Engine is fully initialized and configured. This shortens the switchover time if the active Supervisor Engine fails or if a manual switchover is performed.

When the switch is powered on, RPR+ runs between the two Supervisor Engines. The Supervisor Engine that boots first, either in slot 1 or 2, becomes the active Supervisor Engine. The MSFC or MSFC2 and the PFC or PFC2 become fully operational. The MSFC and PFC on the redundant Supervisor Engine come out of reset but are not operational.

RPR+ enhances RPR by providing these additional benefits:

- Reduced switchover time. Depending on the configuration, the switchover time is in the range of 30 to 60 seconds.
- Installed modules are not reloaded. Because both the startup configuration and the running configuration are continually synchronized from the active to the redundant Supervisor Engine, installed modules are not reloaded during a switchover.
- Online insertion and removal (OIR) of the redundant Supervisor Engine. RPR+ allows OIR of the redundant Supervisor Engine for maintenance. When the redundant Supervisor Engine is inserted, the active Supervisor Engine detects its presence and begins to transition the redundant Supervisor Engine to a fully initialized state.
- Synchronization of OIR events.
- Manual user-initiated switchover using the **redundancy force-switchover** command.

These events cause an RPR+ switchover:

- Clock synchronization failure between Supervisor Engines
- MSFC or PFC failure on the active Supervisor Engine

During RPR mode operation, the startup configuration and register configuration are synchronized by default between the two Supervisor Engines. In a switchover, the new active Supervisor Engine uses the current configuration.

Note Unless autosynchronization has been disabled, the boot variables are synchronized by default.

When a redundant Supervisor Engine configuration is running in RPR+ mode, these operations trigger synchronization:

- When a redundant Supervisor Engine first comes online, the configuration information is synchronized in bulk from the active Supervisor Engine to the redundant Supervisor Engine. This synchronization overwrites any existing startup configuration file on the redundant Supervisor Engine.
- When configuration changes occur during normal operation, RPR+ performs an incremental synchronization from the active Supervisor Engine to the redundant Supervisor Engine. RPR+ synchronizes user-entered command-line interface (CLI) commands incrementally, line –by line, from the active Supervisor Engine to the redundant Supervisor Engine.

Note Even though the redundant Supervisor Engine is fully initialized, it interacts with the active Supervisor Engine only to receive incremental changes to the configuration files as they occur. You cannot enter CLI commands on the redundant Supervisor Engine. Synchronization of the startup configuration file is enabled by default in RPR+ mode.

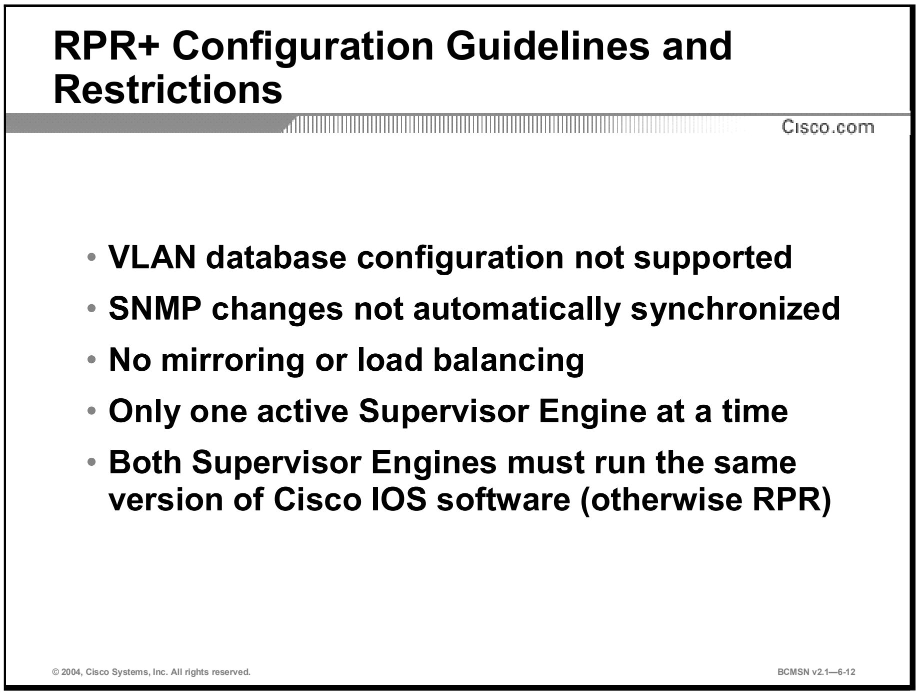

RPR+ Configuration Guidelines and Restrictions

- **VLAN database configuration not supported**
- **SNMP changes not automatically synchronized**
- **No mirroring or load balancing**
- **Only one active Supervisor Engine at a time**
- **Both Supervisor Engines must run the same version of Cisco IOS software (otherwise RPR)**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-12

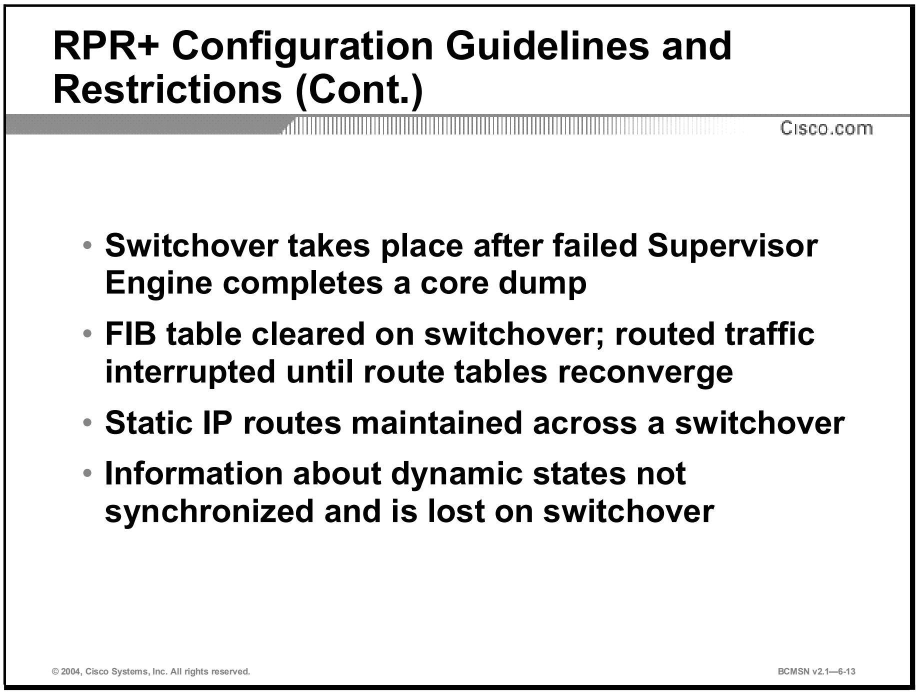

RPR+ Configuration Guidelines and Restrictions (Cont.)

- Switchover takes place after failed Supervisor Engine completes a core dump
- FIB table cleared on switchover; routed traffic interrupted until route tables reconverge
- Static IP routes maintained across a switchover
- Information about dynamic states not synchronized and is lost on switchover

BCMSN v2.1—6-13

These guidelines and restrictions apply to RPR+:

- RPR+ redundancy does not support configuration entered in VLAN database mode. Use global configuration mode with RPR+ redundancy.
- Configuration changes made through Simple Network Management Protocol (SNMP) are not synchronized to the redundant Supervisor Engine. Enter a **copy running-config startup-config** command to synchronize the configuration on the redundant Supervisor Engine.
- Supervisor Engine redundancy does not provide Supervisor Engine mirroring or load balancing. Only one Supervisor Engine is active. Network services are disrupted until the redundant Supervisor Engine takes over and the switch recovers.
- The two Gigabit Ethernet interfaces on the redundant Supervisor Engine are always active.
- With RPR+, both Supervisor Engines must run the same version of Cisco IOS software. If the Supervisor Engines are not running the same version of IOS software, the redundant Supervisor Engine comes online in RPR mode.
- RPR+ switchover takes place after the failed Supervisor Engine completes a core dump. A core dump can take up to 15 minutes. To get faster switchover time, disable core dump on the Supervisor Engines.
- The Forwarding Information Base (FIB) tables are cleared on a switchover. As a result, routed traffic is interrupted until route tables reconverge.
- Static IP routes are maintained across a switchover because they are configured from entries in the configuration file.

- Information about dynamic states maintained on the active Supervisor Engine is not synchronized to the redundant Supervisor Engine and is lost on switchover. Examples of dynamic state information that is lost at switchover include the following:
  - — Frame Relay switch virtual circuit (SVC) connections (Frame Relay-switched data-link connection identifier [DLCI] information is maintained across a switchover because that configuration is in the configuration file.)
  - — All terminated PPP sessions
  - — All ATM SVC information
  - — All terminated TCP and other connection-oriented Layer 3 and Layer 4 sessions
  - — Border Gateway Protocol (BGP) sessions
  - — All automatic protection switching (APS) state information

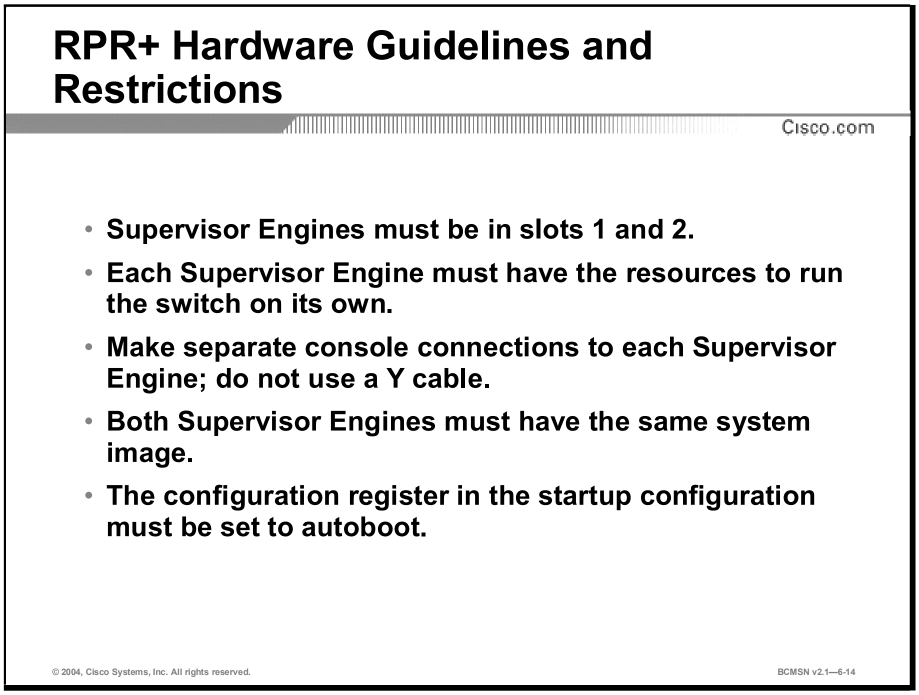

RPR+ Hardware Guidelines and Restrictions

- **Supervisor Engines must be in slots 1 and 2.**
- **Each Supervisor Engine must have the resources to run the switch on its own.**
- **Make separate console connections to each Supervisor Engine; do not use a Y cable.**
- **Both Supervisor Engines must have the same system image.**
- **The configuration register in the startup configuration must be set to autoboot.**

BCMSN v2.1—6-14

For redundant operation, these hardware guidelines and restrictions must be met:

- The active and redundant Supervisor Engines must be in slots 1 and 2.
- Each Supervisor Engine must have the resources to run the switch on its own. This means that all Supervisor Engine resources are duplicated. In other words, each Supervisor Engine has its own Flash memory device and console port connections.
- There are separate console connections to each Supervisor Engine. Do not connect a Y cable to the console ports.
- Both Supervisor Engines must have the same system image.

Note If the redundant Supervisor Engine is running Catalyst software, remove the active Supervisor Engine and boot the switch with only the redundant Supervisor Engine installed. Follow the procedures in the current release notes to convert the redundant Supervisor Engine from Catalyst software.

- The configuration register in the startup configuration must be set to autoboot.

Note There is no support for booting from the network.

If these requirements are met, the switch functions in RPR+ mode by default.

These configuration restrictions apply during the startup synchronization process:

- You cannot perform configuration changes during the startup (bulk) synchronization. If you attempt to make configuration changes during this process, this message is generated:

  ```
  Config mode locked out till standby initializes
  ```

- If configuration changes occur at the same time as a Supervisor Engine switchover, those configuration changes are lost.

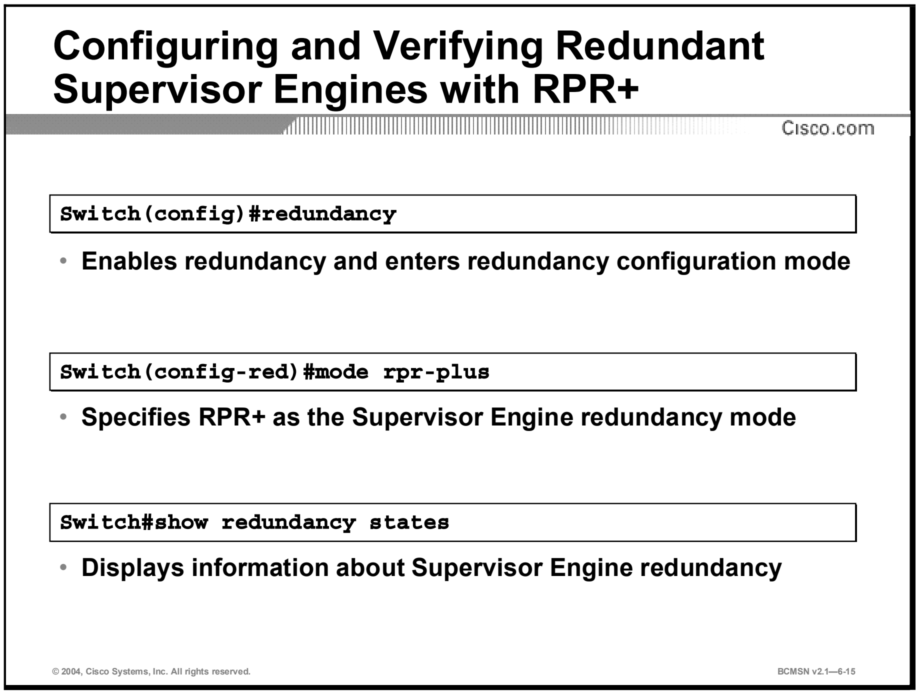

Configuring and Verifying Redundant Supervisor Engines with RPR+

Cisco.com

```
Switch(config)#redundancy
```

- Enables redundancy and enters redundancy configuration mode

```
Switch(config-red)#mode rpr-plus
```

- Specifies RPR+ as the Supervisor Engine redundancy mode

```
Switch#show redundancy states
```

- Displays information about Supervisor Engine redundancy

BCMSN v2.1—6-15

Example: Configuring RPR+

This example shows how to configure and verify RPR+:

```
Switch#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)#redundancy
Switch(config-red)#mode rpr-plus
Switch(config-red)#^Z
Switch#show redundancy states
       my state = 13 -ACTIVE
     peer state = 1  -DISABLED
           Mode = Simplex
           Unit = Primary
        Unit ID = 1

Redundancy Mode (Operational) = Route Processor Redundancy
Plus
Redundancy Mode (Configured)  = Route Processor Redundancy
Plus
     Split Mode = Disabled
   Manual Swact = Disabled  Reason: Simplex mode
 Communications = Down      Reason: Simplex mode

   client count = 11
 client_notification_TMR = 30000 milliseconds
          keep_alive TMR = 4000 milliseconds
        keep_alive count = 0
    keep_alive threshold = 7
           RF debug mask = 0x0
```

Implementing Redundant Supervisor Uplink Modules in Catalyst Switches

The uplink ports on a redundant Supervisor Engine are live, even when the Supervisor Engine is in standby mode. This allows redundant uplink connections from different cards, removing a single-point-of-failure concern. This topic discusses redundant supervisor uplink modules.

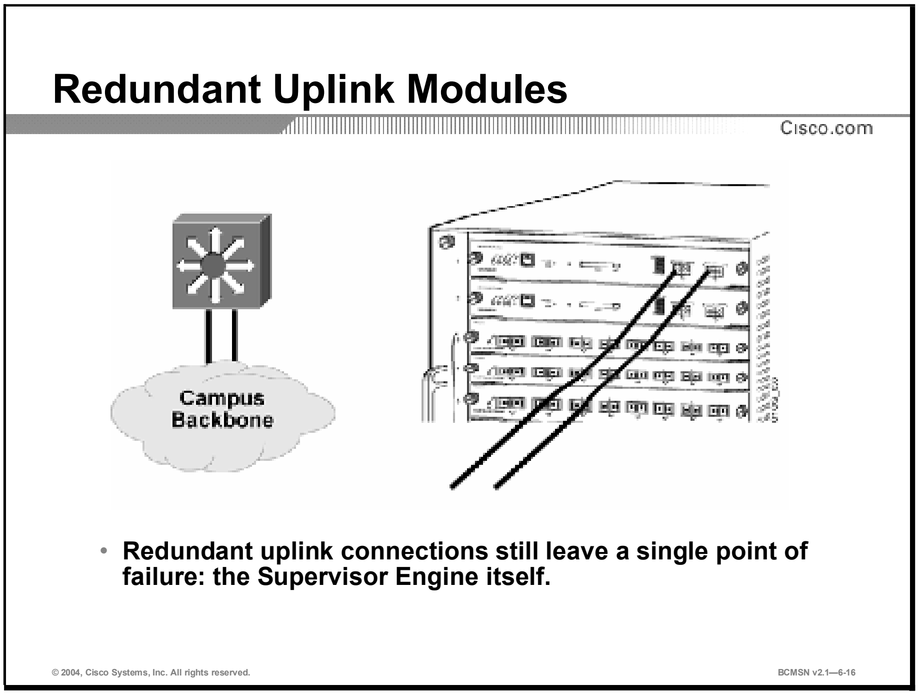

Redundant Uplink Modules

Cisco.com

- **Redundant uplink connections still leave a single point of failure: the Supervisor Engine itself.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-16

A high-availability network design calls for redundancy throughout the network. When setting up links between Building Distribution and Campus Backbone switches, for example, providing redundancy can reduce the chance of total connectivity failure by providing an alternate hardware path to the core. The uplink ports on a Supervisor Engine can be used for this purpose. However, with both links originating from the same Supervisor Engine, the connectivity is completely lost if the Supervisor Engine fails.

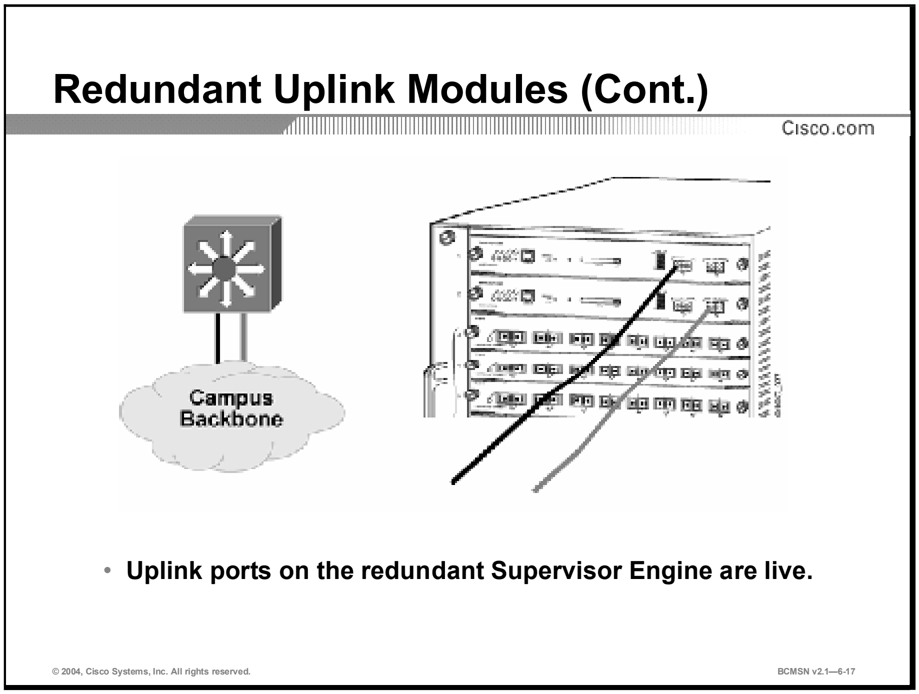

Redundant Uplink Modules (Cont.)

Cisco.com

- **Uplink ports on the redundant Supervisor Engine are live.**

BCMSN v2.1—6-17

By using the uplink ports on the redundant Supervisor Engine, you can ensure that connectivity will be maintained if the link from the primary Supervisor Engine should fail, or if the entire primary Supervisor Engine should fail.

Implementing Redundant Power Supplies

Multiple power supplies in a Catalyst switch can be configured to provide combined power or backup power in a failover situation. This topic explains how to implement redundant power supplies.

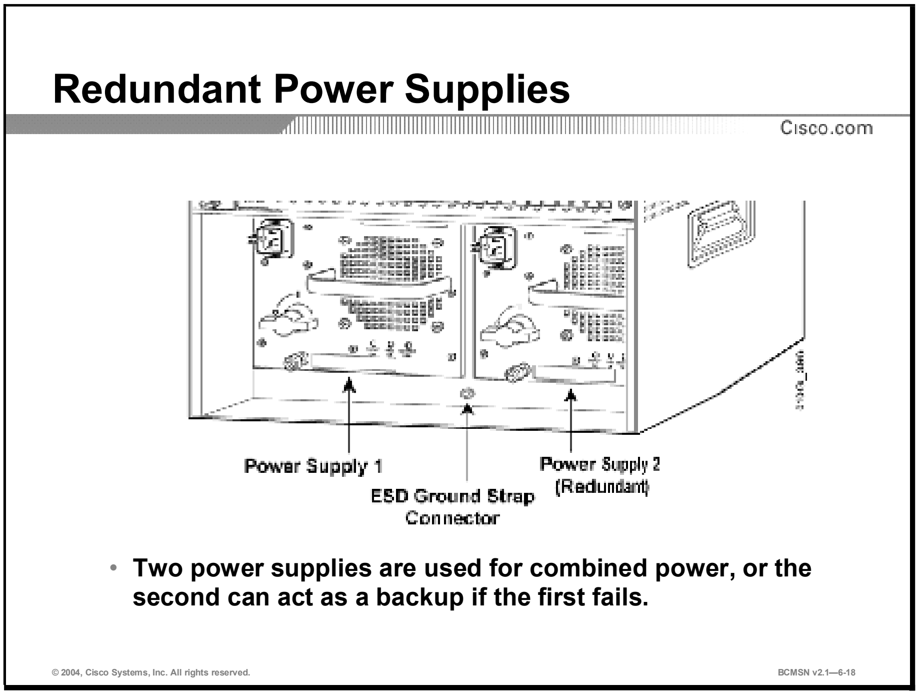

Redundant Power Supplies

Cisco.com

- **Two power supplies are used for combined power, or the second can act as a backup if the first fails.**

BCMSN v2.1—6-18

The Catalyst 6000 series of switches can be configured with dual power supplies. Depending on the specific power supplies installed and the power consumption of the modules installed in the switch, both supplies may be necessary to provide the power to the switch. However, if one of the power supplies is sufficient for meeting the power needs of the switch, the second power supply can be used as a redundant backup in case the primary power supply fails. This configuration helps ensure availability of the switch and, therefore, availability of the connectivity it provides.

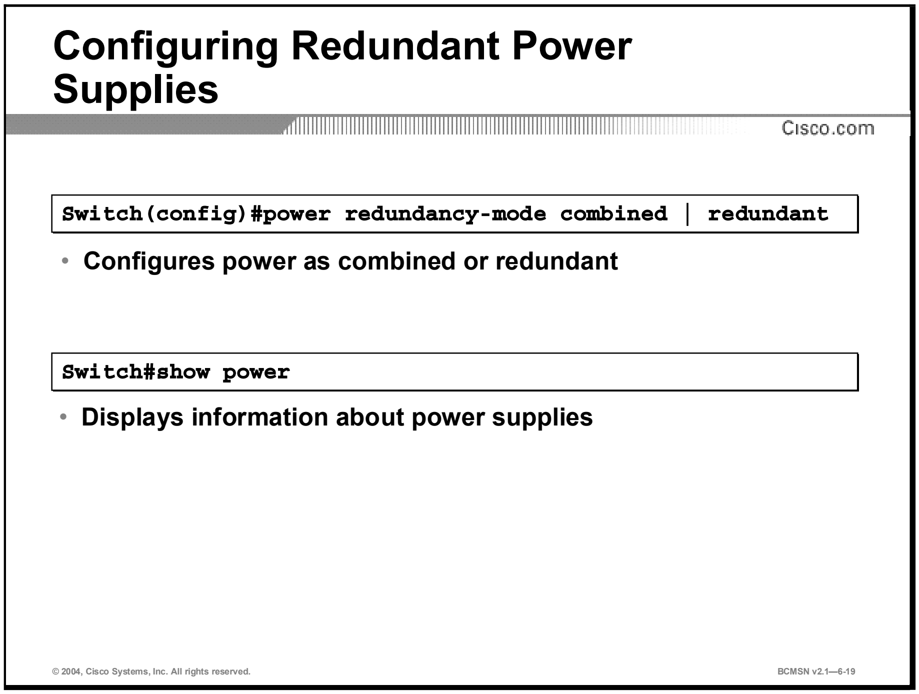

Configuring Redundant Power Supplies

Cisco.com

```
Switch(config)#power redundancy-mode combined | redundant
```

- Configures power as combined or redundant

```
Switch#show power
```

- Displays information about power supplies

BCMSN v2.1—6-19

From global configuration mode, enter the **power redundancy-mode combined | redundant** command to disable or enable redundancy (redundancy is enabled by default). You can change the configuration of the power supplies to redundant or combined at any time.

Specifying the **combined** keyword disables redundancy. In a nonredundant configuration, the power available to the system is the combined power capability of both power supplies. The system powers up as many modules as the combined capacity allows.

However, if one supply should fail and there is not enough power for all previously powered-up modules, the system powers down those modules for which there is not enough power.

Specifying the **redundant** keyword enables redundancy. In a redundant configuration, the total power drawn from both supplies is at no time greater than the capability of one supply. If one supply malfunctions, the other supply can take over the entire system load. When you install and turn on two power supplies, each concurrently provides approximately half of the required power to the system. Load sharing and redundancy are enabled automatically; no software configuration is required.

Enter the **show power** command to view the current state of modules and the total power available for modules, as follows:

```
Switch#show power
system power redundancy mode = redundant
system power total = 27.460A
system power used = -6.990A
system power available = 20.470A
FRU-type       #    current   admin state oper
power-supply   1    27.460A   on          on
module         1    -4.300A   on          on
module         2    -4.300A   off         off (admin request)
module         5    -2.690A   on          on
Switch#
```

You can power down a module from the CLI by entering the **no power enable module** *slot* command.

Note When you enter the **no power enable module** *slot* command to power down a module, the module configuration is not saved.

From global configuration mode, enter the **power enable module** *slot* command to turn the power on for a module that was previously powered down.

From global configuration mode, enter the **power cycle module** *slot* command to power cycle (reset) a module; the module powers off for 5 seconds and then powers on.

Note It is very important to maintain power budgets, especially when using inline power modules to support IP telephony.

Summary

This topic summarizes the key points discussed in this lesson.

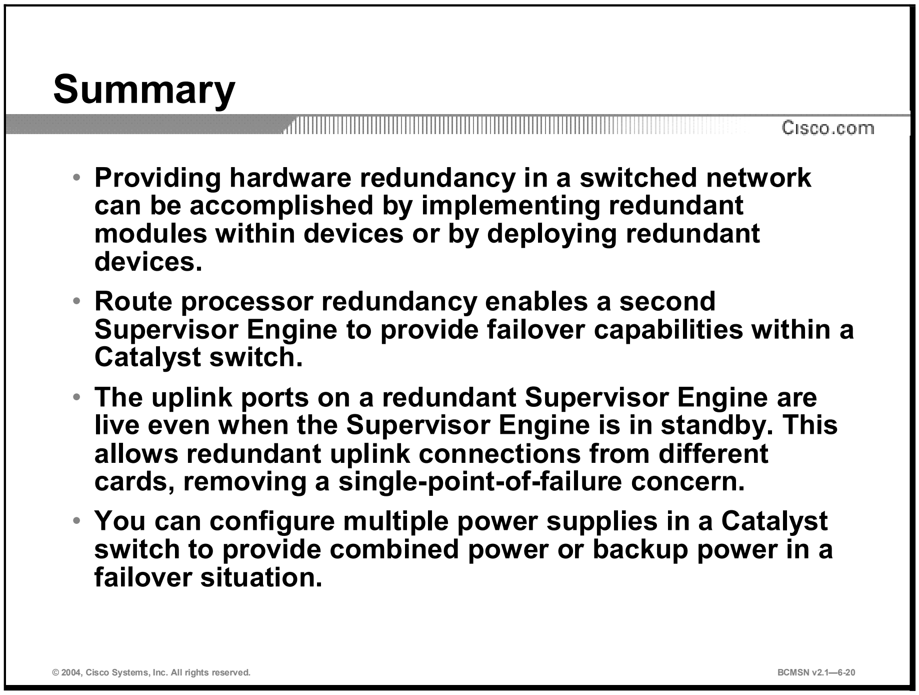

Summary

Cisco.com

- Providing hardware redundancy in a switched network can be accomplished by implementing redundant modules within devices or by deploying redundant devices.
- Route processor redundancy enables a second Supervisor Engine to provide failover capabilities within a Catalyst switch.
- The uplink ports on a redundant Supervisor Engine are live even when the Supervisor Engine is in standby. This allows redundant uplink connections from different cards, removing a single-point-of-failure concern.
- You can configure multiple power supplies in a Catalyst switch to provide combined power or backup power in a failover situation.

BCMSN v2.1—6-20

References

For additional information, refer to this resource:

- Your Cisco Catalyst switch hardware documentation

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which two drawbacks can apply to attempting to achieve reliability solely through extremely fault-tolerant devices? (Choose two.)

A) increased device cost

B) increased cabling cost

C) reduced device efficiency

D) increased administrative complexity

E) components not contributing to performance

Q2) Which command correctly configures RPR+ redundancy mode on a Catalyst switch?

A) Switch(config)#**mode rpr-plus**

B) Switch(config-red)#**mode rpr-plus**

C) Switch(config-red)#**redundancy rpr-plus**

D) Switch(config)#**redundancy mode rpr-plus**

Q3) What are two advantages of using the uplink ports on a redundant Supervisor Engine for backup connectivity? (Choose two.)

A) load balancing

B) improving security

C) eliminating a single point of failure

D) minimizing administrative complexity

Q4) Which command correctly configures redundant power on a Catalyst switch?

A) Switch#**power redundancy-mode redundant**

B) Switch#**no power redundancy-mode combined**

C) Switch(config)#**power redundancy-mode redundant**

D) Switch(config)#**no power redundancy-mode combined**

Quiz Answer Key

Q1) A, E

Relates to: Introducing Redundancy

Q2) B

Relates to: Implementing Redundant Supervisor Engines in Catalyst Switches

Q3) A, C

Relates to: Implementing Redundant Supervisor Uplink Modules in Catalyst Switches

Q4) C

Relates to: Implementing Redundant Power Supplies

Implementing Router Redundancy in a Multilayer Switched Network

Overview

Redundancy is one method for creating highly available networks. To provide failover in case of router failure in a Layer 3 switch, Cisco supports HSRP, Virtual Router Redundancy Protocol (VRRP), and Gateway Load Balancing Protocol (GLBP).

Relevance

High availability is critical for most enterprises. You should analyze the need for redundancy and select the hardware and software options required for each particular application.

Objectives

Upon completing this lesson, you will be able to:

- Explain how router redundancy operates
- Describe how HSRP operates
- Describe the HSRP states
- Configure and verify HSRP
- Describe VRRP
- Describe GLBP
- Describe, configure, and verify SRM
- Describe, configure, and verify SLB

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Understanding How Router Redundancy Works
- HSRP Operations
- HSRP States
- Configuring and Verifying HSRP
- Introducing VRRP
- Introducing GLBP
- Configuring and Verifying SRM
- Configuring and Verifying SLB
- Summary
- Quiz

Understanding How Router Redundancy Works

Router device redundancy identifies multiple routers as a single virtual router, with one router acting as the forwarding router and the others standing by to take over should the forwarding router fail. This topic describes how router redundancy works.

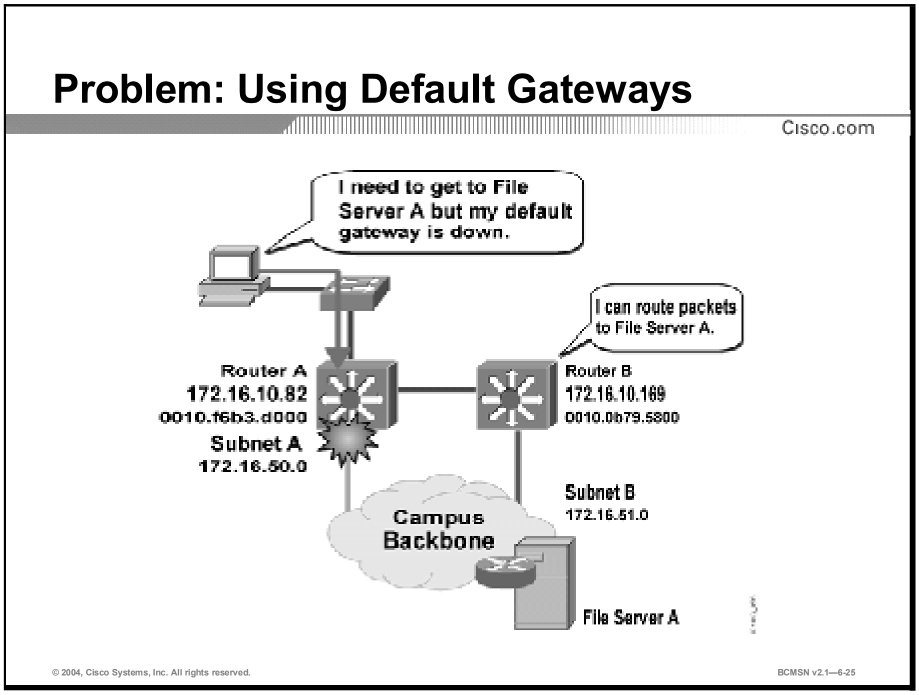

Problem: Using Default Gateways

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-25

Primary and secondary paths between the Building Access submodule and the Building Distribution submodule provide continual access in the event of a link failure at the Building Access layer. Primary and secondary paths between the Building Distribution layer and the Building Core layer provide continual operations should a link fail at the Building Distribution layer.

In this example, router A is responsible for routing packets for subnet A, and router B is responsible for handling packets for subnet B. If router A becomes unavailable to the end user, fast converging routing protocols can respond within seconds. After convergence, router B is prepared to transfer packets that would otherwise have gone through router A.

However, it is not the responsibility of the workstation, servers, and printers to exchange dynamic routing information, nor is routing by such devices a good idea. These devices typically are configured with a single default gateway IP address. If the router that is the default gateway fails, the device is limited to communicating only on the local IP network segment and is effectively disconnected from the rest of the network. Even if a redundant router exists that could serve as a default gateway, there is no dynamic method by which these devices can switch to a new default gateway IP address.

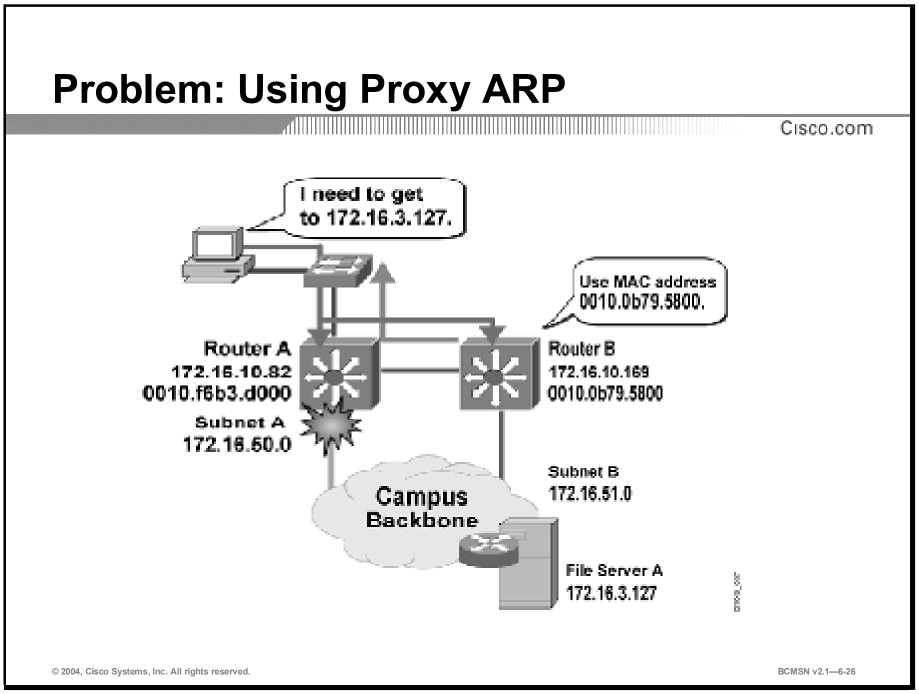

Problem: Using Proxy ARP

Cisco.com

BCMSN v2.1—6-26

Address Resolution Protocol (ARP), proxy ARP, and Reverse Address Resolution Protocol (RARP) are defined in RFCs 826, 1027, and 903, respectively. IP hosts use ARP to select a router. ARP is used to associate IP addresses with media or MAC addresses. You can configure a router as a proxy ARP server.

Cisco IOS software uses proxy ARP (as defined in RFC 1027) to help hosts that have no knowledge of routing to determine the media addresses of hosts on other networks or subnets. For example, if the router receives an ARP request for a host that is not on the same interface as the ARP request sender, and if the router has all of its routes to that host through other interfaces, it generates a proxy ARP reply packet giving its own local data-link address. The host that sent the ARP request then sends packets to the router. The router then forwards the packets to the intended host. Proxy ARP is enabled by default.

With proxy ARP, the end-user station behaves as if the destination device was connected to the same segment of the network. If the responsible router fails, the source end station continues to send packets for the destination to the MAC address of that router. Those packets subsequently are discarded.

To acquire the MAC address of the failover router, the source end station must either initiate another ARP request or be rebooted. In either case, the source end station cannot communicate with the destination for a significant period of time, even though the routing protocol has converged. The ARP protocol uses the ARP update plus the ARP flush time entry to calculate the interval in which the source cannot communicate with the destination.

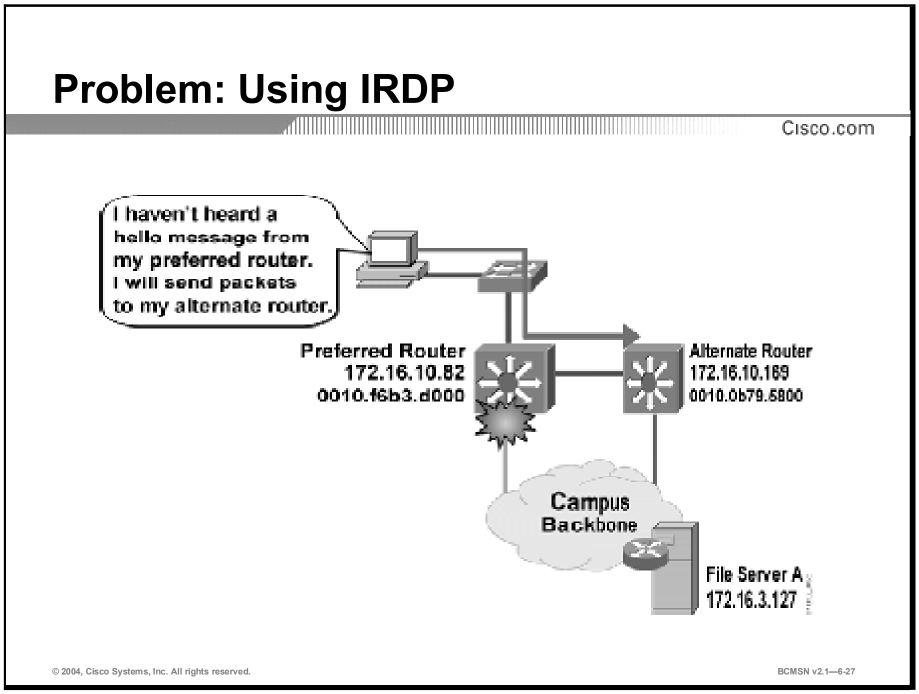

Problem: Using IRDP

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-27

Some IP hosts use the Internet Control Message Protocol (ICMP) Router Discovery Protocol (IRDP) to find a new path when the primary router becomes unavailable. IRDP is not a routing protocol like Routing Information Protocol (RIP) or the Interior Gateway Routing Protocol (IGRP). IRDP is an extension to ICMP that provides a mechanism for routers to advertise useful default routes. IRDP offers several advantages over other methods of discovering addresses of neighboring routers. IRDP does not require hosts to recognize routing protocols, nor does IRDP require manual configuration by an administrator.

A host that uses IRDP listens for hello multicast messages from the preferred default router. The IRDP-based advertisements are considered valid only for a predefined lifetime value. If a new advertisement is not seen during that lifetime, the router address is considered invalid and the host removes the corresponding default route. The IRDP protocol allows for varying timing values. A lifetime value is included in the header of every IRDP advertisement and applies to all addresses included in the packet. A host will use the router address only for the number of lifetime seconds after the most recent advertisement.

Advertisements are sent every 7 to 10 minutes; the default lifetime is 30 minutes. However, the router has complete control over the interval and lifetime values and, thus, it can control the period of time the addresses are considered valid.

IRDP has two separate interval times: a minimum and a maximum advertisement interval. All unsolicited advertisements are sent in the window of time defined by these two values. IRDP is covered in greater detail in RFC 1256.

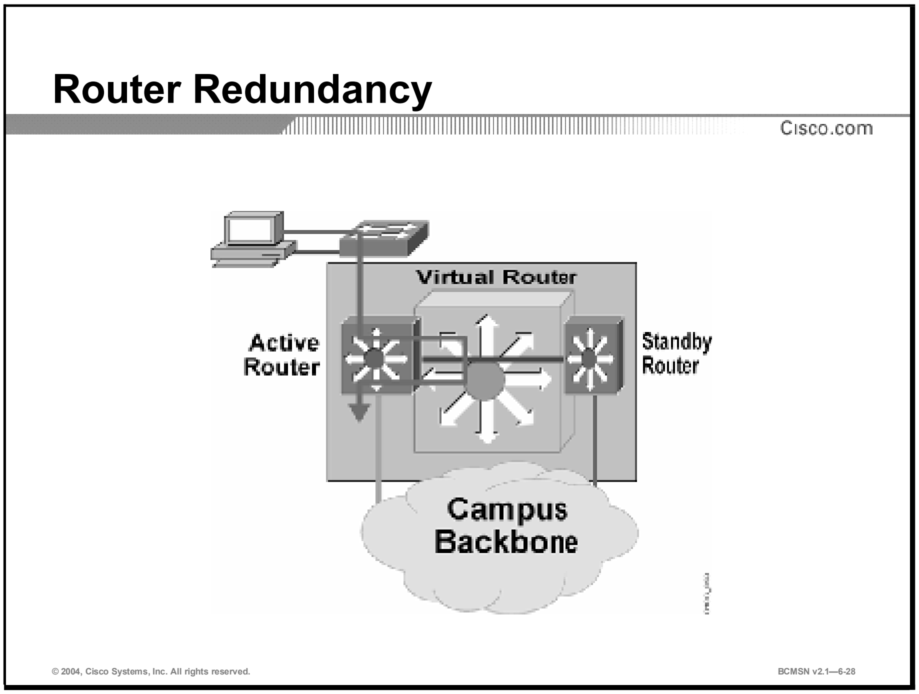

Router Redundancy

Cisco.com

BCMSN v2.1—6-28

With router device redundancy, a protocol is used to identify two or more routers as the devices responsible for a single virtual router MAC address and IP address. Host devices send traffic to the virtual router. The actual router that handles forwarding that traffic is transparent to the end stations. The redundancy protocol provides the mechanism for determining which router should take the active role in forwarding traffic, and when that role must be taken over by one of the other routers. The transition from one forwarding router to another is transparent to the end devices.

HSRP Operations

HSRP defines a standby group of routers, with one router as the active router. This topic discusses HSRP operations.

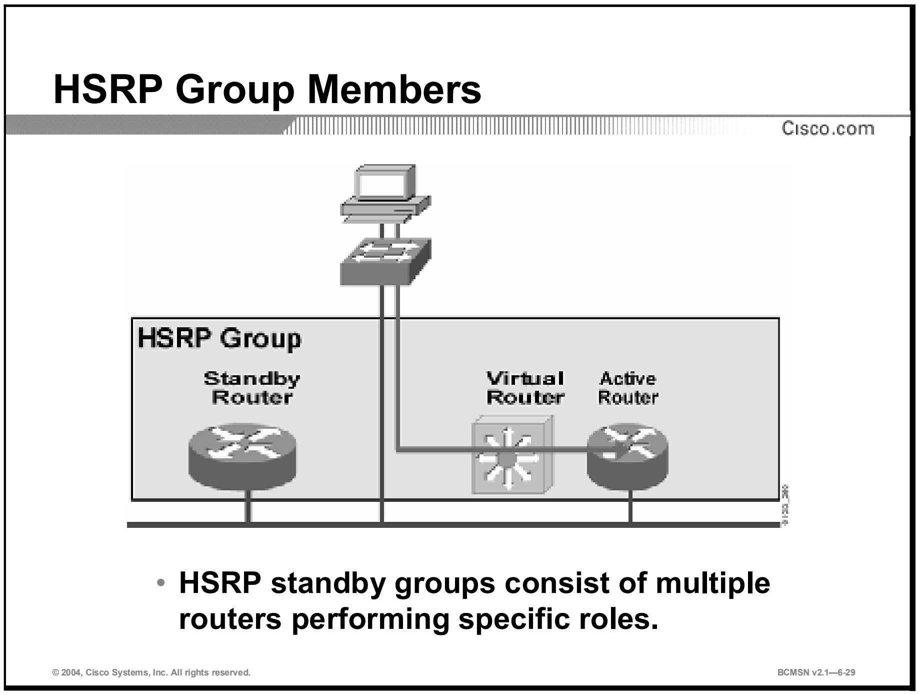

HSRP Group Members

Cisco.com

- **HSRP standby groups consist of multiple routers performing specific roles.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-29

HSRP supplies a method of providing nonstop path redundancy for IP by sharing protocol and MAC addresses between redundant gateways. The protocol consists of a virtual MAC address and a protocol address that are shared between two routers. The protocol also includes a process that monitors both LAN and serial interfaces via a multicast protocol. These terms apply to HSRP:

Term	Definition
Active router	The router that is currently forwarding packets for the virtual router
Standby router	The primary backup router
Standby group	The set of routers participating in HSRP that jointly emulate a virtual router
Hello interval time	The interval between successive HSRP hello messages from a given router
Hold interval time	The interval between the receipt of a hello message and the presumption that the sending router has failed

An HSRP standby group comprises these entities:

- One active router
- One standby router
- One virtual router
- Other routers

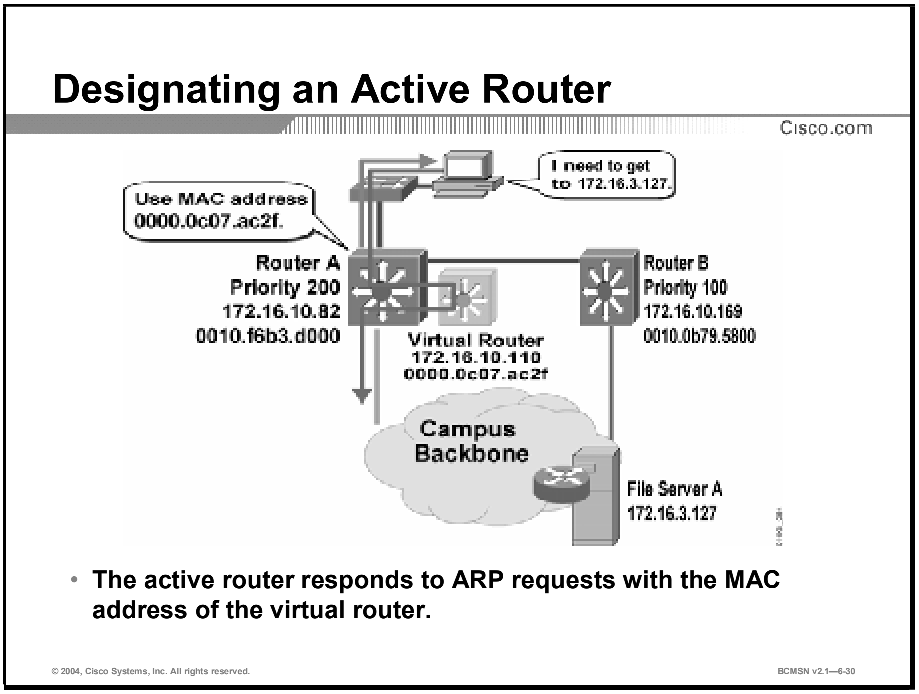

Designating an Active Router

Cisco.com

- **The active router responds to ARP requests with the MAC address of the virtual router.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-30

Within the standby group, one router is elected to be the active router. The active router forwards the packets sent to the virtual router. The router with the highest standby priority in the group becomes the active router. The default priority for an HSRP router is 100; however, the end user can change this option.

Note	When **preempt** is not configured, the first router to come up is the active router.

The active router responds to traffic for the virtual router. If an end station sends a packet to the virtual router MAC address, the active router receives and processes that packet. If an end station sends an ARP request with the virtual router IP address, the active router replies with the virtual router MAC address.

In this example, router A has a priority of 200 and router B has a default priority of 100. Router A assumes the active router role and forwards all frames addressed to the well-known MAC address of 0000.0c07.ac*xx*, where *xx* is the HSRP group identifier.

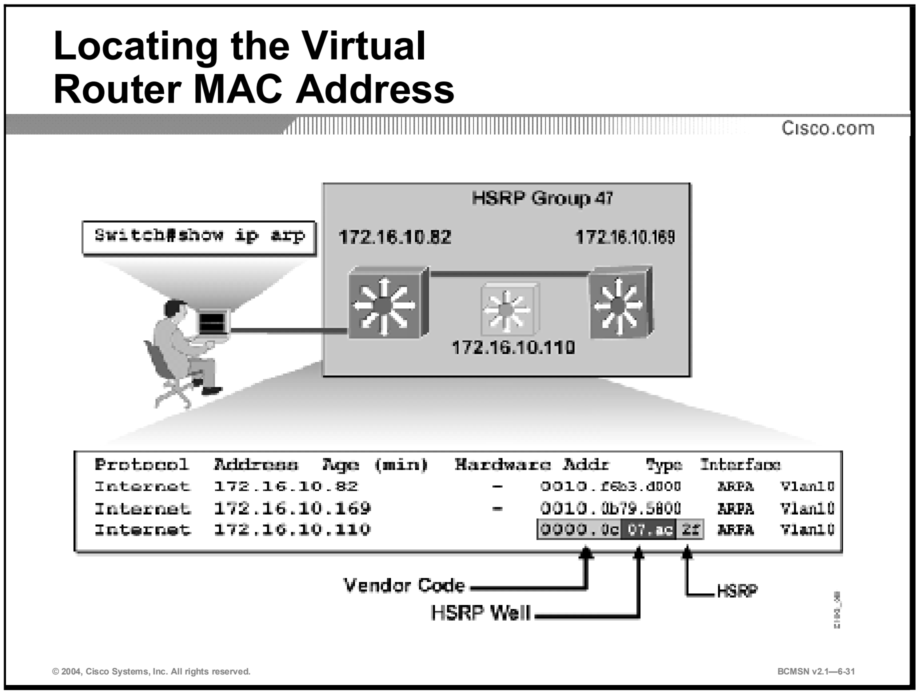

Locating the Virtual Router MAC Address

```
Switch#show ip arp
```

```
Protocol  Address         Age (min)  Hardware Addr    Type  Interface
Internet  172.16.10.82           -   0010.f6b3.d000   ARPA  Vlan10
Internet  172.16.10.169          -   0010.0b79.5800   ARPA  Vlan10
Internet  172.16.10.110              0000.0c07.ac2f   ARPA  Vlan10
```

BCMSN v2.1—6-31

ARP establishes correspondences between network addresses, such as an IP address and a hardware Ethernet address. Each router maintains a table of resolved addresses. The router checks this ARP cache before attempting to contact a device to determine if the address has already been resolved. The IP address and corresponding MAC address of the virtual router is maintained in the ARP table of each router in an HSRP standby group.

As shown in the figure, the command **show ip arp** displays the ARP cache on a router. Running this command displays the following information:

Field	Definition
Protocol	Protocol for network address in the Address field
Address	The network address that corresponds to hardware address
Age (min)	Age, in minutes, of the cache entry
Hardware Addr	The MAC address that corresponds to network address
Type	Type of encapsulation: Advanced Research Project Agency (ARPA Ethernet), Subnetwork Access Protocol (SNAP RFC 1042), or Session Announcement Protocol (SAP IEEE 802.3)
Interface	Interface to which this address mapping has been assigned

In the example, the output displays an ARP cache for a Route Switch Module (RSM) that is a member of HSRP standby group 47 in VLAN10. The virtual router for VLAN10 is identified as 172.16.10.110. The well-known MAC address that corresponds to this IP address is 0000.0c07.ac2f, where 2f is the HSRP group identifier for standby group 47. The HSRP group number is the standby group number (47) converted to hexadecimal (2f).

Note You can also obtain the HSRP virtual router IP and MAC address using the **show standby** command. The **show standby** command is discussed later in this lesson.

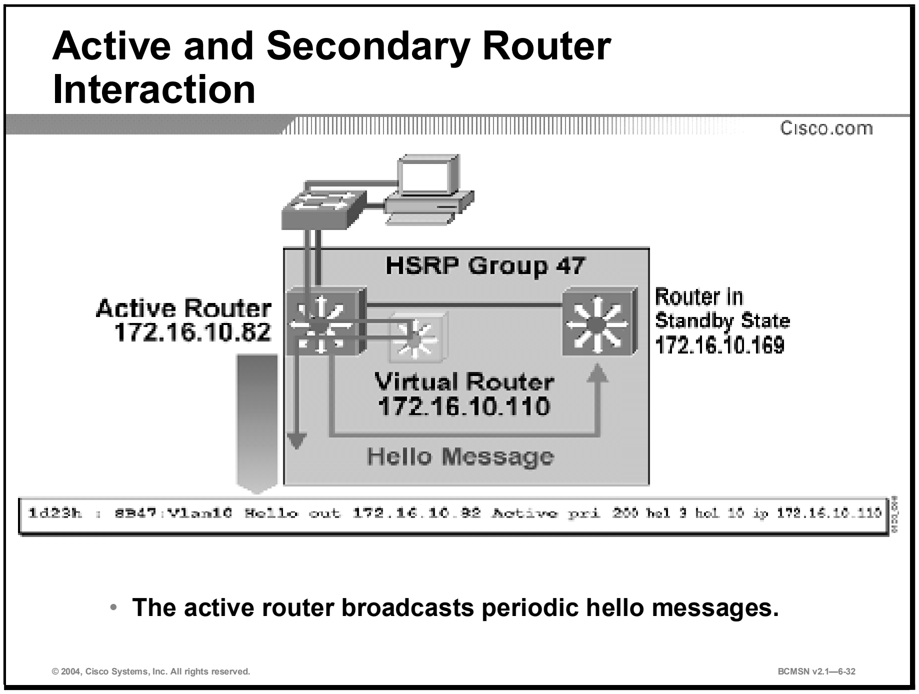

Active and Secondary Router Interaction

```
1d23h : SB47:Vlan10 Hello out 172.16.10.82 Active pri 200 hel 3 hol 10 ip 172.16.10.110
```

- **The active router broadcasts periodic hello messages.**

BCMSN v2.1—6-32

Each HSRP group contains the following:

- An active router
- A standby router
- A virtual router

The function of the active router is to forward packets sent to the virtual router. The active router assumes and maintains its active role through the transmission of hello messages.

Another router in the group is elected as the standby router. The function of the standby router is to monitor the operational status of the HSRP group and quickly assume packet-forwarding responsibility if the active router becomes inoperable. The standby router also transmits hello messages to inform all other routers in the group of the role and status of the standby router .

The function of the virtual router is to present a consistently available router to the end user. The virtual router is assigned its own IP and MAC address; however, the virtual router does not forward packets.

An HSRP standby group may contain other routers. These routers monitor the hello messages but do not respond. These routers do forward any packets addressed to the IP addresses of the routers, but do not forward packets for the virtual router. Those routers issue speak messages every hello interval time, which is the interval between the receipt of a hello message and the presumption that the sending router has failed.

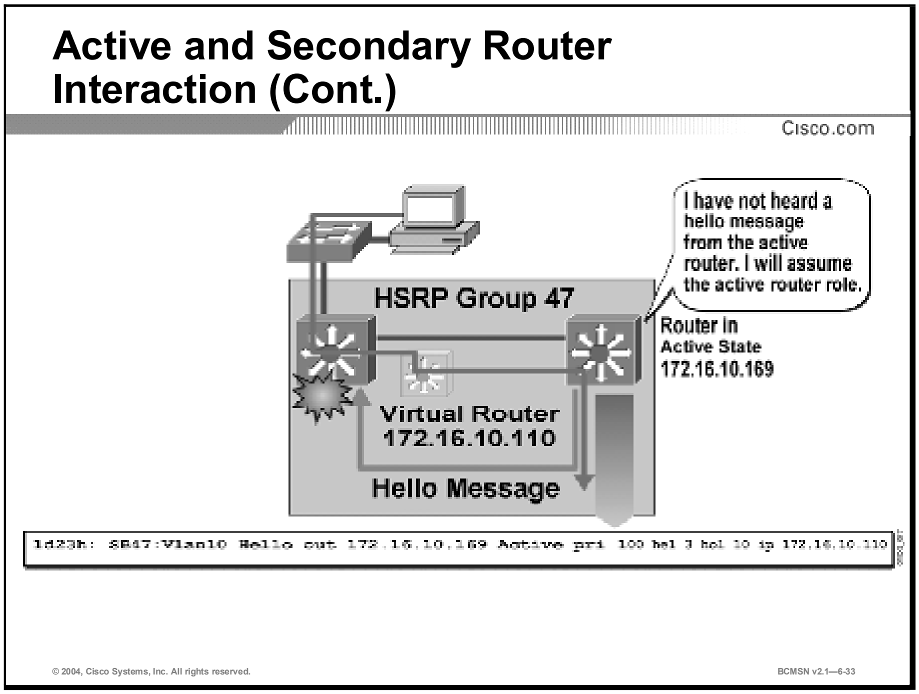

Active and Secondary Router Interaction (Cont.)

Cisco.com
HSRP Group 47
I have not heard a hello message from the active router. I will assume the active router role.
Router in Active State 172.16.10.169
Virtual Router 172.16.10.110
Hello Message

```
1d23h: SB47:Vlan10 Hello out 172.16.10.169 Active pri 100 hel 3 hol 10 ip 172.16.10.110
```

BCMSN v2.1—6-33

When the active router fails, the other HSRP routers stop receiving hello messages, and the standby router assumes the role of the active router. This occurs when the hold time expires. Therefore, the length of time it takes to fail over is dependent on the hold time.

Because the new active router assumes both the IP and MAC addresses of the virtual router, the end stations see no disruption in service. The end-user stations continue to send packets to the virtual router MAC address, and the new active router delivers the packets to the destination.

In the event that both the active and standby routers fail, all routers in the group contend for the active and standby router roles. When the active router fails, the standby takes over. If there are other routers participating in the group, those routers then contend to be the new standby router.

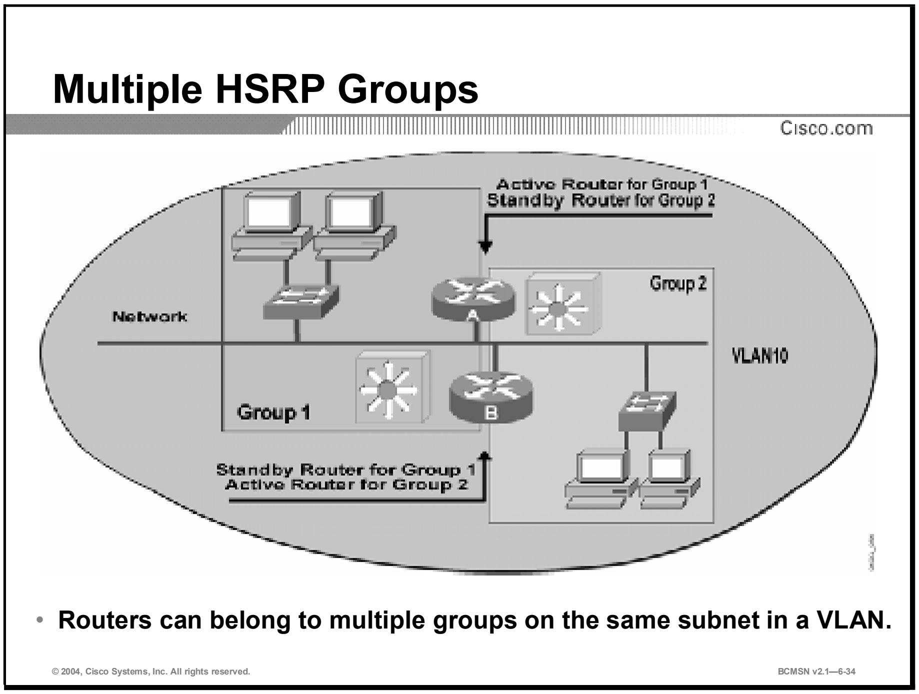

Multiple HSRP Groups

Cisco.com

Active Router for Group 1
Standby Router for Group 2
Group 2
Network
A
VLAN10
Group 1
B
Standby Router for Group 1
Active Router for Group 2

- **Routers can belong to multiple groups on the same subnet in a VLAN.**

BCMSN v2.1—6-34

To facilitate load sharing, a single router may be a member of multiple HSRP standby groups on a single segment. Multiple standby groups further enable redundancy and load sharing within networks. While a router is actively forwarding traffic for one HSRP group, the router can be in standby or listen state for another group. Each standby group emulates a single virtual router. There can be up to 255 standby groups on any LAN.

Caution Increasing the number of groups in which a router participates increases the load on the router. This can have an impact on the performance of the router.

In the figure, both router A and router B are members of groups 1 and 2. However, router A is the active forwarding router for group 1 and the standby router for group 2. Router B is the active forwarding router for group 2 and the standby router for group 1.

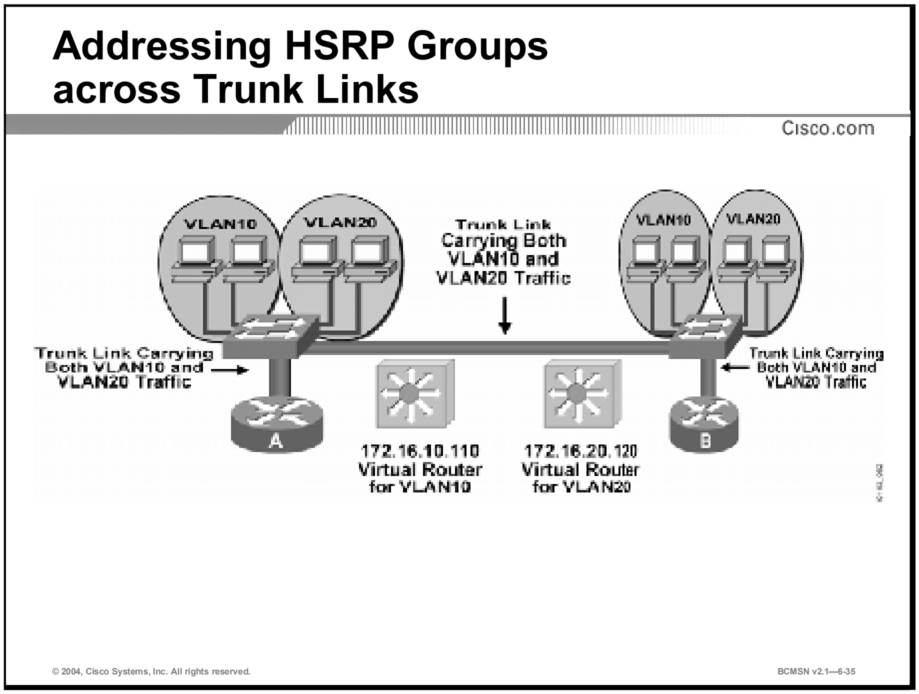

Addressing HSRP Groups across Trunk Links

Cisco.com

VLAN10
VLAN20
Trunk Link Carrying Both VLAN10 and VLAN20 Traffic
VLAN10
VLAN20
Trunk Link Carrying Both VLAN10 and VLAN20 Traffic
Trunk Link Carrying Both VLAN10 and VLAN20 Traffic
A
172.16.10.110 Virtual Router for VLAN10
172.16.20.120 Virtual Router for VLAN20
B

BCMSN v2.1—6-35

Routers can simultaneously provide redundant backup and perform load sharing across different IP subnets.

For each standby group, an IP address and a single well-known MAC address with a unique group identifier is allocated to the group.

The IP address of a group is in the range of addresses belonging to the subnet that is in use on the LAN. However, the IP address of the group must differ from the addresses allocated as interface addresses on all routers and hosts on the LAN, including virtual IP addresses assigned to other HSRP groups.

In the figure, two HSRP-enabled routers participate in two separate VLANs using Inter-Switch Link (ISL) or 802.1Q. Running HSRP over trunking allows users to configure redundancy between multiple routers that are configured as front ends for VLAN IP subnets. By configuring HSRP over ISLs, users can eliminate situations in which a single point of failure causes traffic interruptions. This feature inherently provides some improvement in overall networking resilience by providing load balancing and redundancy capabilities between subnets and VLANs.

Note A Route Processor (RP) theoretically can support up to 32,650 subinterfaces; however, the actual number of supported interfaces is limited by the capacity of the RP and the number of VLANs.

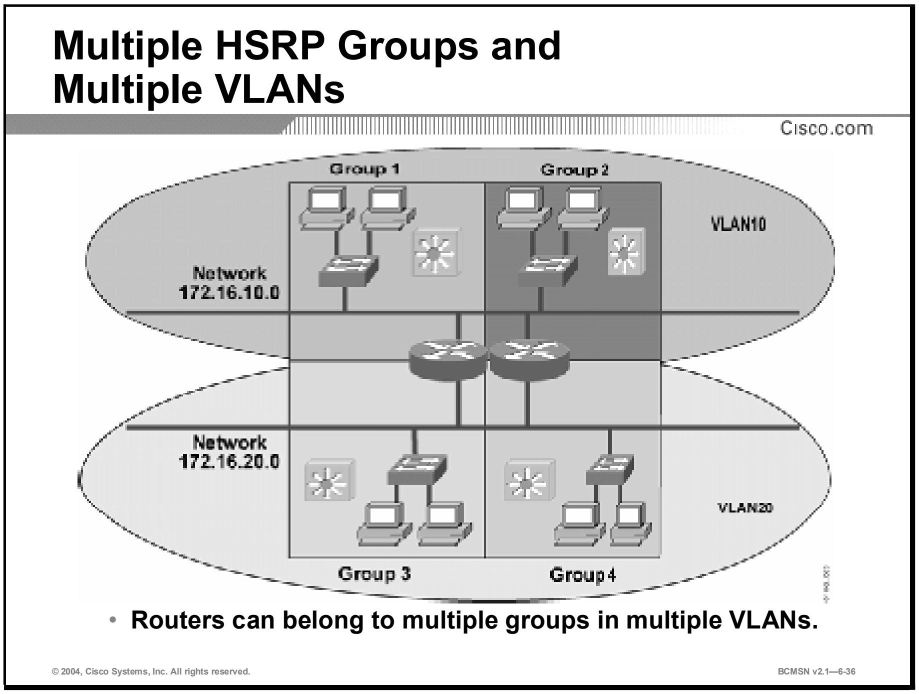

Multiple HSRP Groups and Multiple VLANs

- **Routers can belong to multiple groups in multiple VLANs.**

BCMSN v2.1—6-36

Routers can belong to multiple groups within multiple VLANs. As members of multiple hot standby groups, routers can simultaneously provide redundant backup and perform load sharing across different IP subnets.

Although multiple routers can exist in an HSRP group, only the active router forwards the packets sent to the virtual router.

Note A separate HSRP group is configured for each VLAN subnet.

HSRP States

A router in an HSRP standby group can be in one of the following states: initial, learn, listen, speak, standby, or active. This topic discusses the different HSRP states.

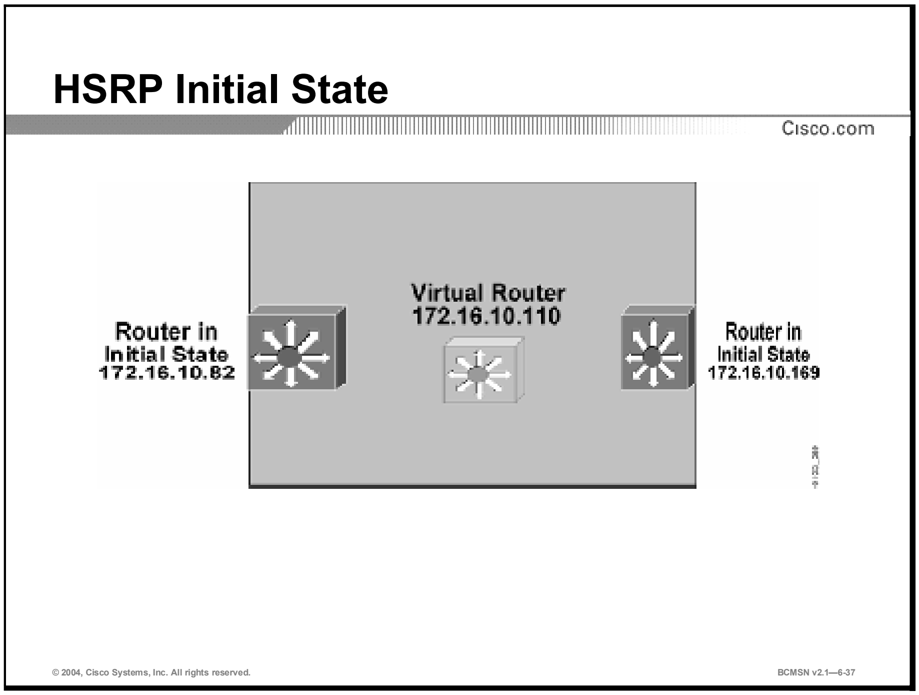

HSRP Initial State

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-37

HSRP defines six states in which an HSRP-configured router may exist. When a router exists in one of these states, the router performs the necessary actions required in that state.

The HSRP states are as follows:

- Initial state
- Listen state
- Learn state
- Speak state
- Standby state
- Active state

Not all HSRP routers will transition through all states. For example, a router that is not the standby or active router will not enter the standby or active states.

All routers begin in the initial state. This is the starting state and indicates that HSRP is not running. This state is entered via a configuration change or when an interface is initiated.

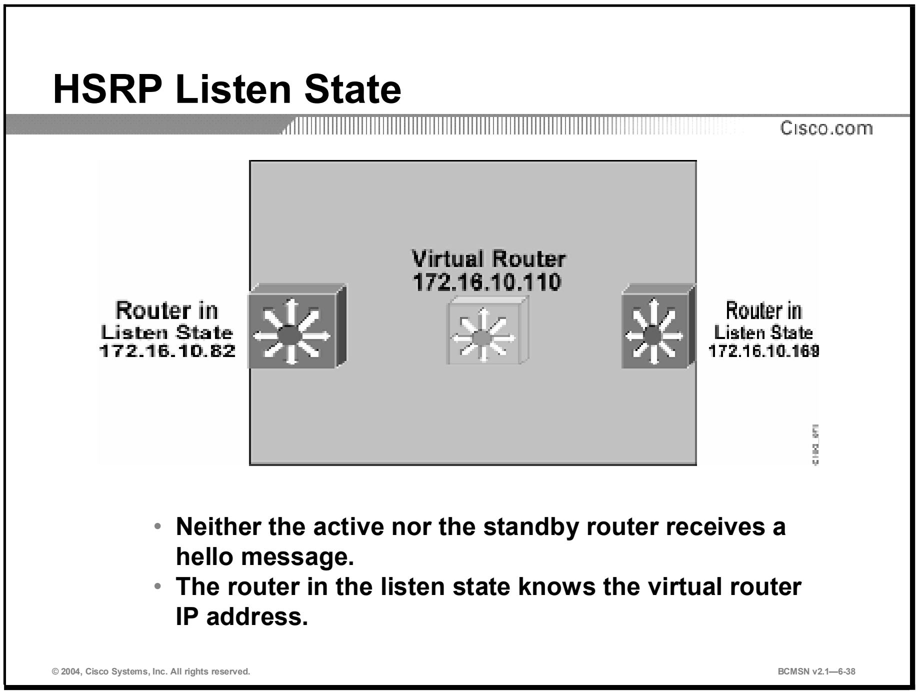

HSRP Listen State

Cisco.com

- Neither the active nor the standby router receives a hello message.
- The router in the listen state knows the virtual router IP address.

BCMSN v2.1—6-38

In the listen state, the router knows the IP address of the virtual router, but is neither the active router nor the standby router. The router listens for hello messages from those routers for its configured hold time; the purpose of this is to determine if there are active or standby routers. Then this router will join based on its configuration.

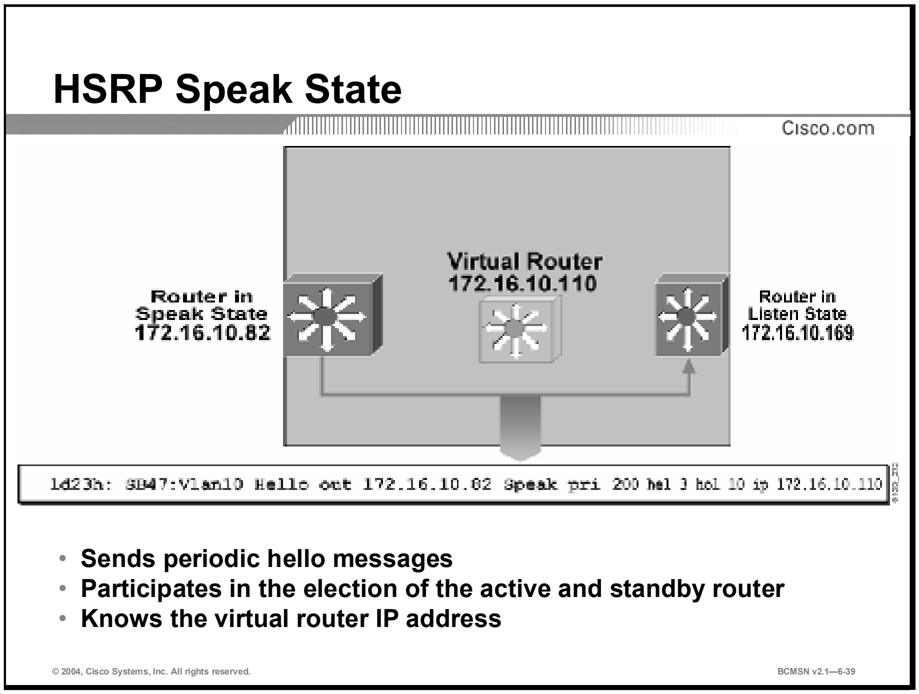

HSRP Speak State

```
1d23h: SB47:Vlan10 Hello out 172.16.10.82 Speak pri 200 hel 3 hol 10 ip 172.16.10.110
```

- **Sends periodic hello messages**
- **Participates in the election of the active and standby router**
- **Knows the virtual router IP address**

BCMSN v2.1—6-39

In the speak state, the router sends periodic hello messages and is actively participating in the election of the active router or standby router or both. A router cannot enter the speak state unless the router has the IP address of the virtual router. The router will remain in the speak state unless it becomes an active or standby router.

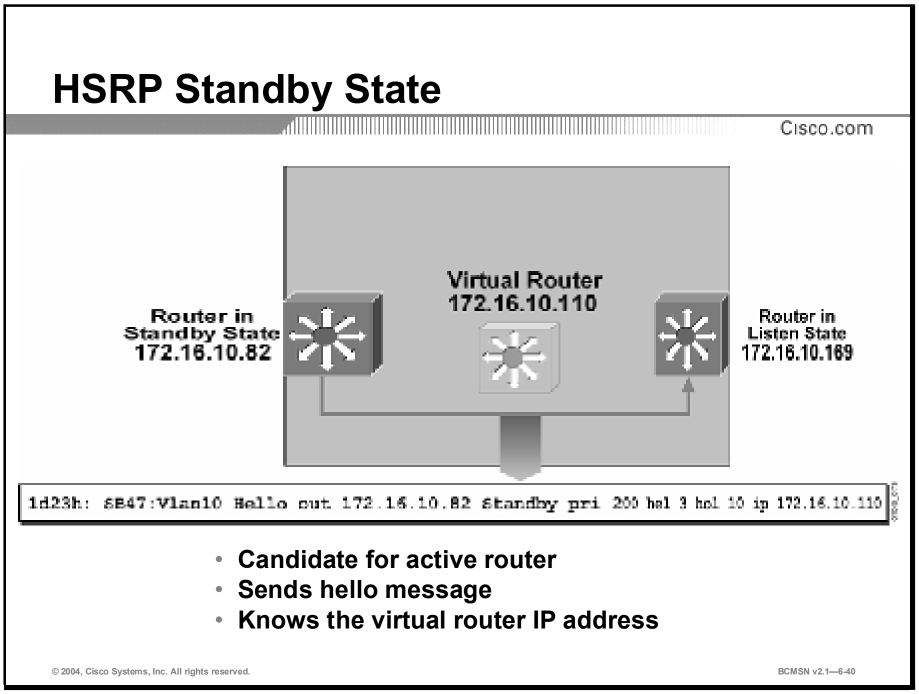

HSRP Standby State

Cisco.com

```
1d23h: SB47:Vlan10 Hello out 172.16.10.82 Standby pri 200 hel 3 hol 10 ip 172.16.10.110
```

- Candidate for active router
- Sends hello message
- Knows the virtual router IP address

BCMSN v2.1—6-40

In the standby state, the router is a candidate to become the next active router and sends periodic hello messages. There must be one standby router in the HSRP group.

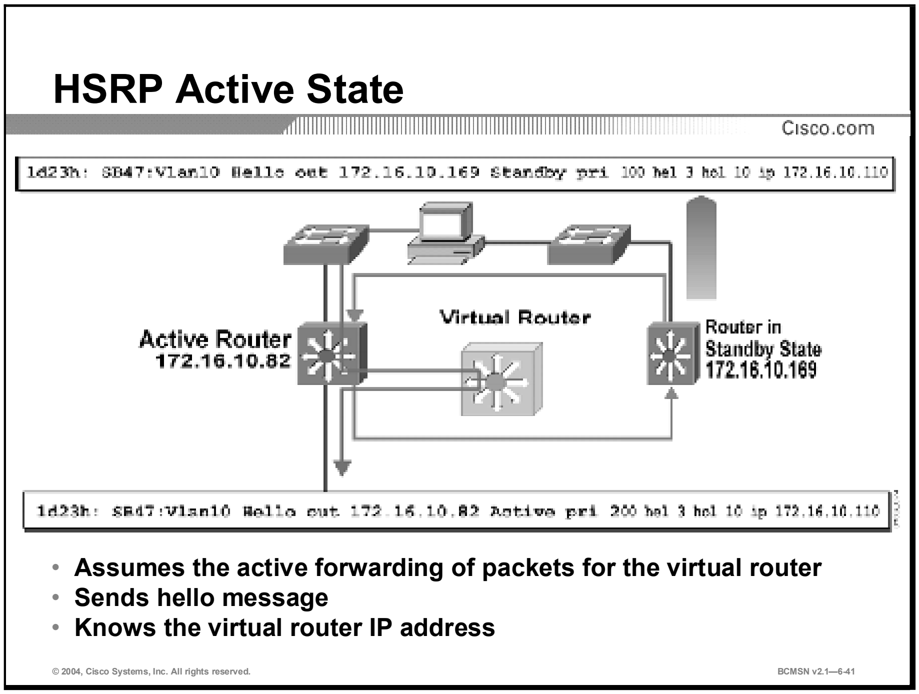

HSRP Active State

Cisco.com

```
1d23h: SB47:Vlan10 Hello out 172.16.10.169 Standby pri 100 hel 3 hol 10 ip 172.16.10.110
```

```
1d23h: SB47:Vlan10 Hello out 172.16.10.82 Active pri 200 hel 3 hol 10 ip 172.16.10.110
```

- Assumes the active forwarding of packets for the virtual router
- Sends hello message
- Knows the virtual router IP address

BCMSN v2.1—6-41

In the active state, the router is currently forwarding packets that are sent to the virtual MAC address of the group. The active router sends periodic hello messages. There must be one active router in the HSRP group.

Configuring and Verifying HSRP

Use the **standby** command from interface configuration mode to configure HSRP. This topic explains how to configure and verify HSRP.

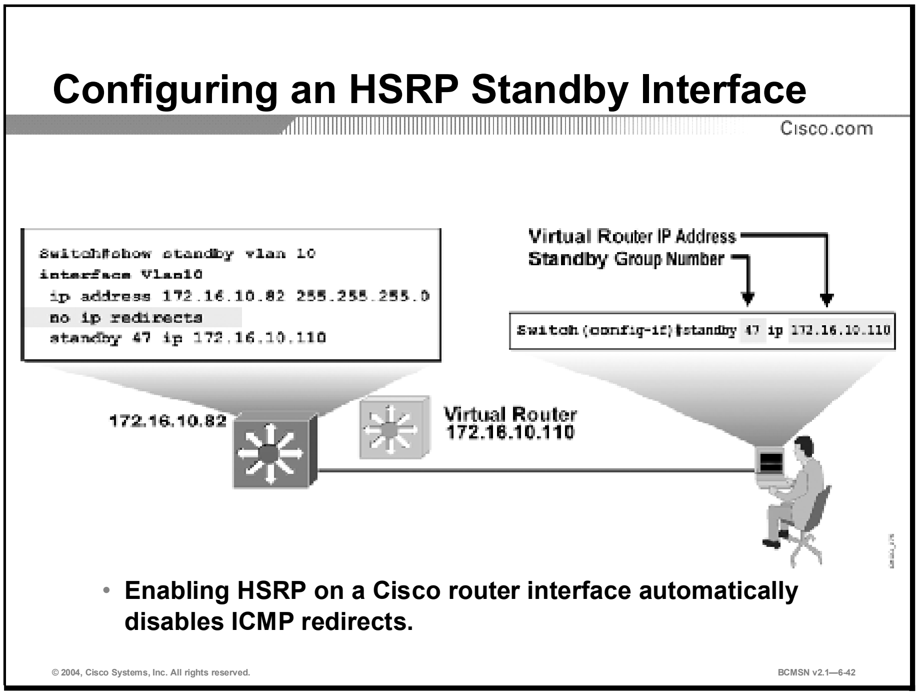

Configuring an HSRP Standby Interface

```
Switch#show standby vlan 10
interface Vlan10
 ip address 172.16.10.82 255.255.255.0
 no ip redirects
 standby 47 ip 172.16.10.110
```

```
Switch(config-if)#standby 47 ip 172.16.10.110
```

- **Enabling HSRP on a Cisco router interface automatically disables ICMP redirects.**

BCMSN v2.1—6-42

To configure a router as a member of an HSRP standby group, enter this command in interface configuration mode:

```
Switch(config-if)#standby group-number ip ip-address
```

Variable	Definition
group-number	(Optional) Indicates the HSRP group to which this interface belongs. Specifying a unique group number in the standby commands enables the creation of multiple HSRP groups. The default group is 0.
Ip-address	Indicates the IP address of the virtual HSRP router.

While running HSRP, it is important that the end-user stations do not discover the actual MAC addresses of the routers in the standby group. Any protocol that informs a host of the router actual address must be disabled. To ensure that the actual addresses of the participating HSRP routers are not discovered, enabling HSRP on a Cisco router interface automatically disables ICMP redirects on that interface.

After the **standby ip** command is issued, the interface changes to the appropriate state. When the router successfully executes the command, the router issues an HSRP message.

To remove an interface from an HSRP group, enter the **no standby** *group* **ip** command.

Example: Displaying HSRP Standby Group

The following example states that interface VLAN10 is a member of the HSRP standby group 47, the virtual router IP address for that group is 172.16.10.110, and that ICMP redirects are disabled:

```
Switch#show running-config
Building configuration...

Current configuration:
!
(text deleted)
interface Vlan10
ip address 172.16.10.82 255.255.255.0
no ip redirects
standby 47 ip 172.16.10.110
!
```

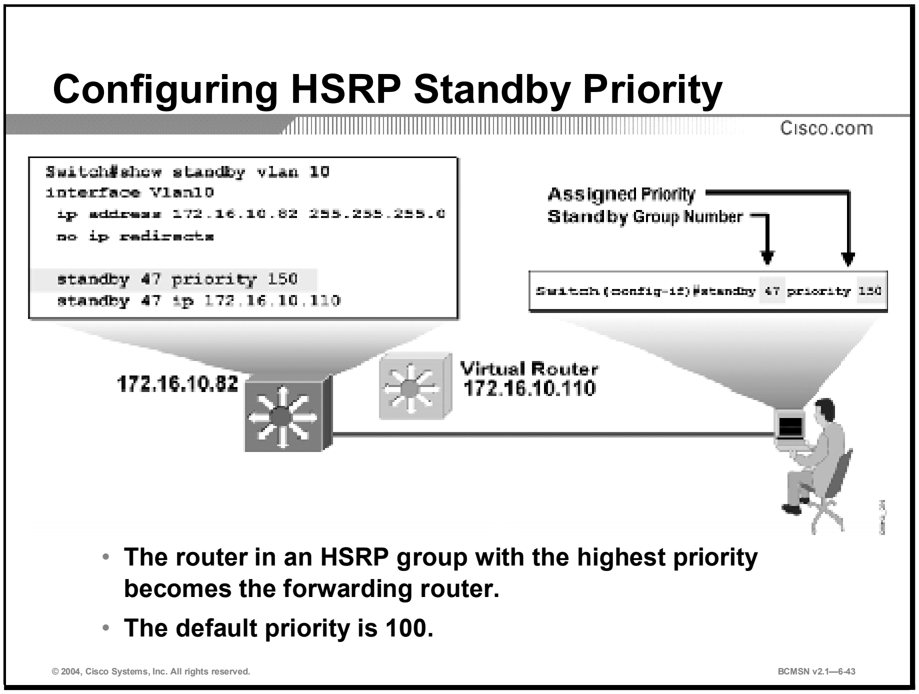

Configuring HSRP Standby Priority

```
Switch#show standby vlan 10
interface Vlan10
 ip address 172.16.10.82 255.255.255.0
 no ip redirects

 standby 47 priority 150
 standby 47 ip 172.16.10.110
```

```
Switch(config-if)#standby 47 priority 150
```

- The router in an HSRP group with the highest priority becomes the forwarding router.
- The default priority is 100.

BCMSN v2.1—6-43

Each standby group has its own active and standby routers. The network administrator can assign a priority value to each router in a standby group, allowing the administrator to control the order in which active routers for that group are selected.

To set the priority value of a router, enter this command in interface configuration mode:

```
Switch(config-if)#standby group-number priority priority-value
```

Variable	Definition
group-number	Indicates the HSRP standby group. This number can be in the range of 0 to 255.
priority-value	Indicates the number that prioritizes a potential hot standby router. The range is 0 to 255; the default is 100.

During the election process, the router in an HSRP group with the highest priority becomes the forwarding router.

To reinstate the default standby priority value, enter the **no standby priority** command.

Example: Displaying HSRP Standby Priority

The following example states that interface VLAN10 has a priority value of 150 in HSRP standby group 47. If this priority value is the highest number in that HSRP standby group, the routing device on which this interface resides is the active router for that group.

```
Switch#show running-config
Building configuration...

Current configuration:
!
(text deleted)
interface Vlan10
ip address 172.16.10.32 255.255.255.0
no ip redirects
standby 47 priority 150
standby 47 ip 172.16.10.110
```

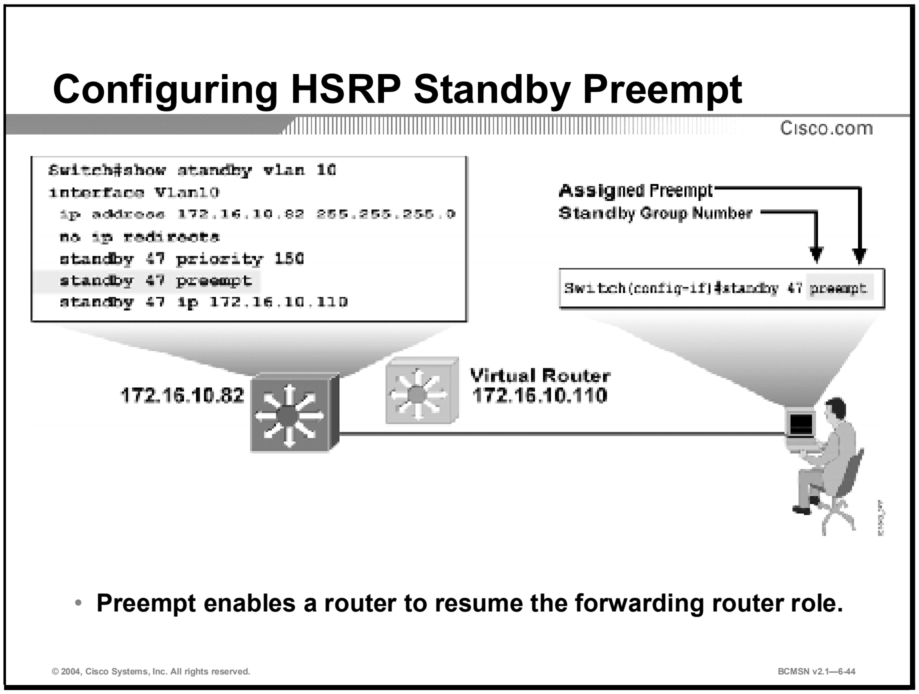

Configuring HSRP Standby Preempt

```
Switch#show standby vlan 10
interface Vlan10
 ip address 172.16.10.82 255.255.255.0
 no ip redirects
 standby 47 priority 150
 standby 47 preempt
 standby 47 ip 172.16.10.110
```

```
Switch(config-if)#standby 47 preempt
```

- **Preempt enables a router to resume the forwarding router role.**

The standby router automatically assumes the active router role when the active router fails or is removed from service. This new active router remains the forwarding router even when the former active router with the higher priority regains service in the network.

The former active router can be configured to resume the forwarding router role from a router with a lower priority. To enable a router to resume the forwarding router role, enter the following command in interface configuration mode:

```
Switch(config-if)#standby [group-number] preempt [{delay}
[minimum delay]
[sync delay]]
```

When the **standby preempt** command is issued, the interface changes to the appropriate state.

To remove the interface from preemptive status, enter the **no standby** *group* **preempt** command.

Example: Displaying HSRP Preempt

The following example states that interface VLAN10 is configured to resume its role as the active router in HSRP group 47, assuming that interface VLAN10 on this router has the highest priority in that standby group.

```
Switch#show running-config
Building configuration...

Current configuration:
!
(text deleted)
interface Vlan10
ip address 172.16.10.82 255.255.255.0
no ip redirects
standby 47 priority 150
standby 47 preempt
standby 47 ip 172.16.10.110
```

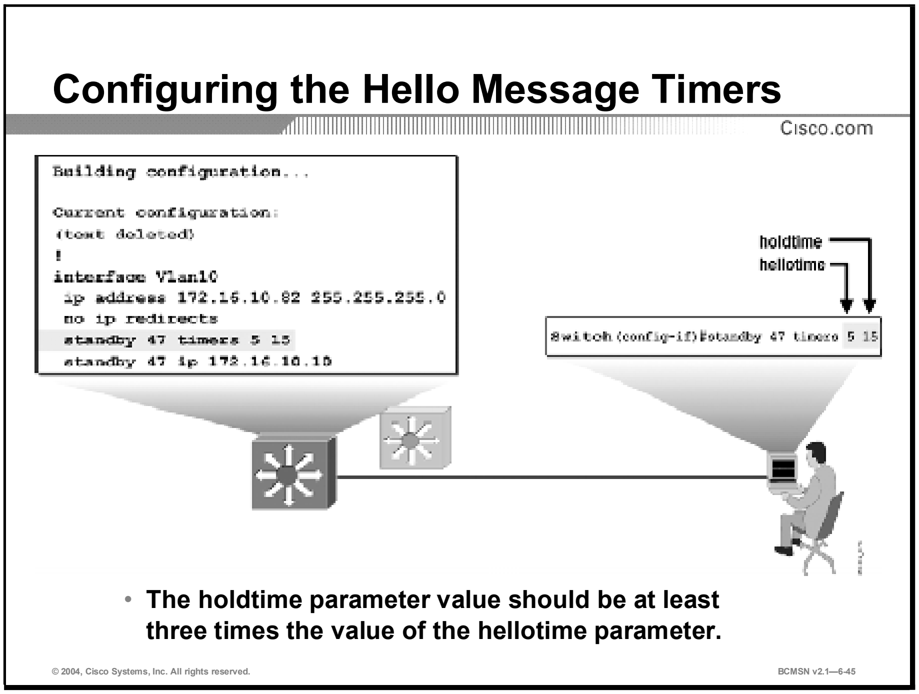

Configuring the Hello Message Timers

```
Building configuration...

Current configuration:
(text deleted)
!
interface Vlan10
 ip address 172.16.10.82 255.255.255.0
 no ip redirects
 standby 47 timers 5 15
 standby 47 ip 172.16.10.15
```

```
Switch(config-if)#standby 47 timers 5 15
```

- **The holdtime parameter value should be at least three times the value of the hellotime parameter.**

An HSRP-enabled router sends hello messages to indicate that the router is running and is capable of becoming either the active or standby router. The hello message contains the priority of the router and also hellotime and holdtime parameter values. The hellotime parameter value indicates the interval between the hello messages that the router sends. The holdtime parameter value indicates the amount of time that the current hello message is considered valid. The standby timer includes an “msec” parameter for faster failovers.

If an active router sends a hello message, receiving routers consider that hello message to be valid for one holdtime.

Note The holdtime value should be at least three times the value of the hellotime. The holdtime value must be greater than the value of the hellotime.

Both the hellotime and the holdtime parameters are configurable. To configure the time between hello messages and the time before other group routers declare the active or standby router to be nonfunctioning, enter this command in interface configuration mode:

```
Switch(config-if)#standby group-number timers hellotime
holdtime
```

Variable	Description
group-number	(Optional) Group number on the interface to which the timers apply. The default is 0.
hellotime	Hello interval in seconds. This is an integer from 1 through 255. The default is 3 seconds.
holdtime	Time, in seconds, before the active or standby router is declared to be down. This is an integer from 1 through 255. The default is 10 seconds.

To reinstate the default standby timer values, enter the **no standby** *group* **timers** command.

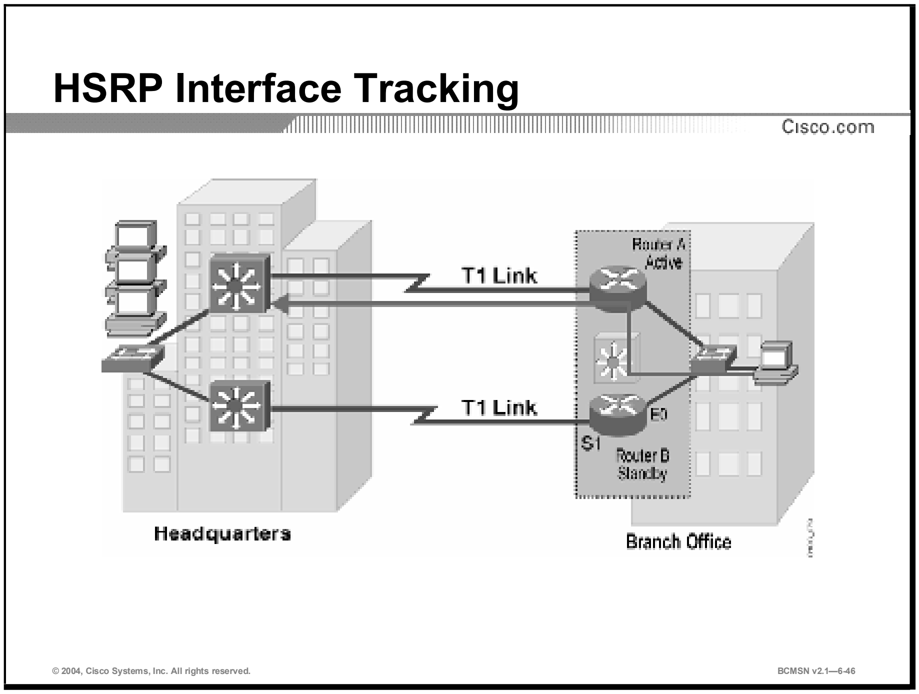

HSRP Interface Tracking

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-46

In some situations, the status of an interface directly affects which router needs to become the active router. This is particularly true when each of the routers in an HSRP group has a different path to resources within the campus network.

In this example, router A and router B reside in a branch office. These two routers each support a T1 link to the headquarters. Router A has the higher priority and is the active forwarding router for standby group 47. Router B is the standby router for that group. Router A and B are exchanging hello messages through their E0 interfaces.

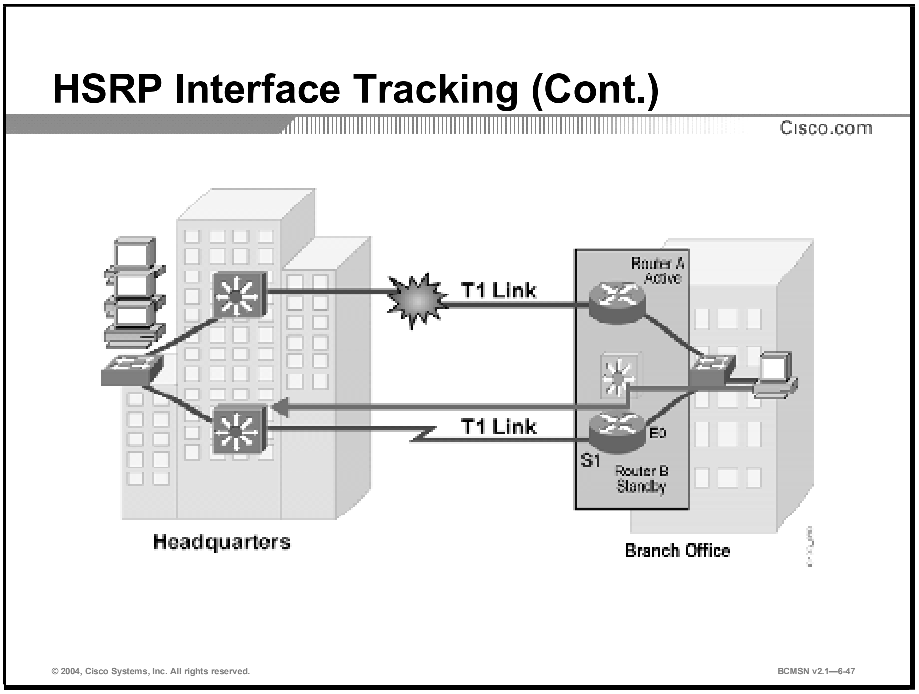

HSRP Interface Tracking (Cont.)

Cisco.com

BCMSN v2.1—6-47

The T1 link between the active forwarding router for the standby group and the headquarters experiences a failure. Without HSRP enabled, router A would detect the failed link and send an ICMP redirect to router B. However, when HSRP is enabled, ICMP redirects are disabled. Therefore, neither router A nor the virtual router sends an ICMP redirect. In addition, although the S1 interface on router A is no longer functional, router A still communicates hello messages out interface E0, indicating that router A is still the active router. Packets sent to the virtual router for forwarding to the headquarters cannot be routed.

Interface tracking enables the priority of a standby group router to be automatically adjusted, based on availability of the interfaces of that router. When a tracked interface becomes unavailable, the HSRP priority of the router is decreased. The HSRP tracking feature reduces the likelihood that a router with an unavailable key interface will remain the active router.

In this example, the E0 interface on router A tracks the S1 interface. If the link between the S1 interface and the headquarters fails, the router automatically decrements its priority on that interface and stops transmitting hello messages out interface E0. Router B assumes the active router role when no hello messages are detected for the specific hold time period.

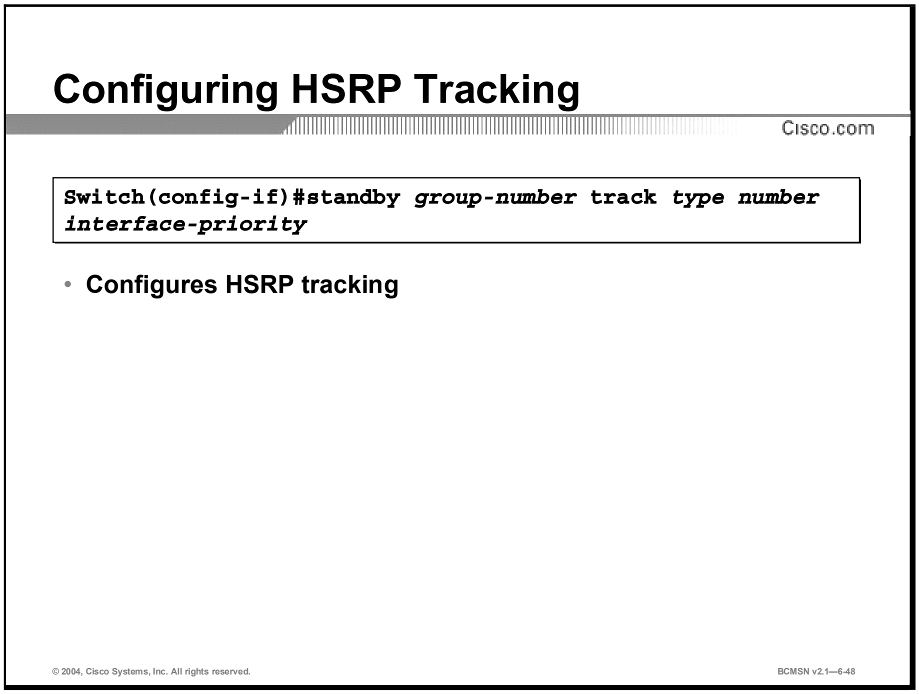

Configuring HSRP Tracking

```
Switch(config-if)#standby group-number track type number
interface-priority
```

- **Configures HSRP tracking**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-48

To configure HSRP tracking, enter this command in interface configuration mode:

```
Switch(config-if)#standby group-number track type number
interface-priority
```

Variable	Description
group-number	(Optional) Indicates the group number on the interface to which the tracking applies. The default number is 0.
type	Indicates the interface type (combined with the interface number) that will be tracked.
number	Indicates the interface number (combined with the interface type) that will be tracked.
interface-priority	(Optional) Indicates the amount by which the hot standby priority for the router is decremented when the interface becomes disabled. The priority of the router is incremented by this amount when the interface becomes available. The default value is 10.

To disable interface tracking, enter the **no standby** *group* **track** command.

The command to configure HSRP tracking on a multilayer switch is the same as on the external router, except that the interface type can be identified as a switch virtual interface (**vlan** followed by the *vlan number* assigned to that interface) or by a physical interface.

The internal routing device uses the same command as the external routing device to disable interface tracking.

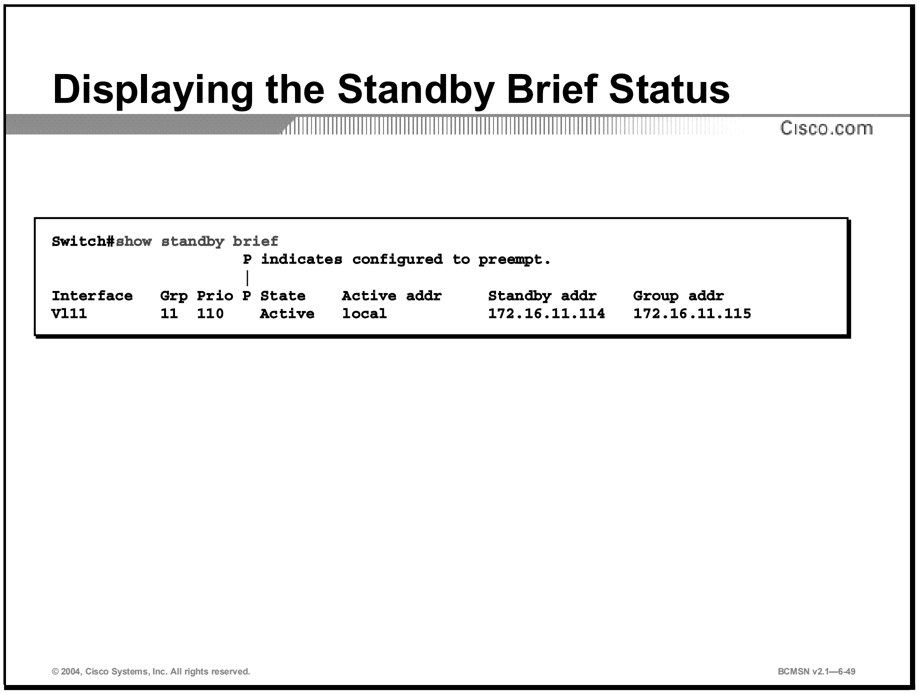

Displaying the Standby Brief Status

```
Switch#show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active addr     Standby addr    Group addr
Vl11        11  110    Active   local           172.16.11.114   172.16.11.115
```

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—6-49

To display the status of the HSRP router, enter one of these commands:

Switch#**show standby** [*interface* [*group*]] [**active** | **init** |
listen | **standby**] [**brief**]
Switch#**show standby delay** [*type-number*]

If the optional interface parameters are not indicated, the **show standby** command displays HSRP information for all interfaces.

Example: Displaying HSRP Status

This example shows the output of the **show standby** command:

```
Switch#show standby Vlan10 47
Vlan11 - Group 47
  Local state is Active, priority 150, may preempt
  Hellotime 3 holdtime 10
  Next hello sent in 00:00:02.944
  Hot standby IP address is 172.16.10.110 configured
  Active router is local
  Standby router is 172.16.10.82 expires in 00:00:08
  Standby virtual mac address is 0000.0c07.ac2f
  Tracking interface states for 1 interface, 1 up:
    Up  Vlan51 Priority decrement: 40
```

This is an example of the output resulting when you specify the **brief** parameter:

```
Switch#show standby brief
Interface   Grp   Prio P State    Active addr     Standby addr
Group addr
Vl10        47  150  P Active     local           172.16.10.82
172.16.11.10
Vl12        12  100    Standby    172.16.102.82   local
172.16.12.10
```

Note When specifying a group, you must designate an interface.

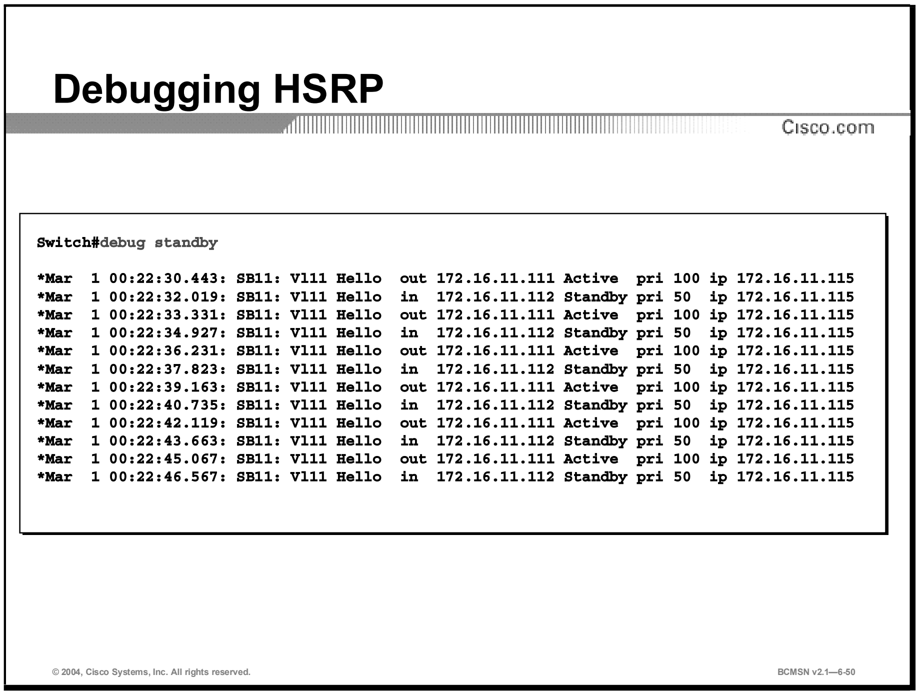

Debugging HSRP

Cisco.com

```
Switch#debug standby

*Mar  1 00:22:30.443: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip 172.16.11.115
*Mar  1 00:22:32.019: SB11: Vl11 Hello  in  172.16.11.112 Standby pri 50  ip 172.16.11.115
*Mar  1 00:22:33.331: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip 172.16.11.115
*Mar  1 00:22:34.927: SB11: Vl11 Hello  in  172.16.11.112 Standby pri 50  ip 172.16.11.115
*Mar  1 00:22:36.231: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip 172.16.11.115
*Mar  1 00:22:37.823: SB11: Vl11 Hello  in  172.16.11.112 Standby pri 50  ip 172.16.11.115
*Mar  1 00:22:39.163: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip 172.16.11.115
*Mar  1 00:22:40.735: SB11: Vl11 Hello  in  172.16.11.112 Standby pri 50  ip 172.16.11.115
*Mar  1 00:22:42.119: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip 172.16.11.115
*Mar  1 00:22:43.663: SB11: Vl11 Hello  in  172.16.11.112 Standby pri 50  ip 172.16.11.115
*Mar  1 00:22:45.067: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip 172.16.11.115
*Mar  1 00:22:46.567: SB11: Vl11 Hello  in  172.16.11.112 Standby pri 50  ip 172.16.11.115
```

BCMSN v2.1—6-50

The IOS implementation of HSRP supports the **debug** command. Enabling the debug facility displays the HSRP state changes and debugging information regarding transmission and receipt of HSRP packets. To enable HSRP debugging, enter this command in privileged EXEC mode:

```
Switch#debug standby events icmp
```

Caution Because debugging output is assigned high priority in the CPU process, this command can render the system unusable.

Example: HSRP Debugging on Negotiating for Role of Active Router

This example displays the **debug standby** command output as the router with the IP address 172.16.1.111 initializes and negotiates for the role of the active router:

```
*Mar  8 20:34:10.221: SB11: Vl11 Init: a/HSRP enabled
*Mar  8 20:34:10.221: SB11: Vl11 Init -> Listen
*Mar  8 20:34:20.221: SB11: Vl11 Listen: c/Active timer expired
(unknown)
*Mar  8 20:34:20.221: SB11: Vl11 Listen -> Speak
*Mar  8 20:34:20.221: SB11: Vl11 Hello  out 172.16.11.111 Speak
pri 100 ip 172.16.11.115
*Mar  8 20:34:23.101: SB11: Vl11 Hello  out 172.16.11.111 Speak
pri 100 ip 172.16.11.115
*Mar  8 20:34:25.961: SB11: Vl11 Hello  out 172.16.11.111 Speak
pri 100 ip 172.16.11.115
*Mar  8 20:34:28.905: SB11: Vl11 Hello  out 172.16.11.111 Speak
pri 100 ip 172.16.11.115
*Mar  8 20:34:30.221: SB11: Vl11 Speak: d/Standby timer expired
```

```
(unknown)
*Mar  8 20:34:30.221: SB11: Vl11 Standby router is local
*Mar  8 20:34:30.221: SB11: Vl11 Speak -> Standby
*Mar  8 20:34:30.221: SB11: Vl11 Hello  out 172.16.11.111 Standby
pri 100 ip 172.16.11.115
*Mar  8 20:34:30.221: SB11: Vl11 Standby: c/Active timer expired
(unknown)
*Mar  8 20:34:30.221: SB11: Vl11 Active router is local
*Mar  8 20:34:30.221: SB11: Vl11 Standby router is unknown, was
local
*Mar  8 20:34:30.221: SB11: Vl11 Standby -> Active
*Mar  8 20:34:30.221: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state
Standby -> Active
*Mar  8 20:34:30.221: SB11: Vl11 Hello  out 172.16.11.111 Active
pri 100 ip 172.16.11.115
```

To disable the debugging feature, enter either the **no debug standby** command or **the no debug all** command.

Example: HSRP Debugging on First and Only Router on Subnet

In this example, because DSW111 is the only router on the subnet and because it is not configured for preempt, this router will go through all the HSRP states before becoming the active router. Notice at time stamp Mar 8 20:34:10.221 that the interface comes up and DSW111 enters the listen state. The router stays in the listen state for the hold time of 10 seconds. DSW111 then goes into the speak state at time stamp Mar 8 20:34:20.221 for 10 seconds. When the router is speaking, it sends its state out every 3 seconds according to its hello interval. After 10 seconds in speak state, the router has determined that there is no standby router at time-stamp Mar 8 20:34:30.221 and enters the standby state. The router has also determined that there is not an active router; therefore, the router immediately enters the active state at time-stamp Mar 8 20:34:30.221. From then on, the active router will send its active state hello message every 3 seconds. Because there are no other routers on this broadcast domain, there are no hellos being received.

```
DSW111(config)#interface vlan 11
DSW111(config-if)#no shut

*Mar  8 20:34:08.925: %SYS-5-CONFIG_I: Configured from console by console
*Mar  8 20:34:10.213: %LINK-3-UPDOWN: Interface Vlan11, changed state to up
*Mar  8 20:34:10.221: SB:  Vl11 Interface up
*Mar  8 20:34:10.221: SB11: Vl11 Init: a/HSRP enabled
*Mar  8 20:34:10.221: SB11: Vl11 Init -> Listen
*Mar  8 20:34:11.213: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan11,
changed state to up
*Mar  8 20:34:20.221: SB11: Vl11 Listen: c/Active timer expired (unknown)
*Mar  8 20:34:20.221: SB11: Vl11 Listen -> Speak
*Mar  8 20:34:20.221: SB11: Vl11 Hello  out 172.16.11.111 Speak   pri 100 ip
172.16.11.115
*Mar  8 20:34:23.101: SB11: Vl11 Hello  out 172.16.11.111 Speak   pri 100 ip
172.16.11.115
*Mar  8 20:34:25.961: SB11: Vl11 Hello  out 172.16.11.111 Speak   pri 100 ip
172.16.11.115
*Mar  8 20:34:28.905: SB11: Vl11 Hello  out 172.16.11.111 Speak   pri 100 ip
172.16.11.115
*Mar  8 20:34:30.221: SB11: Vl11 Speak: d/Standby timer expired (unknown)
*Mar  8 20:34:30.221: SB11: Vl11 Standby router is local
*Mar  8 20:34:30.221: SB11: Vl11 Speak -> Standby
*Mar  8 20:34:30.221: SB11: Vl11 Hello  out 172.16.11.111 Standby pri 100 ip
172.16.11.115
*Mar  8 20:34:30.221: SB11: Vl11 Standby: c/Active timer expired (unknown)
*Mar  8 20:34:30.221: SB11: Vl11 Active router is local
*Mar  8 20:34:30.221: SB11: Vl11 Standby router is unknown, was local
*Mar  8 20:34:30.221: SB11: Vl11 Standby -> Active
*Mar  8 20:34:30.221: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state Standby -
> Active
```

```
*Mar  8 20:34:30.221: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
*Mar  8 20:34:33.085: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
*Mar  8 20:34:36.025: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
*Mar  8 20:34:38.925: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
```

Example: HSRP on NonPreempt Configured Router Coming Up

Router DSW111 is configured with a priority of 100. This priority is higher than the priority of the current active router DSW112 (172.16.11.112), which has a priority of 50. Note that router DSW111 is *not* configured with preempt. Only when a router is configured with preempt will a router with a higher priority immediately become the active router. After router DSW111 goes through the HSRP initialization states, it will come up as the standby router.

```
DSW111(config)#interface vlan 11
DSW111(config-if)#no shut

*Mar  1 00:12:16.871: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:16.871: SB11: Vl11 Active router is 172.16.11.112
*Mar  1 00:12:16.891: %SYS-5-CONFIG_I: Configured from console by console
*Mar  1 00:12:18.619: %LINK-3-UPDOWN: Interface Vlan11, changed state to up
*Mar  1 00:12:18.623: SB:  Vl11 Interface up
*Mar  1 00:12:18.623: SB11: Vl11 Init: a/HSRP enabled
*Mar  1 00:12:18.623: SB11: Vl11 Init -> Listen
*Mar  1 00:12:19.619: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan11, changed
state to up
*Mar  1 00:12:19.819: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:19.819: SB11: Vl11 Listen: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:12:22.815: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:22.815: SB11: Vl11 Listen: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:12:25.683: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:25.683: SB11: Vl11 Listen: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:12:28.623: SB11: Vl11 Listen: d/Standby timer expired (unknown)
*Mar  1 00:12:28.623: SB11: Vl11 Listen -> Speak
*Mar  1 00:12:28.623: SB11: Vl11 Hello  out 172.16.11.111 Speak   pri 100 ip
172.16.11.115
*Mar  1 00:12:28.659: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:28.659: SB11: Vl11 Speak: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:12:31.539: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:31.539: SB11: Vl11 Speak: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:12:31.575: SB11: Vl11 Hello  out 172.16.11.111 Speak   pri 100 ip
172.16.11.115
*Mar  1 00:12:34.491: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:34.491: SB11: Vl11 Speak: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:12:34.547: SB11: Vl11 Hello  out 172.16.11.111 Speak   pri 100 ip
172.16.11.115
*Mar  1 00:12:37.363: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:37.363: SB11: Vl11 Speak: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:12:37.495: SB11: Vl11 Hello  out 172.16.11.111 Speak   pri 100 ip
172.16.11.115
*Mar  1 00:12:38.623: SB11: Vl11 Speak: d/Standby timer expired (unknown)
*Mar  1 00:12:38.623: SB11: Vl11 Standby router is local
*Mar  1 00:12:38.623: SB11: Vl11 Speak -> Standby
*Mar  1 00:12:38.623: SB11: Vl11 Hello  out 172.16.11.111 Standby pri 100 ip
172.16.11.115
```

```
*Mar  1 00:12:40.279: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:40.279: SB11: Vl11 Standby: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:12:41.551: SB11: Vl11 Hello  out 172.16.11.111 Standby pri 100 ip
172.16.11.115
*Mar  1 00:12:43.191: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:43.191: SB11: Vl11 Standby: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:12:44.539: SB11: Vl11 Hello  out 172.16.11.111 Standby pri 100 ip
172.16.11.115
*Mar  1 00:12:46.167: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:46.167: SB11: Vl11 Standby: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:12:47.415: SB11: Vl11 Hello  out 172.16.11.111 Standby pri 100 ip
172.16.11.115
*Mar  1 00:12:49.119: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:12:49.119: SB11: Vl11 Standby: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:12:50.267: SB11: Vl11 Hello  out 172.16.11.111 Standby pri 100 ip
172.16.11.115
```

Example: HSRP on Preempt Configured Router Coming Up

DSW111 is configured with a priority of 100. This priority is higher than the priority of the active router DSW112 (172.16.11.112). DSW111 is also configured with preempt. Only when a router is configured with preempt will that router with a higher priority transition into the active state. At time-stamp Mar 1 00:16:43.099, the interface VLAN11 on DSW111 comes up and transitions into the listen state. At time-stamp Mar 1 00:16:43.295, DSW111 receives a hello message from the active router (DSW112). DSW111 determines that the active router has a lower priority.. At time-stamp Mar 1 00:16:43.295, DSW111 immediately sends out a coup message, indicating that DSW111 is transitioning into the active router. DSW112 enters the speak state and eventually becomes the standby router.

```
DSW111(config)#interface vlan 11
DSW111(config-if)#no shut

*Mar  1 00:16:41.295: %SYS-5-CONFIG_I: Configured from console by console
*Mar  1 00:16:43.095: %LINK-3-UPDOWN: Interface Vlan11, changed state to up
*Mar  1 00:16:43.099: SB:  Vl11 Interface up
*Mar  1 00:16:43.099: SB11: Vl11 Init: a/HSRP enabled
*Mar  1 00:16:43.099: SB11: Vl11 Init -> Listen
*Mar  1 00:16:43.295: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip
172.16.11.115
*Mar  1 00:16:43.295: SB11: Vl11 Active router is 172.16.11.112
*Mar  1 00:16:43.295: SB11: Vl11 Listen: h/Hello rcvd from lower pri Active router
(50/172.16.11.112)
*Mar  1 00:16:43.295: SB11: Vl11 Active router is local, was 172.16.11.112
*Mar  1 00:16:43.295: SB11: Vl11 Coup   out 172.16.11.111 Listen  pri 100 ip
172.16.11.115
Mar  1 00:16:43.295
*Mar  1 00:16:43.299: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state Listen -> Active
*Mar  1 00:16:43.299: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
*Mar  1 00:16:43.303: SB11: Vl11 Hello  in  172.16.11.112 Speak   pri 50 ip
172.16.11.115
*Mar  1 00:16:44.095: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan11, changed
state to up
*Mar  1 00:16:46.187: SB11: Vl11 Hello  in  172.16.11.112 Speak   pri 50 ip
172.16.11.115
*Mar  1 00:16:46.207: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
*Mar  1 00:16:49.095: SB11: Vl11 Hello  in  172.16.11.112 Speak   pri 50 ip
172.16.11.115
*Mar  1 00:16:49.195: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
*Mar  1 00:16:52.079: SB11: Vl11 Hello  in  172.16.11.112 Speak   pri 50 ip
172.16.11.115
```

```
*Mar  1 00:16:52.147: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
*Mar  1 00:16:53.303: SB11: Vl11 Hello  in  172.16.11.112 Standby pri 50 ip
172.16.11.115
*Mar  1 00:16:53.303: SB11: Vl11 Standby router is 172.16.11.112
*Mar  1 00:16:55.083: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
*Mar  1 00:16:56.231: SB11: Vl11 Hello  in  172.16.11.112 Standby pri 50 ip
172.16.11.115
*Mar  1 00:16:58.023: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
*Mar  1 00:16:59.223: SB11: Vl11 Hello  in  172.16.11.112 Standby pri 50 ip
172.16.11.115
*Mar  1 00:17:00.983: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
*Mar  1 00:17:02.211: SB11: Vl11 Hello  in  172.16.11.112 Standby pri 50 ip
172.16.11.115
*Mar  1 00:17:03.847: SB11: Vl11 Hello  out 172.16.11.111 Active  pri 100 ip
172.16.11.115
```

Introducing VRRP

Virtual Router Redundancy Protocol (VRRP) enables a group of routers to form a single virtual router. This topic introduces VRRP.

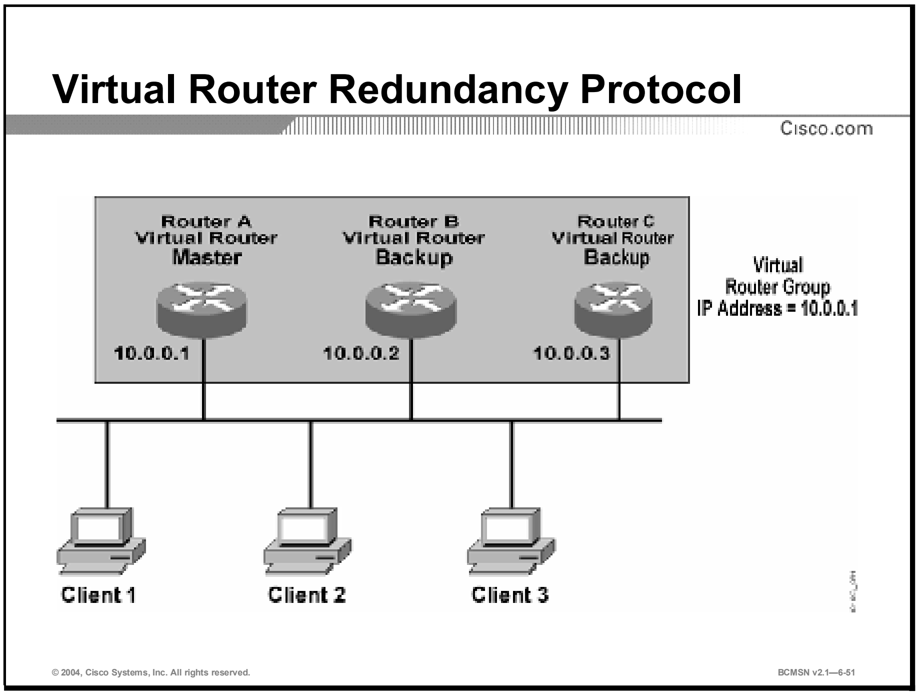

Virtual Router Redundancy Protocol

Cisco.com

BCMSN v2.1—6-51

There are several ways a LAN client can determine which router should be the first hop to a particular remote destination. The client can use a dynamic process or static configuration. The following are examples of dynamic router discovery:

- **Proxy ARP:** The client uses ARP to get the destination address. A router will send its own MAC address as a response to the ARP request.
- **Routing protocol:** The client listens to dynamic routing protocol updates (for example, from RIP) and forms its own routing table.
- **IRDP client:** The client runs an ICMP router discovery client.

The drawback to dynamic discovery protocols is that they incur some configuration and processing overhead on the LAN client. Also, in the event of a router failure, the process of switching to another router can be slow.

An alternative to dynamic discovery protocols is to statically configure a default router on the client. This approach simplifies client configuration and processing, but creates a single point of failure. If the default gateway fails, the LAN client is limited to communicating on the local IP network segment only and is cut off from the rest of the network.

The VRRP standard can solve the static configuration problem. VRRP enables a group of routers to form a single virtual router. The LAN clients can then be configured with the virtual router as their default gateway. The virtual router, representing a group of routers, is also known as a VRRP group.

VRRP is supported on Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces, and on Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs) and VLANs.

For example, the figure shows a LAN topology in which VRRP is configured. In this example, routers A, B, and C are VRRP routers (routers running VRRP) that make up a virtual router. The IP address of the virtual router is the same as that configured for the Ethernet interface of router A (10.0.0.1).

Because the virtual router uses the IP address of the physical Ethernet interface of router A, router A assumes the role of the master virtual router and is also known as the IP address owner. As the master virtual router, router A controls the IP address of the virtual router and is responsible for forwarding packets sent to this IP address. Clients 1 through 3 are configured with the default gateway IP address of 10.0.0.1.

Routers B and C function as backup virtual routers. If the master virtual router fails, the router configured with the higher priority will become the master virtual router and provide uninterrupted service for the LAN hosts. When router A recovers, it becomes the master virtual router again.

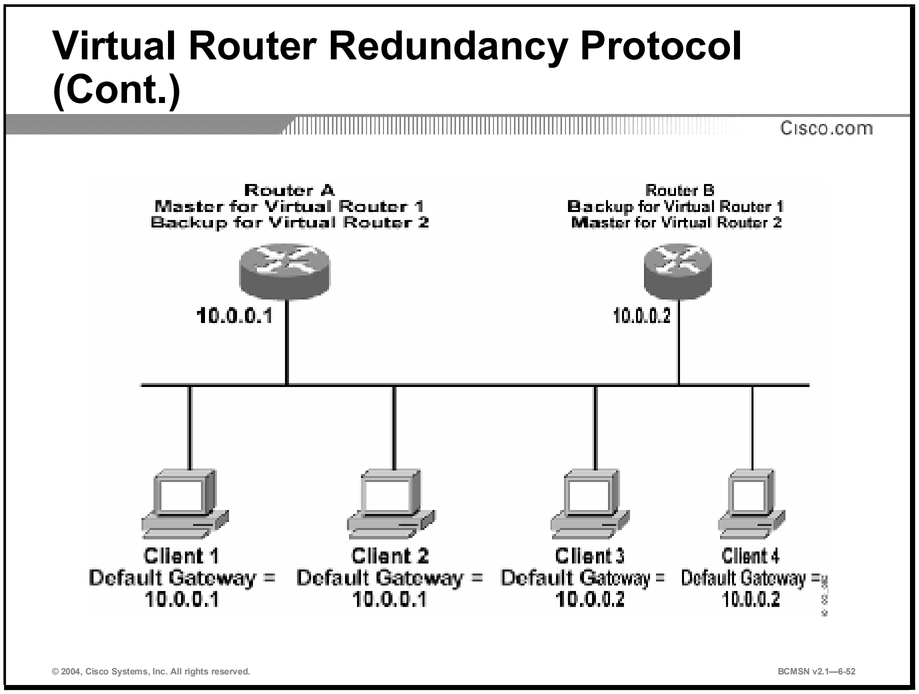

Virtual Router Redundancy Protocol (Cont.)

BCMSN v2.1—6-52

This figure shows a LAN topology in which VRRP is configured so that routers A and B share the traffic to and from clients 1 through 4, and that routers A and B act as backup virtual routers to each other if either router fails.

In this topology, two virtual routers are configured. For virtual router 1, router A is the owner of IP address 10.0.0.1 and is the master virtual router for clients configured with the default gateway IP address of 10.0.0.1.1.. Router B is the backup virtual router to router A.

For virtual router 2, router B is the owner of IP address 10.0.0.2 and is the master virtual router for clients configured with the default gateway IP address of 10.0.0.2. Router A is the backup virtual router to router B.

VRRP offers these redundancy features:

- VRRP provides redundancy for the real IP address of a router, or for a virtual IP address shared among the VRRP group members.
- If a real IP address is used, the owning router must be the master.
- If a virtual IP address is used, an election process takes place. The router with the highest priority wins the role of master.
- A VRRP group has one master router and one or more backup routers. The master router uses VRRP messages to inform group members of the IP addresses of the backup routers.

If a real IP address is used by the VRRP group, the owning router must be the master in the VRRP group. The priority of the master must be 255. Otherwise, the router with the highest priority wins the election and becomes the master. Backup values range from 1 to 254; the default value is 100. A priority of zero (0) indicates that the master is releasing responsibility for the virtual router.

The IP protocol number assigned by the Internet Assigned Numbers Authority (IANA) for VRRP is 112. Only the master sends the advertisement on multicast 224.0.0.18. The suggested default interval is 1 second, depending on implementation. You will back up the router and preempt the master if its priority is higher.

A VRRP flow message is similar in concept to an HSRP coup message. You can prohibit preemption through configuration. A master with a priority of zero triggers a transition to a backup router. The result is similar to when you use the HSRP resign message.

If the IP address owner is available, it will always become the master. Otherwise, the router with the highest priority becomes the master.

The dynamic failover when the active (master) becomes unavailable uses two timers within VRRP: the advertisement interval and the master-down interval. The advertisement interval is the time interval between advertisements (seconds). The default is 1 second. The master-down interval is the time interval for backup to declare the master down (seconds).

Introducing GLBP

This topic introduces GLBP.

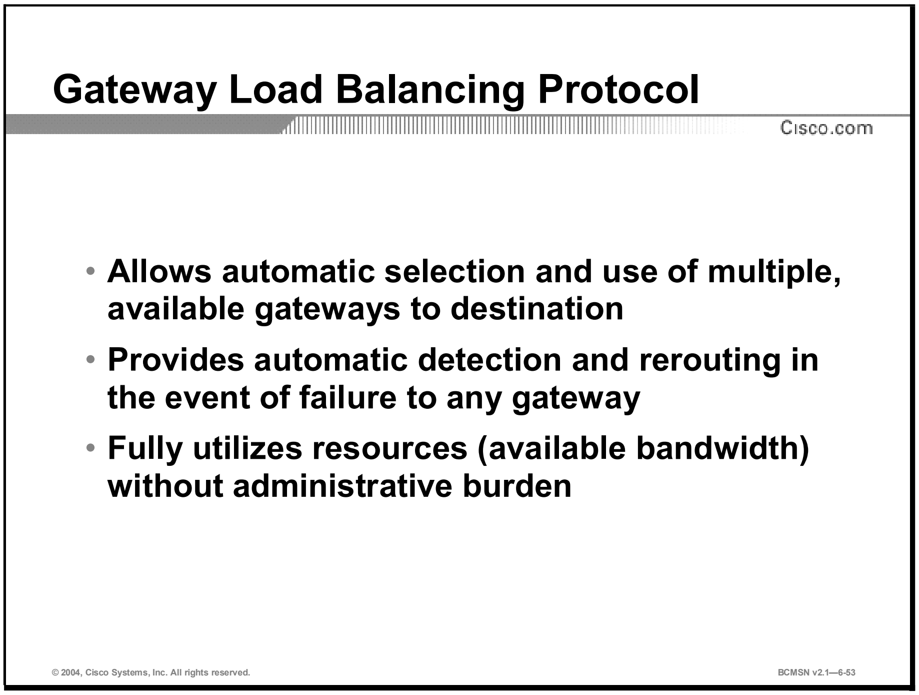

Gateway Load Balancing Protocol

- Allows automatic selection and use of multiple, available gateways to destination
- Provides automatic detection and rerouting in the event of failure to any gateway
- Fully utilizes resources (available bandwidth) without administrative burden

BCMSN v2.1—6-53

While HSRP and VRRP provide gateway resiliency, the standby members of the redundancy group are underutilized along with their upstream bandwidth. Only the active router for the HSRP group handles traffic for the virtual MAC. Bandwidth and resources associated with the standby router are not fully utilized.

Cisco designed GLBP to allow automatic selection and simultaneous use of multiple, available gateways, and to provide automatic detection and failover to a redundant path in the event of failure to any active gateway. With GLBP, you can fully utilize resources without the extra administrative burden of configuring multiple groups and managing multiple default gateway configurations.

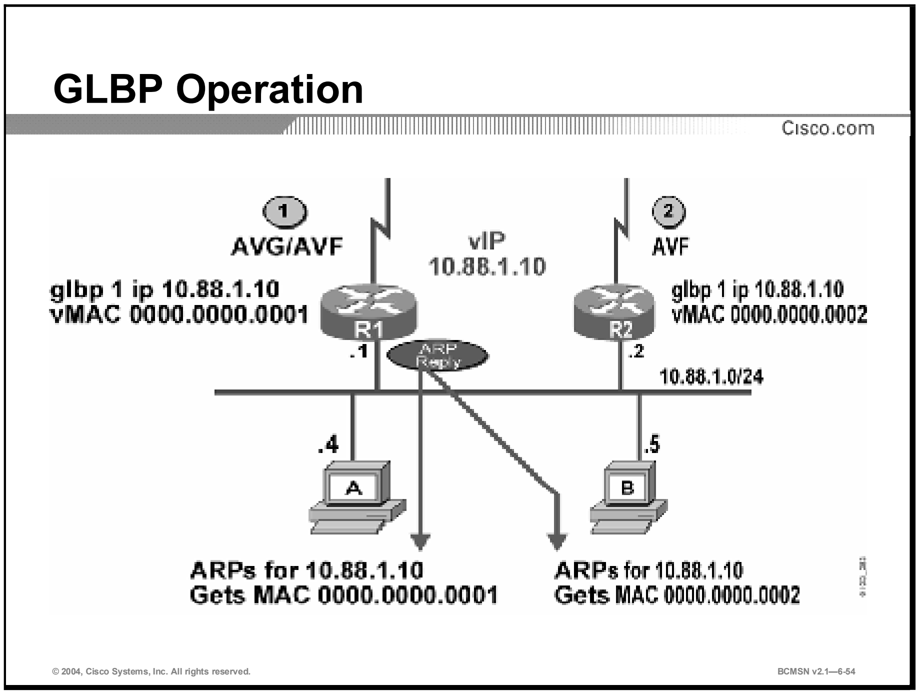

GLBP Operation

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-54

GLBP supports these operational modes for load balancing:

- **Weighted load-balancing algorithm:** The amount of load directed to an active virtual forwarder is dependent upon the weighting value advertised by the gateway containing that active virtual forwarder.
- **Host-dependent load-balancing algorithm:** A host is guaranteed to use the same virtual MAC address as long as that virtual MAC address is participating in the GLBP group.
- **Round-robin load-balancing algorithm:** Each virtual forwarder MAC address takes turns being included in address resolution replies for the virtual IP address.

GLBP allows automatic selection and simultaneous use of multiple, available gateways. The members of a GLBP group elect one gateway to be the active virtual gateway (AVG) for that group. Other members of the group provide backup for the AVG if it should become unavailable. The AVG assigns a virtual MAC address to each member of the GLBP group. These gateways become the active virtual forwarder (AVF) for that virtual MAC address and forwards packets sent to the virtual MAC address.

In the default mode, GLBP will attempt to balance traffic on a per-host basis using a round-robin scheme, as shown in the figure. When a device sends an ARP message for the gateway IP address, the AVG will return a MAC address based on a load-balancing algorithm. When a second device sends an ARP message, the AVG returns the next virtual MAC from the list of available gateways.

Note Only the AVG responds to ARP requests. The AVG is responsible for managing the members of the GLBP group.

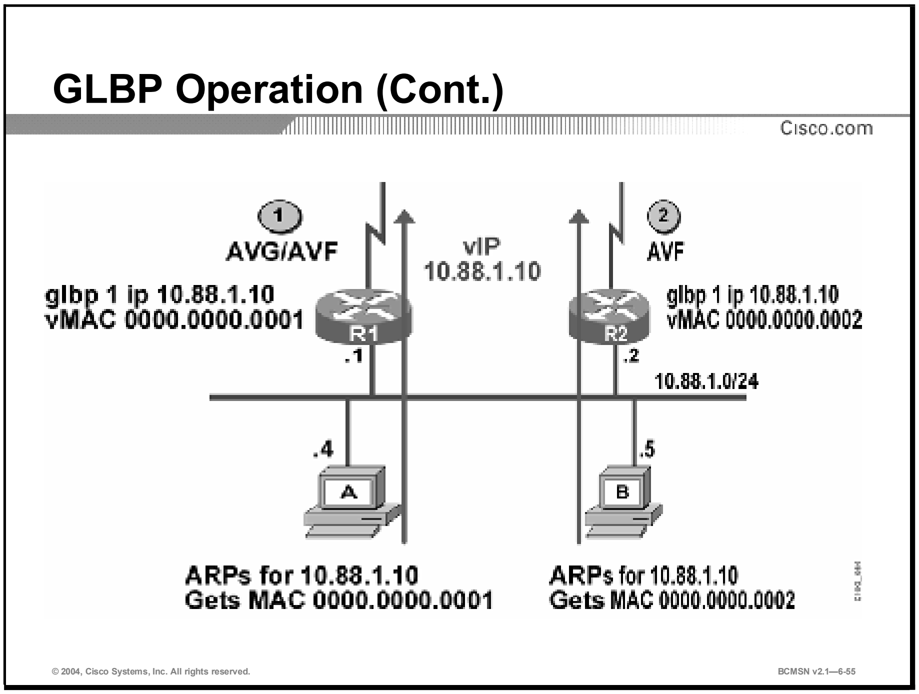

GLBP Operation (Cont.)

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-55

Now, clients A and B send their off-net traffic to separate routers; this is because they each have cached a different MAC address for the single virtual gateway IP address (in this case, 10.88.1.10). Each GLBP router is an active virtual forwarder for the MAC address that it has been assigned.

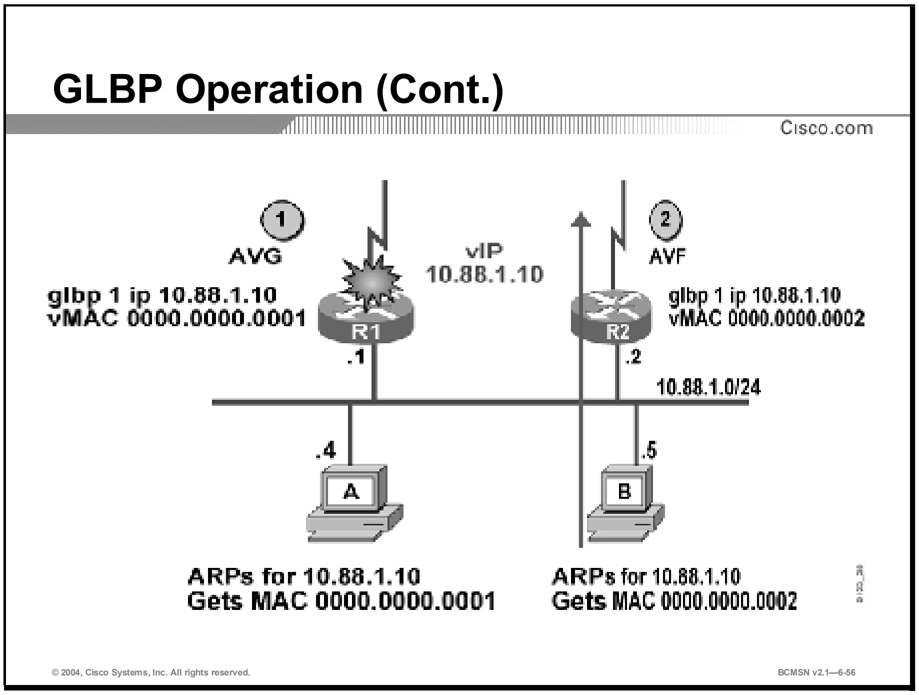

GLBP Operation (Cont.)

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-56

GLBP can be configured to track interfaces, just as with HSRP. In the figure, the link from router R1 is lost. GLBP detects the failure.

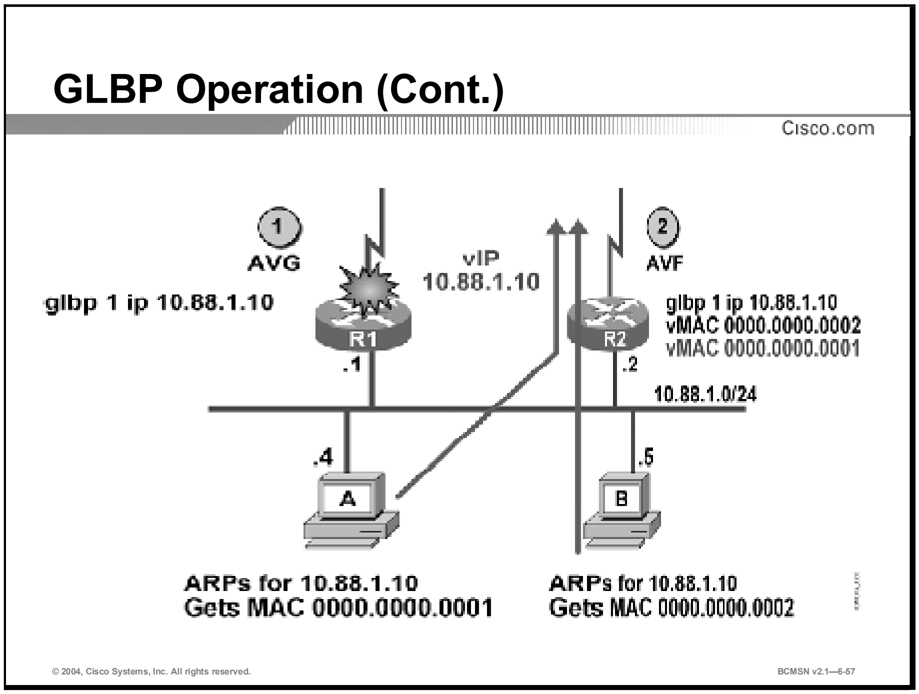

GLBP Operation (Cont.)

Cisco.com

ARPs for 10.88.1.10
Gets MAC 0000.0000.0001

ARPs for 10.88.1.10
Gets MAC 0000.0000.0002

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-57

In this example, the job of forwarding packets for virtual MAC address 0000.0000.0001 will be taken over by the secondary virtual forwarder for the MAC, router R2.

A GLBP group can have up to four member routers acting as gateways. The GLBP automatically manages the virtual MAC address assignment, determines who handles the forwarding, and ensures that each station has a forwarding path in the event of failures to gateways or tracked interfaces. If failures occur, the load-balancing ratio is adjusted among the remaining active virtual forwarders so that resources are used in the most efficient way.

Configuring and Verifying SRM

With Single Router Mode (SRM) redundancy, only the designated router MSFC2 is visible to the network at any given time. This topic explains how to configure and verify SRM.

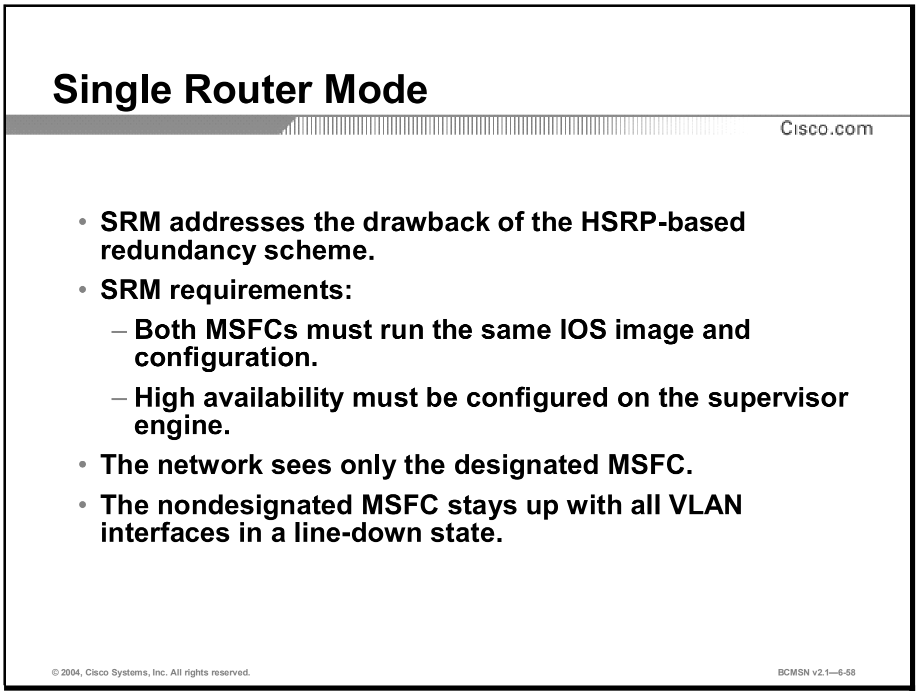

Single Router Mode

- **SRM addresses the drawback of the HSRP-based redundancy scheme.**
- **SRM requirements:**
  - **Both MSFCs must run the same IOS image and configuration.**
  - **High availability must be configured on the supervisor engine.**
- **The network sees only the designated MSFC.**
- **The nondesignated MSFC stays up with all VLAN interfaces in a line-down state.**

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—6-58

SRM redundancy is an alternative to internally redundant (dual) MSFC2 configurations where both MSFC2s are active at the same time. In SRM redundancy, only the designated router is visible to the network at any given time. The nondesignated router is booted up completely and participates in configuration synchronization, which is automatically enabled when entering SRM. All of the configuration following the **alt** keyword is ignored in SRM. Because of this, the configuration of the nondesignated router is exactly the same as the configuration of the designated router. However, the interfaces of the nondesignated router are kept in a line-down state and are not visible to the network. Processes, such as routing protocols, are created on the nondesignated router and the designated router. However, because all nondesignated router interfaces are in a line-down state, they do not send or receive updates from the network.

When the designated router fails, the nondesignated router changes its state from a nondesignated router to a designated router, and its interface state changes to link up. This new designated router builds its routing table while the existing Supervisor Engine switch processor entries are used to forward Layer 3 traffic. After the newly designated router builds its routing table, the entries in the switch processor are updated.

SRM requirements include the following:

- Both MSFCs must run the same IOS image and configuration.
- High availability needs to be configured on the Supervisor Engine.
- The network sees only the designated MSFC.
- The nondesignated MSFC stays up with all VLAN interfaces in a line-down state but is completely booted.
- The configuration is allowed only on the designated MSFC.

When SRM is enabled, the nondesignated router is online, but has all of its interfaces down. Therefore, it does not hold any routing table information. This means that if the designated router fails, there will be some delay before the nondesignated router coming online will have a complete route table. To help account for this, the information being used before the failure by the Supervisor Engine for Layer 3 forwarding is maintained and updated with any new information from the new designated router.

SRM and SUP II/PFC2/MSFC 2 Failure Scenarios

If the SRM and SUP II/PFC 2/MSFC 2 begin to fail:

1. The designated router is failing.
2. The new designated router brings up its VLAN interfaces.
3. FIB entries are maintained on the active SUP, and traffic is switched using the old FIB table for two minutes. After failure of the designated router, the new designated router is not allowed to update the supervisor (SUP) for two minutes while it is building its routing table.
4. After two minutes, the new Cisco Express Forwarding (CEF) table (CEF table of the new designated router) is downloaded to the SUP II, whether or not the routing protocol has completed its convergence.
5. As routing protocol neighbors have their adjacencies cleared, there may still be a forwarding outage on other devices after the switchover.

If the SRM and SUP IA/PFC/MSFC (1 or 2) begin to fail, the following is true:

- The designated router is failing.
- The new designated router brings up its VLAN interfaces.
- The existing Multilayer Switching (MLS) shortcuts are maintained on the SUP. Layer 3 traffic continues to be routed using the old shortcut.
- Any new flow that needs to be created is created by the new designated router immediately with these steps:
  - A packet is a candidate for the Layer 3 shortcut.
  - The packet is forwarded to the new designated router.
  - If the new designated router already has a route to destination, it routes the packet and the new shortcut is created on the SUP.
  - If the new designated router does not yet have a route to the destination (remember, the new designated router may still be busy computing the routing table), the packet is dropped.

The table describes the advantages and disadvantages of SRM.

Advantages	Disadvantages
Conserves IP addresses	Uses an old FIB image of the routing table even though the router that creates it is not on line anymore. There is a risk during the table-update-delay time to route packet to a nonvalid route.
Reduces routing protocol peering	Can be disruptive to the network, because the routing table needs to be calculated from the start on the new designated router.
Configuration much simpler; no risk of running unsupported mismatched configurations	

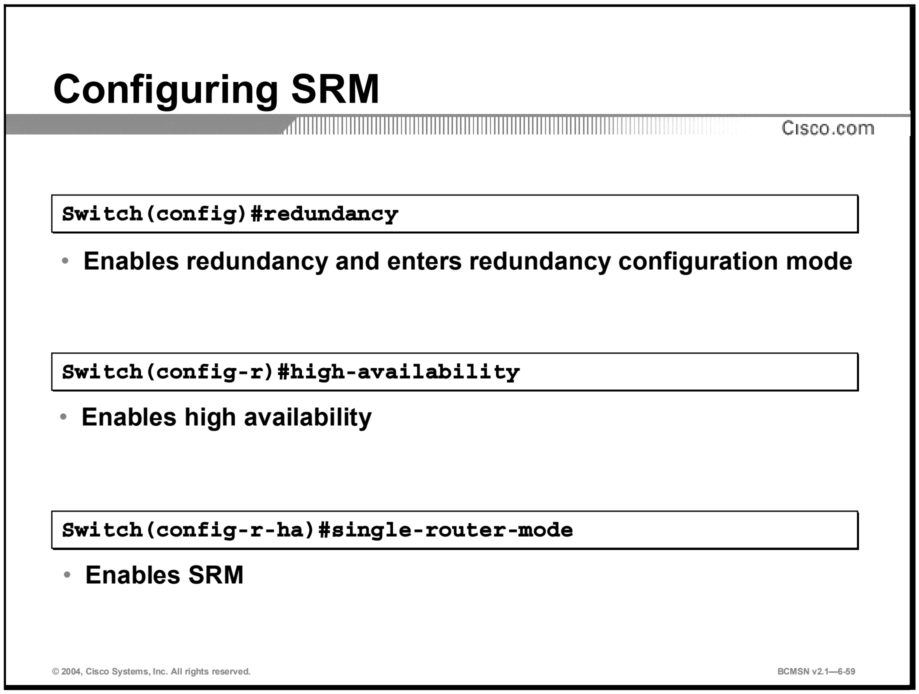

Configuring SRM

Cisco.com

```
Switch(config)#redundancy
```

- Enables redundancy and enters redundancy configuration mode

```
Switch(config-r)#high-availability
```

- Enables high availability

```
Switch(config-r-ha)#single-router-mode
```

- Enables SRM

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-59

Note Before going from Dual Router Mode (DRM) to SRM redundancy, Cisco recommends that you use the **copy running-config** command on the MSFCs to save the non-SRM configuration to boot flash memory. When going to SRM redundancy, the alternative configuration (the configuration following the **alt** keyword) is lost. Therefore, before enabling SRM redundancy, save the DRM configuration to boot flash memory by entering the following command on both MSFCs: **copy running-config bootflash:nosrm_dual_router_config**.

The following procedure assumes that the designated router is the MSFC2 in slot 1 and the nondesignated router is the MSFC2 in slot 2; the active Supervisor Engine is in slot 1 and the standby Supervisor Engine is in slot 2.

To configure SRM, follow these steps:

Step 1 Enter the **show version** command to ensure that both Supervisor Engines are running Supervisor Engine software Release 6.3(1) or later.

Step 2 Enter the **set system highavailability enable** command to enable high availability on the active Supervisor Engine. Enter the **show system highavailability** command to verify that high availability is enabled.

Step 3 If you have a console connection, enter the **switch console** command to access the designated router. If connected through a Telnet session, enter the **session mod** command to access the designated router.

Step 4 Copy the Cisco IOS Release 12.1(8a)E2 or later image to the boot flash of the designated and nondesignated routers.

Step 5 Set the boot image and configuration register on the designated and nondesignated routers to boot the new image on a reload.

For the designated router, enter **boot system flash bootflash:image_name** and ensure that this image is the first in the boot list.

Clear any existing **boot system** commands that appear in the running configuration (**show running-config**) using the **no** form of the **boot system** command.

For the nondesignated router, set the configuration register to autoboot by entering **config-register 0x102**.

Step 6 Enter the **reload** command to reload the designated and nondesignated routers.

Note If you already have SRM-capable IOS images loaded, you do not need to perform Step 6.

Step 7 Disable configuration synchronization (**config-sync**) on the designated router using the **no** form of the command. Enter the **write memory** command. This lets you have access to configuration mode on both designated and nondesignated routers.

Step 8 Enable SRM on the designated router first, and then enable SRM on the nondesignated router as follows:

```
Switch(config)#redundancy
Switch(config-r)#high-availability
Switch(config-r-ha)#single-router-mode
```

Step 9 Enter the **write memory** command on the designated router to ensure that the nondesignated router startup configuration has SRM enabled.

This display summarizes the configuration commands used on the designated router and nondesignated router to enable SRM redundancy:

```
Time   Designated Router                           Nondesignated
Router
  t0:    conf t->red->hi->no config-sync
  t1:                                              conf t->red->hi-
>no config-sync
  t2:    conf t->red->hi->single-router-mode
  t3:                                              conf t->red->hi-
>single-router-m
  t4:    write mem
  t5:                                              reload
```

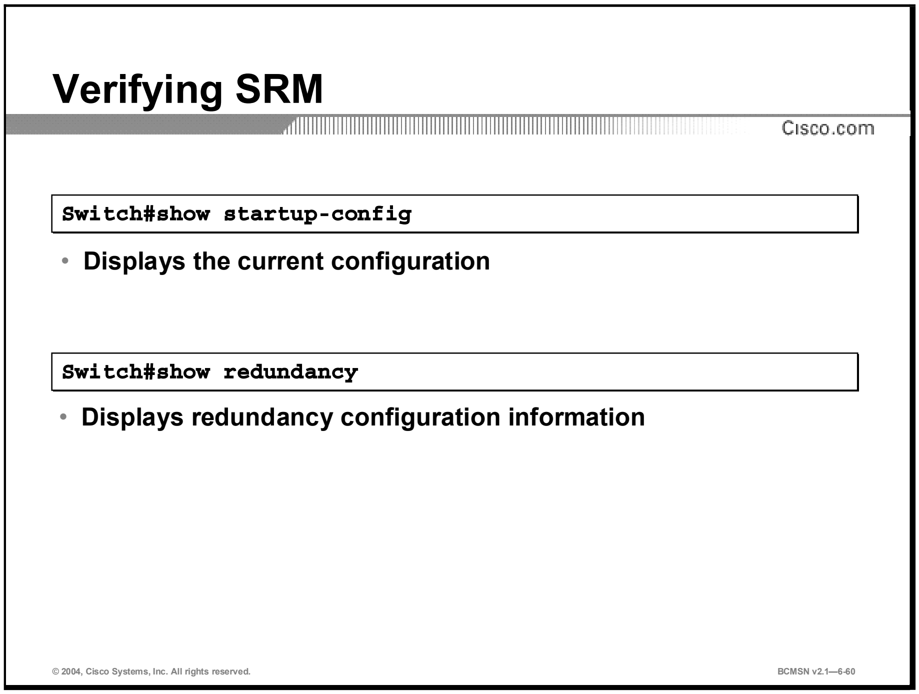

Verifying SRM

Cisco.com

```
Switch#show startup-config
```

- Displays the current configuration

```
Switch#show redundancy
```

- Displays redundancy configuration information

BCMSN v2.1—6-60

To verify the SRM configuration, complete the following steps:

Step 1 Enter the **show startup-config** command on the nondesignated router to ensure that the nondesignated router has the following configuration statements:

```
redundancy
high-availability
single-router-mode
```

Step 2 Enter the **show redundancy** command on the designated and nondesignated routers to ensure that both have this configuration statement:

```
Single Router Mode RuntimeStatus: enabled
```

Step 3 Enter the **reload** command to reload the nondesignated router. When asked whether the configuration should be saved, enter **no**.

Configuring and Verifying SLB

IOS server load balancing (SLB) intelligently load balances TCP/IP traffic across multiple servers. This topic explains how to configure and verify SLB.

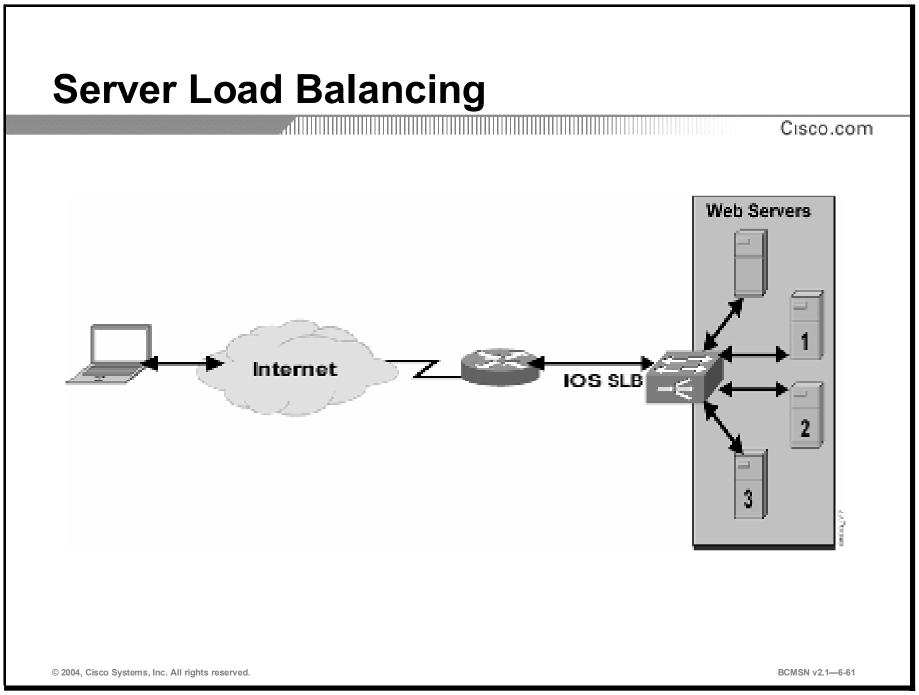

Server Load Balancing

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—6-61

IOS SLB intelligently load balances TCP/IP traffic across multiple servers, as shown in the figure. IOS SLB appears as one virtual server to the requesting clients. All traffic is directed toward a virtual IP address (virtual server) via Domain Name System (DNS). Those requests are distributed over a series of real IP addresses on servers (real servers). The definition of a virtual IP address is an address that is in DNS and that most likely has a domain name. A real IP address is physically located on a real server behind IOS SLB. IOS SLB provides these benefits:

- High performance is achieved by distributing client requests across a cluster of servers.
- Administration of server applications is easier. Clients know only about virtual servers; no administration is required for real server changes.
- Security of the real server is provided because its address is never announced to the external network. Users are familiar only with the virtual IP address. Additionally, filtering of unwanted traffic can be based on both IP address and IP port numbers.
- Ease of maintenance with no downtime is achieved by allowing physical (real) servers to be transparently placed in or out of service.

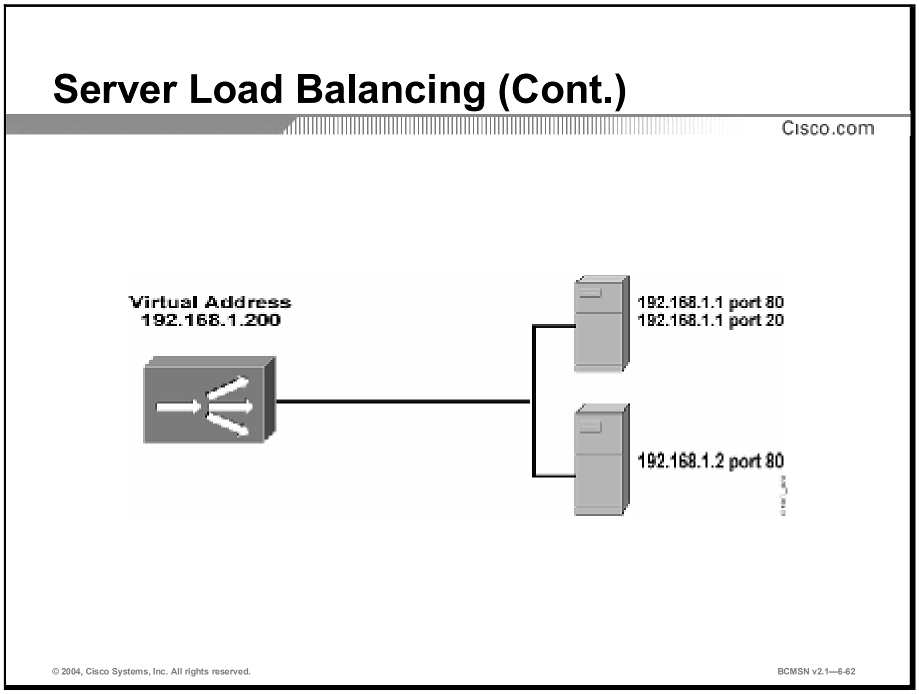

Server Load Balancing (Cont.)

Cisco.com

BCMSN v2.1—6-62

IOS SLB allows users to represent a group of network servers (a server farm) as a single server instance, balance the traffic to the servers, and limit traffic to individual servers. The single server instance that represents a server farm is referred to as a “virtual server”. The servers that comprise the server farm are referred to as “real servers”, as shown in the figure.

In this environment, clients are configured to connect to the IP address of the virtual server. When a client initiates a connection to the virtual server, the SLB function chooses a real server for the connection based on a configured load-balancing algorithm. Representing server farms as virtual servers facilitates scalability and availability for the user. The addition of new servers and the removal or failure of existing servers can occur at any time without affecting the availability of the virtual server.

IOS SLB can operate in one of two redirection modes: directed mode or dispatched mode. In directed mode, the virtual server can be assigned an IP address that is not known to any of the real servers. IOS SLB translates packets exchanged between a client and real server, translating the virtual server IP address to a real server address via Network Address Translation (NAT). In dispatched mode, the real servers know the virtual server address, and IOS SLB redirects packets to the real servers at the MAC layer.

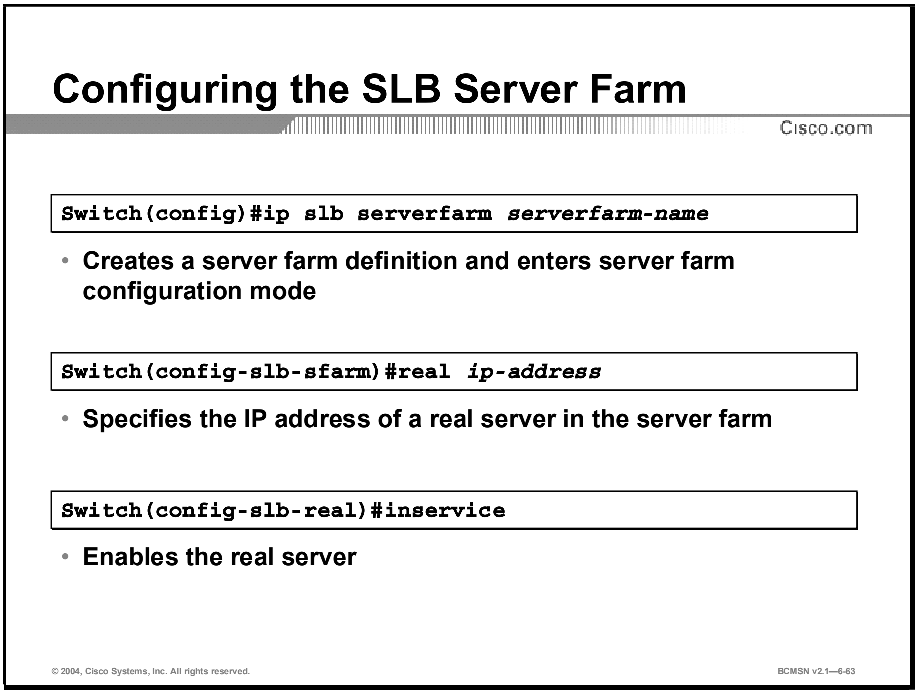

Configuring the SLB Server Farm

```
Switch(config)#ip slb serverfarm serverfarm-name
```

- **Creates a server farm definition and enters server farm configuration mode**

```
Switch(config-slb-sfarm)#real ip-address
```

- **Specifies the IP address of a real server in the server farm**

```
Switch(config-slb-real)#inservice
```

- **Enables the real server**

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—6-63

Configuring IOS SLB involves identifying server farms, configuring groups of real servers in server farms, and configuring the virtual servers that represent the real servers to the clients. The table lists the steps to configure a server farm.

Step	Action	Notes
1.	Create a server farm. `Switch(config)#ip slb serverfarm serverfarm-name`	Adds a server farm definition to the IOS SLB configuration and initiates server farm configuration mode.
2.	Identify a real server. `Switch(config-slb-sfarm)#real ip-address`	Identifies a real server as a member of a server farm and initiates real server configuration mode.
3.	Enable the real server. `Switch(config-slb-real)#inservice`	Enables the real server for use by IOS SLB.

Example: Configuring a Server Farm

These commands configure the server farm named “PUBLIC” and associate the three real servers:

```
Switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip slb serverfarm PUBLIC
Switch(config-slb-sfarm)#real 10.1.1.1
Switch(config-slb-real)#inservice
Switch(config-slb-real)#exit
Switch(config-slb-sfarm)#real 10.1.1.2
Switch(config-slb-real)#inservice
Switch(config-slb-real)#exit
Switch(config-slb-sfarm)#real 10.1.1.3
Switch(config-slb-real)#inservice
Switch(config-slb-real)#end
```

These commands configure the server farm named “RESTRICTED” and associate the two real servers:

```
Switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#ip slb serverfarm RESTRICTED
Switch(config-slb-sfarm)#real 10.1.1.20
Switch(config-slb-real)#inservice
Switch(config-slb-real)#exit
Switch(config-slb-sfarm)#real 10.1.1.21
Switch(config-slb-real)#inservice
Switch(config-slb-real)#end
Switch#
```

This command displays the status of server farms PUBLIC and RESTRICTED, the associated real servers, and their status:

```
Switch#show ip slb real

real                    farm name        weight   state
cons
--------------------------------------------------------------
-------
10.1.1.1                 PUBLIC           8       OPERATIONAL
0
10.1.1.2                 PUBLIC           8       OPERATIONAL
0
10.1.1.3                 PUBLIC           8       OPERATIONAL
0
10.1.1.20                RESTRICTED       8       OPERATIONAL
0
10.1.1.21                RESTRICTED       8       OPERATIONAL
0
```

This command displays the configuration and status of server farms PUBLIC and RESTRICTED:

```
Switch# show ip slb serverfarm

server farm      predictor    nat   reals   bind id
---------------------------------------------------
PUBLIC           ROUNDROBIN   none  3       0
RESTRICTED       ROUNDROBIN   none  2       0
```

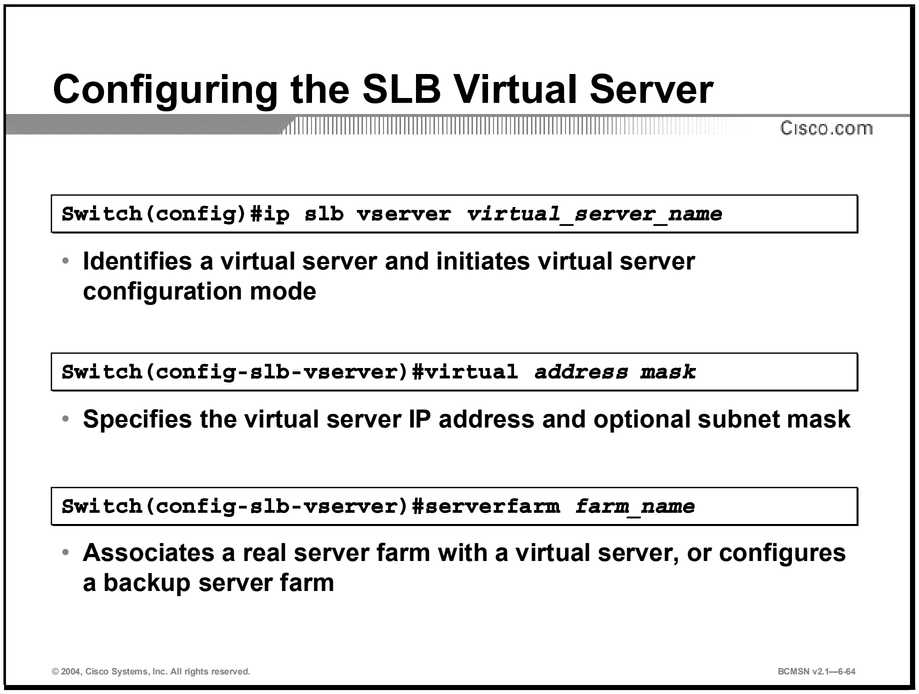

Configuring the SLB Virtual Server

```
Switch(config)#ip slb vserver virtual_server_name
```

- Identifies a virtual server and initiates virtual server configuration mode

```
Switch(config-slb-vserver)#virtual address mask
```

- Specifies the virtual server IP address and optional subnet mask

```
Switch(config-slb-vserver)#serverfarm farm_name
```

- Associates a real server farm with a virtual server, or configures a backup server farm

BCMSN v2.1—6-64

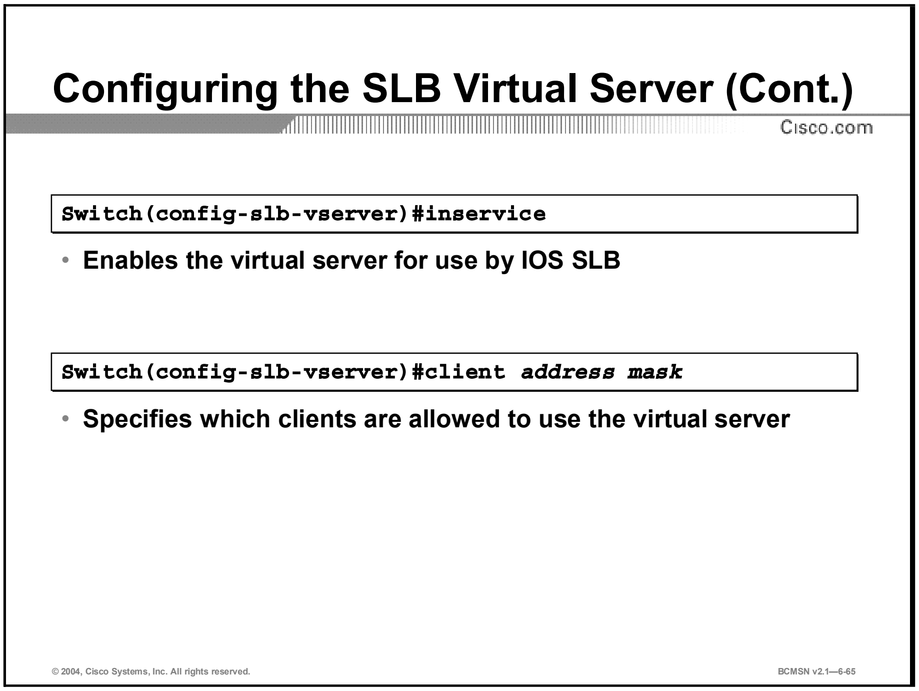

Configuring the SLB Virtual Server (Cont.)

```
Switch(config-slb-vserver)#inservice
```

- Enables the virtual server for use by IOS SLB

```
Switch(config-slb-vserver)#client address mask
```

- Specifies which clients are allowed to use the virtual server

BCMSN v2.1—6-65

The table lists the steps to configure a virtual server.

Step	Action	Notes
1.	Enter virtual server configuration mode. Switch(config)#**ip slb vserver** *virtual_server-name*	Identifies a virtual server and initiates virtual server configuration mode.
2.	Specify the IP address of the virtual server. Switch(config-slb-vserver)#**virtual** *ip-address* [*network-mask*] {**tcp** | **udp**} [*port-number* | **wsp** | **wsp-wtp** | **wsp-wtls** | **wsp-wtp-wtls**] [**service** *service-name*]	Specifies the virtual server IP address and optional subnet mask.
3.	Associate a server farm with the virtual server. Switch(config-slb-vserver)# **serverfarm** *primary-serverfarm-name* [**backup** *backup-serverfarm-name* [**sticky**]]	Associates a real server farm with a virtual server, or configures a backup server farm.
4.	Enable the virtual server. Switch(config-slb-vserver)# **inservice**	Enables the virtual server for use by IOS SLB.
5.	Specify the clients allowed to access the virtual server. Switch(config-slb-vserver)# **client** *ip-address network-mask*	Specifies which clients are allowed to use the virtual server.

Example: Configuring Virtual Servers

These commands configure the virtual servers PUBLIC_HTTP and RESTRICTED_HTTP, with the latter being restricted to access by clients in the network 10.4.4.0:

```
Switch#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip slb vserver PUBLIC_HTTP
Switch(config-slb-vserver)# virtual 10.0.0.1 tcp www
Switch(config-slb-vserver)# serverfarm PUBLIC
Switch(config-slb-vserver)# inservice
Switch(config-slb-vserver)#
.
(Information Deleted)
.
index = 1
Switch(config-slb-vserver)#exit
Switch(config)#ip slb vserver RESTRICTED_HTTP
Switch(config-slb-vserver)#virtual 10.0.0.2 tcp www
Switch(config-slb-vserver)#client 10.4.4.0 255.255.255.0
Switch(config-slb-vserver)#serverfarm RESTRICTED
Switch(config-slb-vserver)#inservice
Switch(config-slb-vserver)#
src = 0 – 0
.
(Information Deleted)
.
```

```
index = 1
Switch(config-slb-vserver)#end
Switch#
```

This command verifies the configuration of the virtual servers PUBLIC_HTTP and RESTRICTED_HTTP:

```
Switch#show ip slb vserver

slb vserver      prot  virtual               state
cons
--------------------------------------------------------------
-----
PUBLIC_HTTP      TCP   10.0.0.1:80           OPERATIONAL     0
RESTRICTED_HTTP  TCP   10.0.0.2:80           OPERATIONAL     0
```

This command verifies the restricted client access and status:

```
Switch#show ip slb cons

vserver         prot client                real
state     nat
--------------------------------------------------------------
-------
RESTRICTED_HTTP TCP  10.4.4.0:80           10.1.1.20
CLOSING   none
```

This command displays detailed information about the restricted client access status:

```
Switch#show ip slb conns client 10.4.4.0 detail

VSTEST_UDP, client = 10.4.4.0:80
  state = CLOSING, real = 10.1.1.20, nat = none
  v_ip = 10.0.0.2:80, TCP, service = NONE
  client_syns = 0, sticky = FALSE, flows attached = 0
```

This command displays detailed information about the IOS SLB network status:

```
Switch#show ip slb stats

Pkts via normal switching:  0
Pkts via special switching: 6
Connections Created:        1
Connections Established:    1
Connections Destroyed:      0
Connections Reassigned:     0
Zombie Count:               0
Connections Reused:         0
```

Summary

This topic summarizes the key points discussed in this lesson.

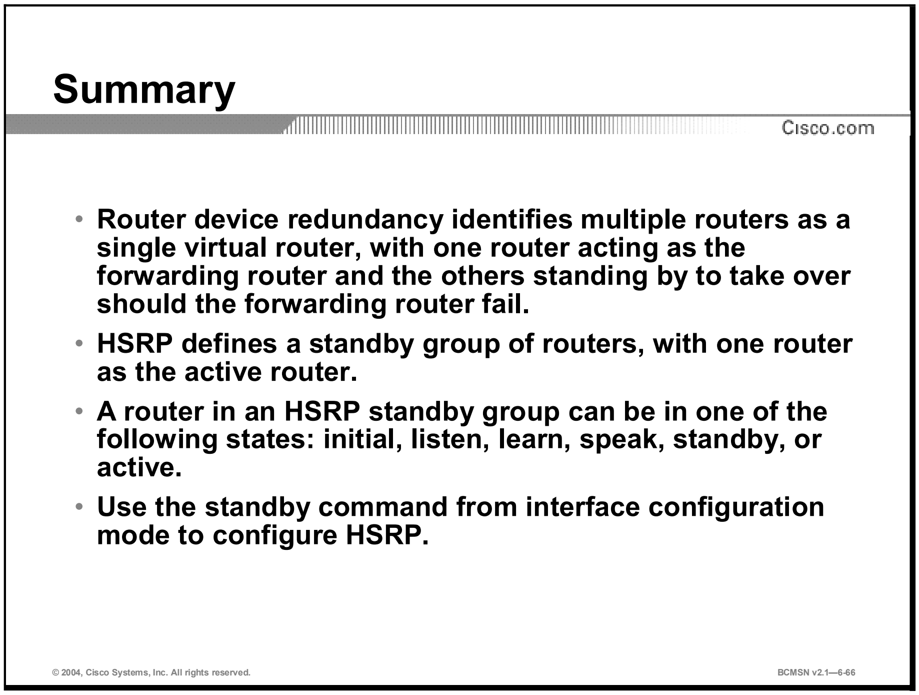

Summary

- Router device redundancy identifies multiple routers as a single virtual router, with one router acting as the forwarding router and the others standing by to take over should the forwarding router fail.
- HSRP defines a standby group of routers, with one router as the active router.
- A router in an HSRP standby group can be in one of the following states: initial, listen, learn, speak, standby, or active.
- Use the standby command from interface configuration mode to configure HSRP.

BCMSN v2.1—6-66

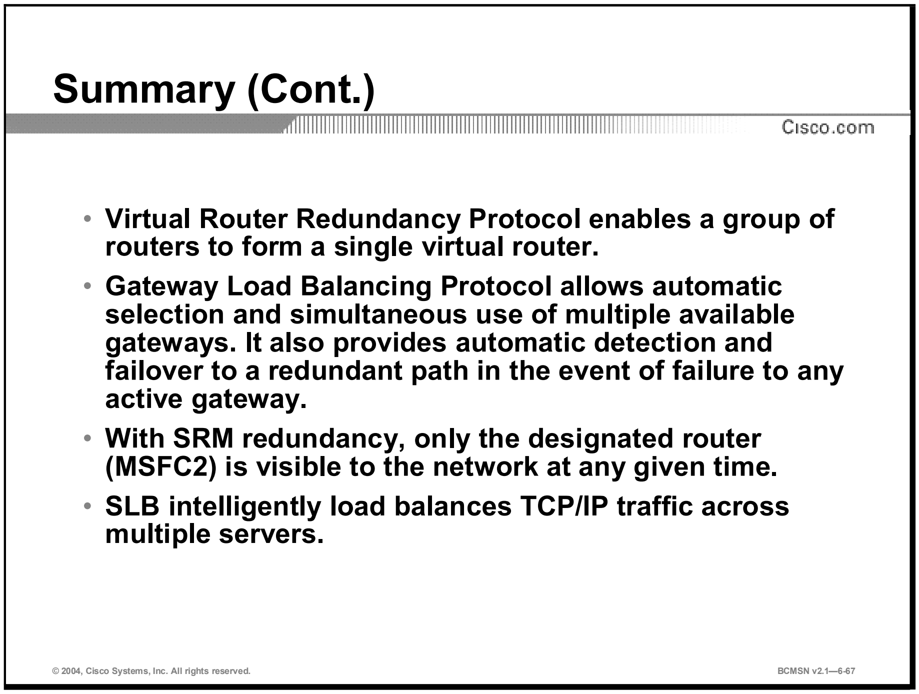

Summary (Cont.)

- Virtual Router Redundancy Protocol enables a group of routers to form a single virtual router.
- Gateway Load Balancing Protocol allows automatic selection and simultaneous use of multiple available gateways. It also provides automatic detection and failover to a redundant path in the event of failure to any active gateway.
- With SRM redundancy, only the designated router (MSFC2) is visible to the network at any given time.
- SLB intelligently load balances TCP/IP traffic across multiple servers.

BCMSN v2.1—6-67

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Next Step

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 6-1: Improving Availability on Multilayer Switched Networks

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which protocol does the Cisco IOS software use to help hosts that have no knowledge of routing determine the media addresses of hosts on other networks or subnets?

A) ARP

B) IRDP

C) RARP

D) proxy ARP

Q2) To which router in an HSRP standby group do end-user stations send packets?

A) the active router

B) the virtual router

C) the standby router

D) the primary router

Q3) In which three HSRP states do routers send hello messages? (Choose three.)

A) initial

B) learn

C) listen

D) speak

E) active

F) standby

Q4) Which command correctly forces a takeover as the active router if the conditions are met?

A) Switch(config)#**standby preempt**

B) Switch(config-if)#**standby preempt**

C) Switch(config)#**standby 15 preempt**

D) Switch(config-if)#**standby 15 preempt**

Q5) What are two drawbacks to dynamic discovery protocols? (Choose two.)

A) slow failover

B) excess traffic

C) slow convergence

D) router configuration overhead

E) LAN client configuration overhead

Q6) Which method of load balancing is used by GLBP in default mode?

- A) fifo
- B) round robin
- C) priority queuing
- D) weighted round robin

Q7) In what state are the nondesignated router interfaces when SRM is configured?

- A) visible
- B) link-up
- C) invisible
- D) link-down

Q8) Which command correctly specifies the IP address of a server to be a member of an SLB server farm?

- A) Switch(config)#**real 10.1.1.20**
- B) Switch(config-slb-sfarm)#**real 10.1.1.20**
- C) Switch(config)#**ip slb serverfarm 10.1.1.20**
- D) Switch(config-slb-sfarm)#**ip slb serverfarm 10.1.1.20**

Quiz Answer Key

Q1) D

Relates to: Understanding How Router Redundancy Works

Q2) B

Relates to: HSRP Operations

Q3) D, E, F

Relates to: HSRP States

Q4) D

Relates to: Configuring and Verifying HSRP

Q5) A, E

Relates to: Introducing VRRP

Q6) B

Relates to: Introducing GLBP

Q7) D

Relates to: Configuring and Verifying SRM

Q8) B

Relates to: Configuring and Verifying SLB

Lesson Assessments

Overview

Use the lesson assessments here to test what you learned in this module. The correct answers and solutions are found in the Lesson Assessment Answer Key.

Outline

This section includes these assessments:

- Quiz 6-1: Implementing Module Redundancy in a Multilayer Switched Network
- Quiz 6-2: Implementing Router Redundancy in a Multilayer Switched Network

Quiz 6-1: Implementing Module Redundancy in a Multilayer Switched Network

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Explain the types of redundancy in a multilayer switched network, including hardware and software redundancy
- Implement redundant supervisor modules in Catalyst switches
- Implement redundant supervisor uplink modules in Catalyst switches
- Implement redundant power supplies

Quiz

Answer these questions:

Q1) Which two advantages can apply to attempting to achieve reliability through redundant topology design? (Choose two.)

A) increased device efficiency

B) decreased infrastructure cost

C) increased aggregate performance

D) decreased administrative complexity

E) reduced impact of non-hardware-failure mechanisms

Q2) What is the switchover time for RPR?

A) 30 to 60 seconds

B) 2 to 4 seconds

C) 1 to 2 minutes

D) 2 to 4 minutes

Q3) What benefit is gained from providing multiple uplink connections as alternate hardware paths on a Catalyst switch?

A) increased security

B) increased efficiency

C) increased availability

D) increased performance

Q4) Which command correctly turns off power to a specific module and then turns it back on in a Catalyst switch?

A) **power cycle module** *slot*

B) **power enable module** *slot*

C) **power restart module** *slot*

D) **no power enable module** *slot*

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 75 percent or better.

Quiz 6-2: Implementing Router Redundancy in a Multilayer Switched Network

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge on how to:

- Explain how router redundancy operates
- Describe how HSRP operates
- Describe the HSRP states
- Configure and verify HSRP
- Describe VRRP
- Describe GLBP
- Describe, configure, and verify SRM
- Describe, configure, and verify SLB

Quiz

Answer these questions:

Q1) Which protocol is an extension to ICMP that provides a mechanism for routers to advertise useful default routes?

A) ARP

B) IRDP

C) RARP

D) proxy ARP

Q2) Which type of HSRP message indicates that the router is running and is capable of becoming either the active or standby router?

A) coup

B) hello

C) resign

D) active

Q3) Which HSRP state indicates that the router knows the virtual IP address, but is neither the active router nor the standby router?

A) initial

B) learn

C) listen

D) speak

Q4) Which command correctly configures the HSRP hello interval time to 5 seconds?

A) Switch(config)#**standby hellotime 5**

B) Switch(config)#**standby 15 timers 5 15**

C) Switch(config-if)#**standby hellotime 5**

D) Switch(config-if)#**standby 15 timers 5 15**

Q5) Assume a scenario in which router A is the initial VRRP master virtual router and routers B and C act as backup virtual routers. If router B takes over as the master virtual router when router A fails, which router is the master virtual router when router A is once again available?

A) router A

B) router B

C) router C

D) the router assigned the highest priority value

Q6) How many members can be assigned to a GLBP group?

A) two

B) four

C) three

D) eight

Q7) Which command correctly enables high availability for SRM on a Catalyst switch?

A) Switch(config)#**redundancy**

B) Switch(config-r)#**high-availability**

C) Switch(config-r-ha)#**single-router-mode**

D) Switch(config-r)#**single-router-mode high-availability**

Q8) Which command correctly assigns a virtual server to an SLB server farm?

A) Switch(config)#**serverfarm PUBLIC**

B) Switch(config-slb-vserver)#**virtual 10.0.0.1**

C) Switch(config-slb-vserver)#**serverfarm PUBLIC**

D) Switch(config-slb-vserver)#**virtual 10.0.0.1 tcp www PUBLIC**

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Lesson Assessment Answer Key

Quiz 6-1: Implementing Module Redundancy in a Multilayer Switched Network

Q1) C, E

Q2) D

Q3) C

Q4) A

Quiz 6-2: Implementing Router Redundancy in a Multilayer Switched Network

Q1) B

Q2) B

Q3) C

Q4) D

Q5) A

Q6) B

Q7) B

Q8) C

Module 7

Examining Cisco AVVID Services and Applications

Overview

IP multicast enables the scalable and efficient distribution of data, voice, and video streams to hundreds, thousands, even millions of users. Cisco IOS multicast enables corporate communications, video conferencing, e-learning, Internet broadcast, Hoot and Holler, and streaming media applications.

Upon completing this module, you will be able to:

- Describe the operation of IP multicast in a multilayer switched network
- Configure, monitor, and troubleshoot IP multicast in a multilayer switched network
- Describe the Cisco IP telephony solution and explain its role in a multilayer switched network

Outline

The module contains these components:

- Examining IP Multicast in a Multilayer Switched Network
- Configuring IP Multicast in a Multilayer Switched Network
- Introducing Cisco IP Telephony
- Lesson Assessments

Examining IP Multicast in a Multilayer Switched Network

Overview

Most campus networks today support intranet applications that operate between one sender and one receiver. This is referred to as "unicast." In the emerging campus network, there is a demand for intranet and multimedia applications in which one sender transmits to a group of receivers simultaneously. These applications include transmitting a corporate message to employees, video and audio broadcasting, interactive video distance learning, transmitting data from a centralized data warehouse to multiple departments, communication of stock quotes to brokers, and collaborative computing.

Relevance

Among the many capabilities of Cisco IOS software are its IP multicast technologies. IP multicast enables the massively scalable and efficient distribution of data, voice, and video streams to hundreds, thousands, even millions of users. IOS multicast enables corporate communications, video conferencing, distance learning, Internet broadcast, Hoot and Holler, and streaming media applications.

Objectives

Upon completing this lesson, you will be able to:

- Explain how IP multicast operates on a multilayer switched network
- Describe the operation of Reverse Path Forwarding
- Describe the operation and issues of PIM dense mode and PIM sparse mode
- Describe the operation of CGMP and of IGMP versions 1, 2, 3, and IGMP v3lite, and explain the impact of these protocols on a Catalyst switch
- Define IGMP snooping

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- IP Multicast Operation Overview
- Reverse Path Forwarding
- PIM Dense Mode and PIM Sparse Mode
- IGMP and CGMP Operation
- IGMP Snooping
- Summary
- Quiz

IP Multicast Operation Overview

This topic discusses IP multicast operations.

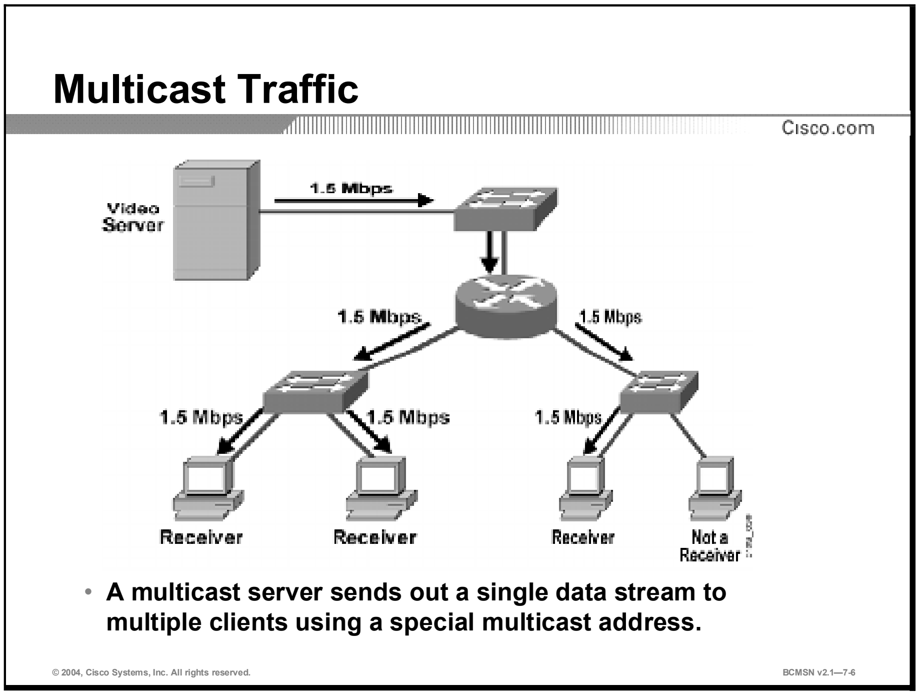

Multicast Traffic

Cisco.com

- **A multicast server sends out a single data stream to multiple clients using a special multicast address.**

BCMSN v2.1—7-6

Multimedia applications offer the integration of sound, graphics, animation, text, and video. These types of applications have become increasingly popular in the network as ways of conducting business become more complex. However, blending several media into a single medium, such as Ethernet, and sending the combined media over a campus data network is a complicated process. This process brings with it the potential for high bandwidth consumption on the network. IP multicast is the transmission of an IP data frame to a host group that is defined by a single flow or a single IP address.

Multimedia traffic can work its way through the network in one of several ways.

- Unicast
- Broadcast
- Multicast

Each one of the above methods of transmission has a different effect on network bandwidth.

With a unicast design, an application sends one copy of each packet to every client unicast address. Unicast transmission has significant scaling restrictions. If the group is large, the same information has to be carried multiple times, even on shared links.

In a broadcast design, an application sends only one copy of each packet using a broadcast address. However, if this technique is used, broadcasts either must be stopped at the broadcast domain boundary with a Layer 3 device or transmitted to all devices in the campus network. Broadcasting a packet to all devices can be inefficient if only a small group in the network actually needs to see the packet.

Broadcast multimedia is dispersed throughout the network just like normal broadcast traffic. As with normal broadcasts, every client must process the broadcast multimedia data frame. However, unlike standard broadcast frames, which are generally small, multimedia broadcasts can reach as high as 7 Mbps or more of data. Even if an end station is not using a multimedia application, the device still processes the broadcast traffic. This requirement can use up most, if not all, of the allocated bandwidth for each device. For this reason, the broadcast multimedia method is rarely implemented.

The most efficient solution is one in which a multimedia server sends one copy of each packet, addressing each packet to a special address. Unlike the unicast environment, a multicast server sends out a single data stream to multiple clients. Unlike the broadcast environment, the client device decides whether to listen to the multicast address. Multicasting saves bandwidth and controls network traffic by forcing the network to replicate packets only when necessary. By eliminating traffic redundancy, multicasting reduces network and host processing.

In the figure, the video server transmits a single video stream for each multicast group. A multicast, or host, group is defined as a set of host devices listening to a specific multicast address. The video stream is then replicated as required by the multicast routers and switches that are in the flow path. This technique allows an arbitrary number of clients to subscribe to the multicast address and receive the broadcast.

In the multicast scenario, only 1.5 Mbps of server-to-network bandwidth is utilized, leaving the remaining bandwidth free for other uses. Within the network, the multicast transmission offers similar efficiency, consuming only 1/*n*th of the bandwidth, where *n* equals the number of users receiving the video stream.

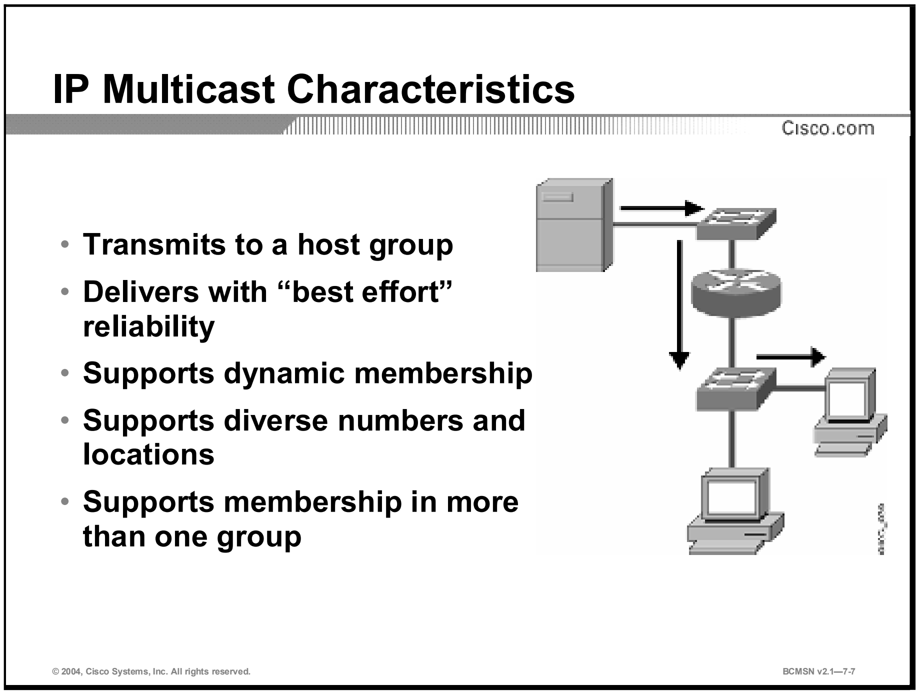

IP Multicast Characteristics

Cisco.com

- Transmits to a host group
- Delivers with “best effort” reliability
- Supports dynamic membership
- Supports diverse numbers and locations
- Supports membership in more than one group

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—7-7

IP multicast is the transmission of an IP data frame to a host group that is defined by a single IP address. IP multicast reduces network traffic by simultaneously delivering a single stream of information to multiple recipients. IP multicasting has these characteristics:

- Facilitates transmission of an IP datagram to a host group consisting of zero or more hosts identified by a single IP destination address
- Delivers a multicast datagram to all members of the destination host group with the same best-effort reliability as regular unicast IP datagrams
- Supports dynamic membership of a host group
- Supports all host groups regardless of the location or number of members
- Supports the membership of a single host in one or more multicast groups
- Upholds multiple data streams at the application level for a single group address
- Supports a single group address for multiple applications on a host

In IP multicasting, the variability in delivery time is limited to the differences in end-to-end network delay along the complete server-to-client path. In a unicast scenario, the server sequences through transmission of multiple copies of the data, so variability in delivery time is large, especially for large transmissions or large distribution lists.

Another unique feature of multicast is that the server does not know the unicast network address of any particular recipient of the transmission. All recipients share the same multicast network address and, therefore, can join a multicast group while maintaining anonymity.

Multicast traffic is handled at the transport layer using the User Datagram Protocol (UDP). Unlike the Transmission Control Protocol (TCP), UDP adds no reliability, flow control, or error recovery functions to IP. Because of the simplicity of UDP, data packet headers contain fewer bytes and consume less network overhead than TCP. Therefore, reliability in multicast must be managed at the receiving end or by other means.

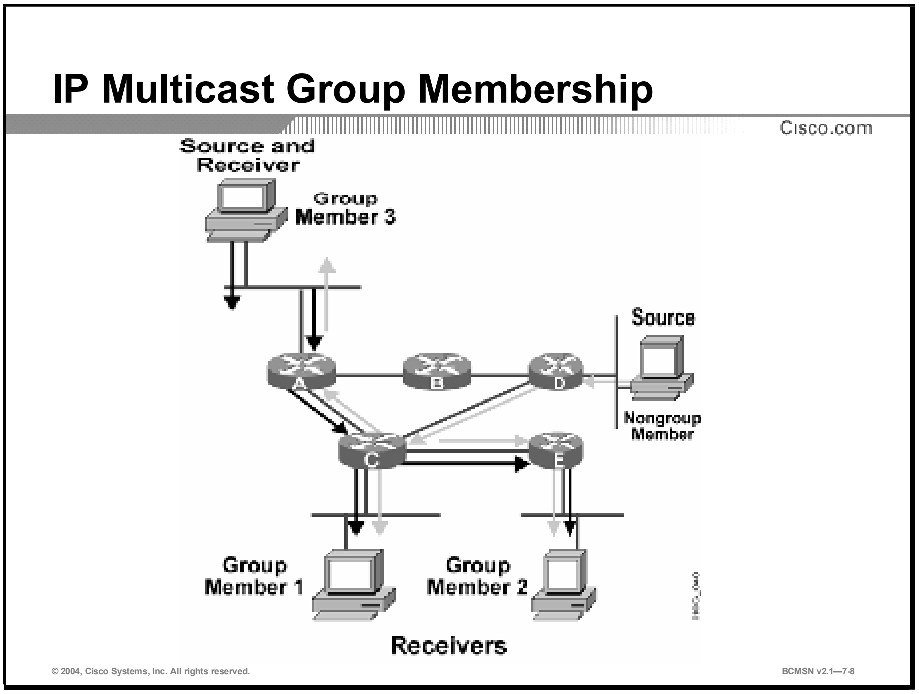

IP Multicast Group Membership

BCMSN v2.1—7-8

IP multicast relies on the concept of a virtual group address. In normal TCP/IP routing, a packet is routed from a source address to a destination address, traversing the IP network on a hop-by-hop basis. In IP multicast, the packet destination address is not assigned to a single destination. Instead, receivers join a group and, when they join, packets addressed to that group begin flowing to them. All members of the group receive the packet; a host must be a member of the group to receive the packet. Multicast sources or senders to the group do not need to be members of that group, and group members are not required to send to the group.

In the figure, packets sent by group member 3 (represented by the dark arrows) are received by group members 1 and 2, but not by the nongroup member. The nonmember host sends packets to the multicast group (represented by the pale arrows), which are received by all three group members. Group members 1 and 2 are not sending any multicast packets.

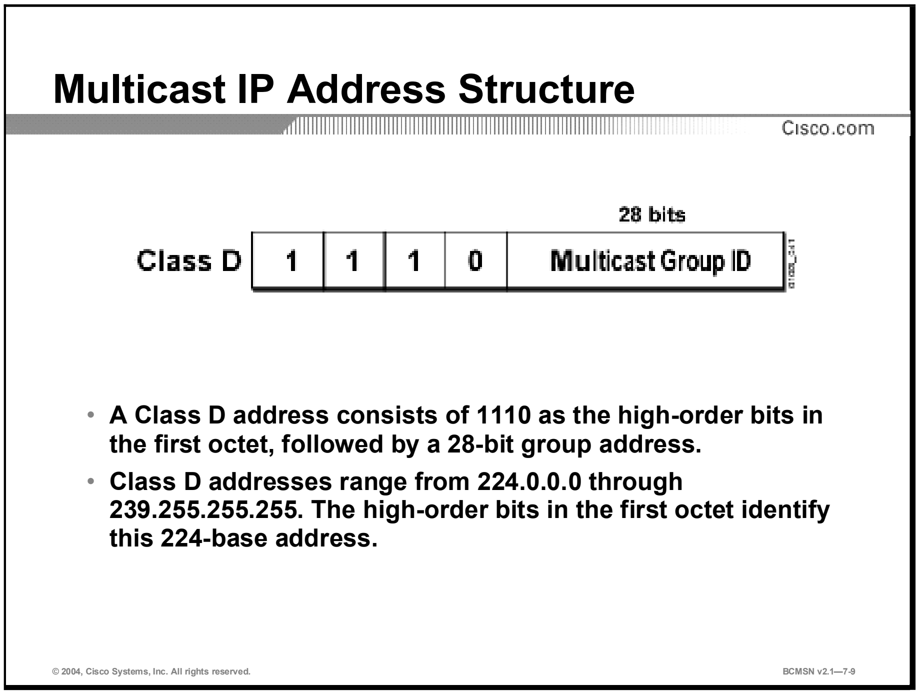

Multicast IP Address Structure

Cisco.com

28 bits

Class D	1	1	1	0	Multicast Group ID

- A Class D address consists of 1110 as the high-order bits in the first octet, followed by a 28-bit group address.
- Class D addresses range from 224.0.0.0 through 239.255.255.255. The high-order bits in the first octet identify this 224-base address.

BCMSN v2.1—7-9

The range of IP addresses is divided into classes based on the high-order bits of a 32-bit IP address. IP multicast uses Class D addresses. A Class D address consists of 1110 as the high-order bits in the first octet, followed by a 28-bit group address. Unlike Class A, B, and C IP addresses, the last 28 bits of a Class D address are unstructured.

These remaining 28 bits of the IP address identify the multicast group ID. This multicast group ID is a single address typically written as decimal numbers in the range 224.0.0.0 through 239.255.255.255. The high-order bits in the first octet identify this 224-base address.

Multicast addresses may be dynamically or statically allocated. Dynamic multicast addressing provides applications with a group address on demand. Because dynamic multicast addresses have a specific lifetime, applications must request this type of address only for as long as the address is needed.

Statically allocated addresses are reserved for specific protocols that require well-known addresses. The Internet Assigned Numbers Authority (IANA) assigns these well-known addresses. These addresses are called permanent host groups and are similar in concept to the well-known TCP and UDP port numbers.

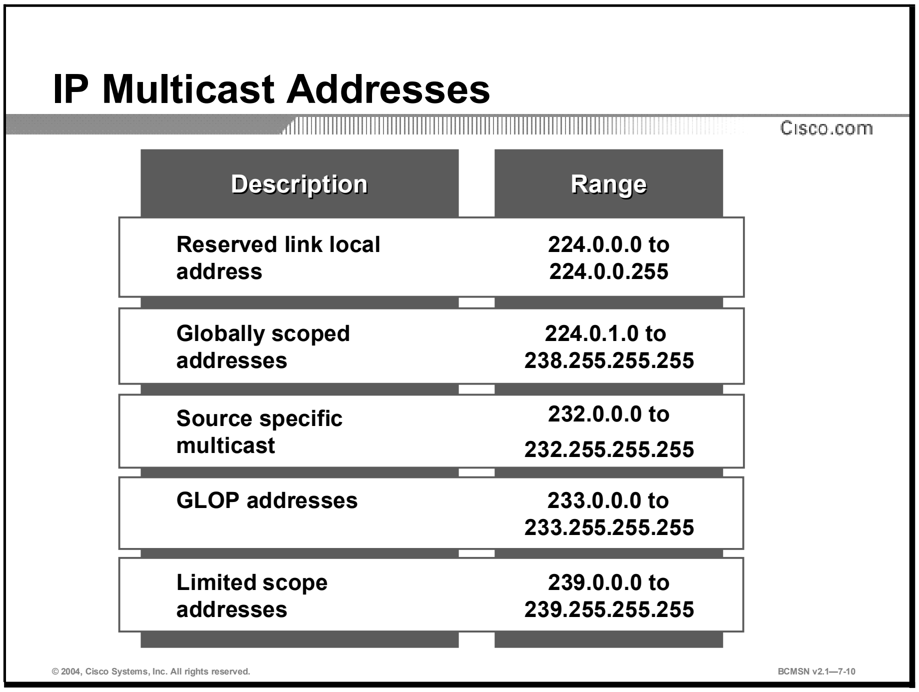

IP Multicast Addresses

Cisco.com

Description	Range
Reserved link local address	224.0.0.0 to 224.0.0.255
Globally scoped addresses	224.0.1.0 to 238.255.255.255
Source specific multicast	232.0.0.0 to 232.255.255.255
GLOP addresses	233.0.0.0 to 233.255.255.255
Limited scope addresses	239.0.0.0 to 239.255.255.255

BCMSN v2.1—7-10

IP multicast addresses specify a set of IP hosts that have joined a group and are interested in receiving multicast traffic designated for that particular group. IPv4 multicast address conventions are described in the figure.

The IANA controls the assignment of IP multicast addresses. IANA has assigned the IPv4 Class D address space to be used for IP multicast. Therefore, all IP multicast group addresses fall in the range from 224.0.0.0 through 239.255.255.255. The Class D address range is used only for the group address or destination address of IP multicast traffic. The source address for multicast datagrams is always the unicast source address.

Reserved Link-Local Addresses

The IANA has reserved addresses in the range 224.0.0.0 to 224.0.0.255 to be used by network protocols on a local network segment. A router should never forward packets with these addresses. Packets with link-local destination addresses are typically sent with a Time to Live (TTL) value of 1 and are not forwarded by a router. Network protocols use these addresses for automatic router discovery and to communicate important routing information. For example, Open Shortest Path First (OSPF) Protocol uses the IP addresses 224.0.0.5 and 224.0.0.6 to exchange link-state information.

Address 224.0.0.1 identifies the all-hosts group. Every multicast-capable host must join this group at the start. If a **ping** command is issued using this address, all multicast-capable hosts on the network must answer the ping request.

Address 224.0.0.2 identifies the all-routers group. Multicast routers must join that group on all multicast-capable interfaces.

Globally Scoped Addresses

Multicast addresses in the range from 224.0.1.0 through 238.255.255.255 are called “globally scoped addresses.” Some of these addresses have been reserved for use by multicast applications through IANA. For example, IP address 224.0.1.1 has been reserved for Network Time Protocol (NTP).

Source Specific Multicast Addresses

Addresses in the 232.0.0.0 to 232.255.255.255 range are reserved for Source Specific Multicast (SSM). SSM is an extension of Protocol Independent Multicast (PIM), which allows for an efficient data delivery mechanism in one-to-many communications.

GLOP Addresses

RFC 2770, *GLOP Addressing in 233/8*, proposes that the 233.0.0.0 to 233.255.255.255 address range be reserved for statically defined addresses by organizations that already have an Autonomous System (AS) number reserved. This practice is called “GLOP addressing.” The AS number of the domain is embedded into the second and third octets of the 233.0.0.0 to 233.255.255.255 address range. For example, the AS 62010 is written in hexadecimal format as “F23A.” Separating the two octets F2 and 3A results in 242 and 58 in decimal format. These values result in a subnet of 233.242.58.0/24, which are globally reserved for AS 62010 to use.

Limited Scope Addresses

Addresses in the 239.0.0.0 to 239.255.255.255 range are called “limited scope addresses” or “administratively scoped addresses.” These addresses are described in RFC 2365, *Administratively Scoped IP Multicast*, as being constrained to a local group or organization. Companies, universities, or other organizations can use limited scope addresses to have local multicast applications that will not be forwarded outside their domain. Typically, routers are configured with filters to prevent multicast traffic in this address range from flowing outside of an AS or any user-defined domain. Within an autonomous system or domain, the limited scope address range can be further subdivided so that local multicast boundaries can be defined. This subdivision is called “address scoping” and allows for address reuse between these smaller domains.

Reverse Path Forwarding

Multicast-capable routers create distribution trees that control the path that IP multicast traffic takes through the network. Forwarding multicast traffic away from the source, rather than to the receiver, is called Reverse Path Forwarding (RPF). This topic explains how IP multicast traffic traverses distribution trees.

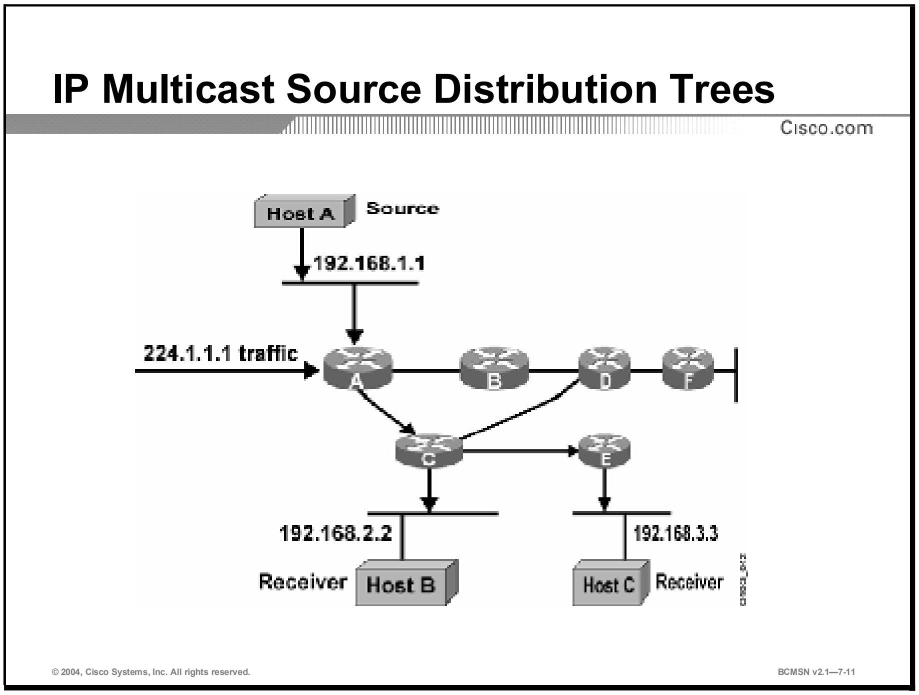

IP Multicast Source Distribution Trees

Cisco.com

BCMSN v2.1—7-11

Multicast-capable routers create distribution trees that control the path that IP multicast traffic takes through the network in order to deliver traffic to all receivers. The two basic types of multicast distribution trees are source trees and shared trees.

The simplest form of a multicast distribution tree is a source tree, with its root at the source and branches forming a spanning tree through the network to the receivers. Because this tree uses the shortest path through the network, it is also referred to as a “shortest path tree (SPT).”

The figure shows an example of an SPT for group 224.1.1.1 rooted at the source (host A) and connecting two receivers (hosts B and C).

The special notation of (S, G), pronounced “S comma G,” enumerates an SPT where S is the IP address of the source and G is the multicast group address. Using this notation, the SPT for the example shown in the figure would be (192.168.1.1, 224.1.1.1).

The (S, G) notation implies that a separate SPT exists for each individual source sending to each group. For example, if host B is also sending traffic to group 224.1.1.1 and hosts A and C are receivers, a separate (S, G) SPT would exist with a notation of (192.168.2.2, 224.1.1.1).

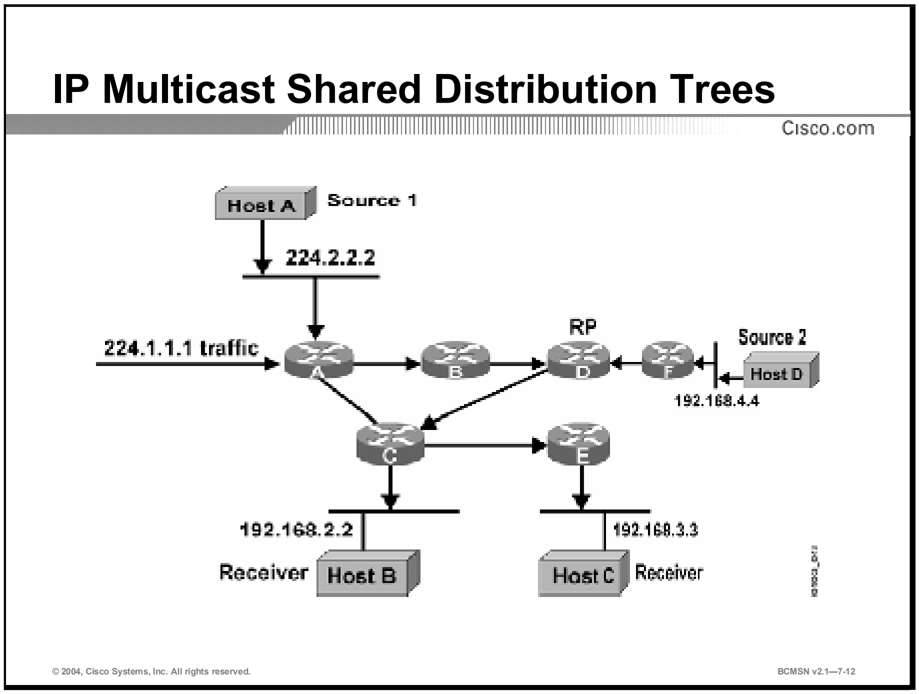

IP Multicast Shared Distribution Trees

Cisco.com

BCMSN v2.1—7-12

Unlike source trees that have their root at the source, shared trees use a single common root placed at some chosen point in the network. This shared root is called a "rendezvous point (RP)." The figure shows a shared unidirectional tree for group 224.2.2.2, with the root located at router D. Source traffic is sent toward the RP on a source tree. The traffic is then forwarded down the shared tree from the RP to reach all of the receivers, unless the receiver is located between the source and the RP, in which case it will be serviced directly.

In this example, multicast traffic from the sources (hosts A and D) travels to the root (router D) and then down the shared tree to the two receivers (hosts B and C). Because all sources in the multicast group use a common shared tree, a wildcard notation written as (*, G), pronounced "star comma G," represents the tree. In this case, "*" means all sources, and G represents the multicast group. Therefore, the shared tree shown in the figure would be written as (*, 224.2.2.2).

Source Trees vs. Shared Trees

Both source trees and shared trees are loop-free. Messages are replicated only where the tree branches.

Members of multicast groups can join or leave at any time; therefore, the distribution trees must be dynamically updated. When all the active receivers on a particular branch stop requesting the traffic for a particular multicast group, the routers prune that branch from the distribution tree and stop forwarding traffic down that branch. If one receiver on that branch becomes active and requests the multicast traffic, the router will dynamically modify the distribution tree and start forwarding traffic again.

Source trees have the advantage of creating the optimal path between the source and the receivers. This advantage guarantees the minimum amount of network latency for forwarding multicast traffic. However, this optimization comes at a cost: the routers must maintain path information for each source. In a network that has thousands of sources and thousands of groups, this overhead can quickly become a resource issue on the routers. Memory consumption from the size of the multicast routing table is a factor that network designers must take into consideration.

Shared trees have the advantage of requiring the minimum amount of network state information in each router. This advantage lowers the overall memory requirements for a network that allows shared trees only. The disadvantage of shared trees is that under certain circumstances the paths between the source and receivers might not be the optimal paths, which might introduce some latency in packet delivery. For example, in the figure, the shortest path between host A (source 1) and host B (a receiver) would be router A and router C. Because router D is used as the root for a shared tree, the traffic must traverse routers A, B, D, and then C. Network designers must carefully consider the placement of the RP when implementing a shared tree-only environment.

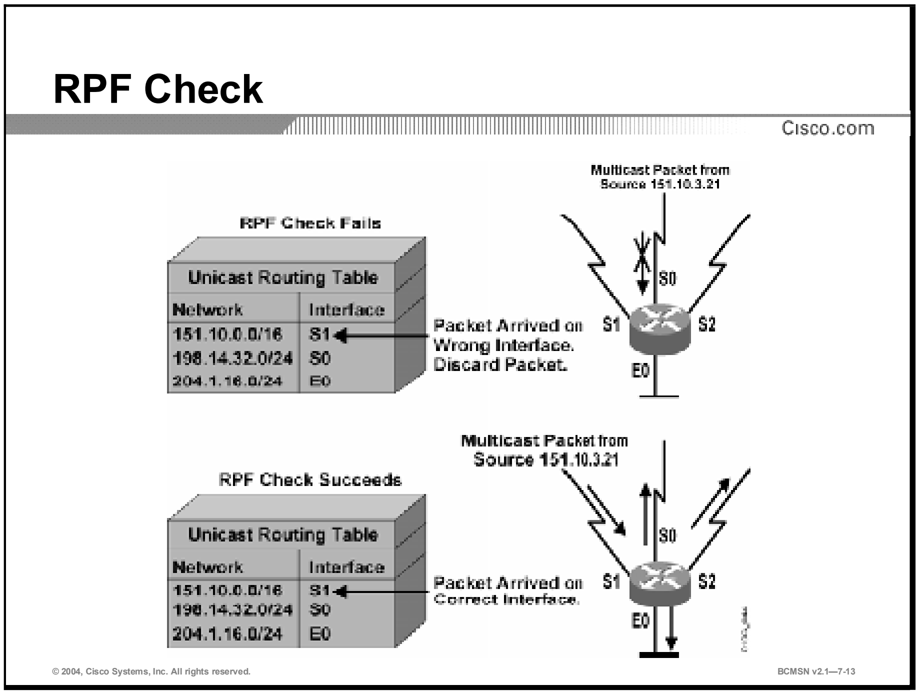

RPF Check

Cisco.com

RPF Check Fails

Unicast Routing Table	
Network	Interface
151.10.0.0/16	S1
198.14.32.0/24	S0
204.1.16.0/24	E0

Packet Arrived on Wrong Interface. Discard Packet.

RPF Check Succeeds

Unicast Routing Table	
Network	Interface
151.10.0.0/16	S1
198.14.32.0/24	S0
204.1.16.0/24	E0

Packet Arrived on Correct Interface.

Multicast Packet from Source 151.10.3.21
S0
S1
S2
E0
Multicast Packet from Source 151.10.3.21
S0
S1
S2
E0

BCMSN v2.1—7-13

In unicast routing, traffic is routed through the network along a single path from the source to the destination host. A unicast router does not consider the source address; it considers only the destination address and how to forward the traffic toward that destination. The router scans through its routing table for the destination address and then forwards a single copy of the unicast packet out the correct interface in the direction of the destination.

In multicast forwarding, the source is sending traffic to an arbitrary group of hosts that is represented by a multicast group address. The multicast router must determine which direction is the upstream direction (toward the source) and which one is the downstream direction (or directions). If there are multiple downstream paths, the router replicates the packet and forwards it down the appropriate downstream paths (best unicast route metric), which is not necessarily all paths.

RPF enables routers to correctly forward multicast traffic down the distribution tree. RPF makes use of the existing unicast routing table to determine the upstream and downstream neighbors. A router will forward a multicast packet only if the router receives the packet on the upstream interface. This RPF check helps to guarantee that the distribution tree will be loop-free.

When a multicast packet arrives at a router, the router will perform an RPF check on the packet. If the RPF check is successful, the packet will be forwarded; otherwise it will be dropped.

For traffic flowing down a source tree, the RPF check procedure works as follows:

Step 1 Router looks up the source address in the unicast routing table to determine if the packet has arrived on the interface located on the reverse path back to the source.

Step 2 If a packet has arrived on the interface leading back to the source, the RPF check is successful and the packet will be forwarded.

Step 3 If the RPF check in Step 2 fails, the packet is quietly dropped.

At the top of the figure, the RPF check fails. A multicast packet from source 151.10.3.21 is received on interface S0. A check of the unicast route table shows that S1 is the interface this router would use to forward unicast data to 151.10.3.21. Because the packet has arrived on S0, the packet will be discarded.

At the bottom of the figure, the RPF check succeeds. This time the multicast packet has arrived on S1. The router checks the unicast routing table and finds that S1 is the correct interface. The RPF check passes and the packet will be forwarded.

PIM Dense Mode and PIM Sparse Mode

PIM is IP routing protocol-independent and can leverage whichever unicast routing protocols are used to populate the unicast routing table. This topic explains PIM.

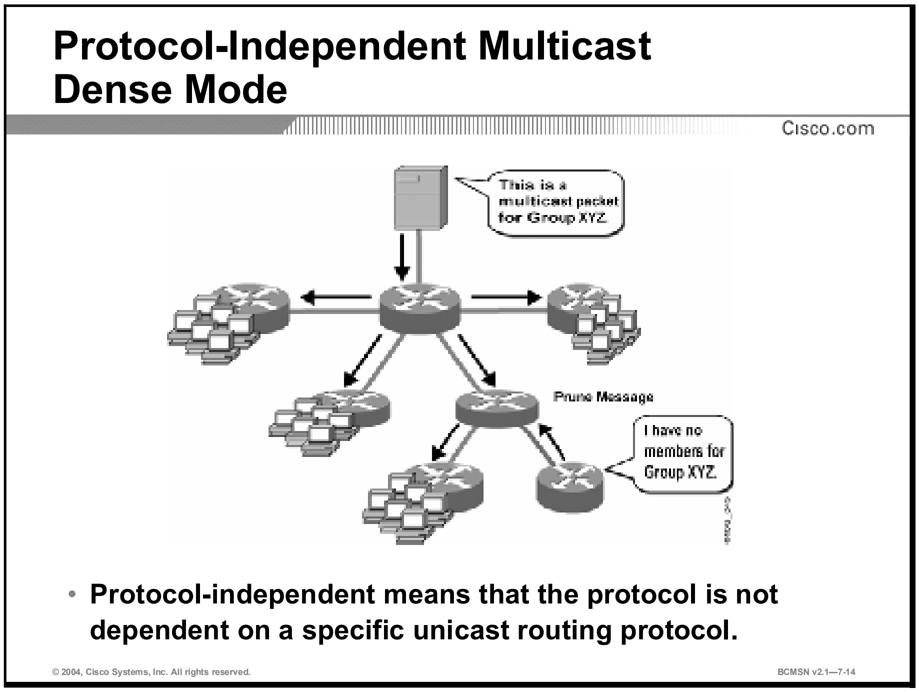

Protocol-Independent Multicast Dense Mode

Cisco.com

This is a multicast packet for Group XYZ.
Prune Message
I have no members for Group XYZ.

- **Protocol-independent means that the protocol is not dependent on a specific unicast routing protocol.**

BCMSN v2.1—7-14

A multicast routing protocol is responsible for the construction of multicast delivery trees and enabling multicast packet forwarding. Different IP multicast routing protocols use different techniques to construct multicast spanning trees and to forward packets. While PIM is independent of a specific routing protocol, a routing table is required to perform the RPF check.

The first approach is based on the assumption that the multicast group members are densely distributed throughout the network and bandwidth is plentiful, meaning that almost all hosts on the network belong to the group. These dense-mode multicast routing protocols rely on periodic flooding of the network with multicast traffic to set up and maintain the distribution tree.

PIM dense mode works best when there are numerous members belonging to each multimedia group. PIM floods the multimedia packet out to all routers in the network and then prunes routers that do not support members of that particular multicast group.

PIM dense mode is most useful under the following circumstances:

- Senders and receivers are in close proximity to one another.
- There are few senders and many receivers.
- The volume of multicast traffic is high.
- The stream of multicast traffic is constant.

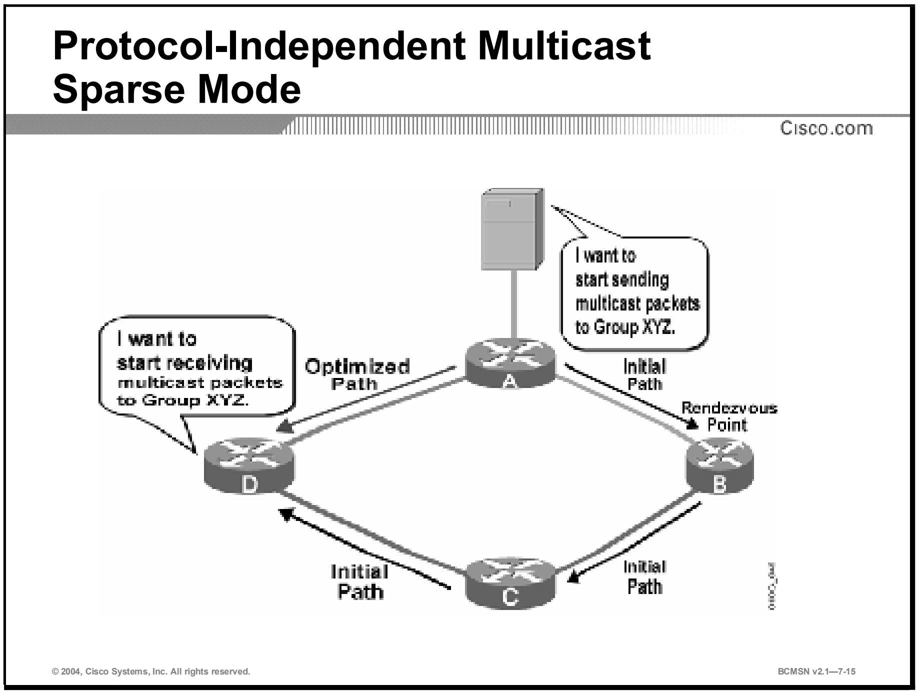

Protocol-Independent Multicast Sparse Mode

Cisco.com

I want to start sending multicast packets to Group XYZ.
I want to start receiving multicast packets to Group XYZ.
Optimized Path
A
Initial Path
Rendezvous Point
D
B
Initial Path
C
Initial Path

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—7-15

The second approach to multicast routing is based on the assumption that the multicast group members are sparsely distributed throughout the network and bandwidth is not necessarily widely available.

It is important to note that sparse mode does not imply that the group has few members, just that they are widely dispersed. In this case, flooding would unnecessarily waste network bandwidth and could cause serious performance problems. Therefore, sparse-mode multicast routing protocols must rely on more selective techniques to set up and maintain multicast trees. Sparse-mode protocols begin with an empty distribution tree and add branches only as the result of explicit requests to join the distribution.

Sparse-mode PIM is optimized for environments where there are many multipoint data streams. Sparse multicast is most useful when:

- There are few receivers in a group.
- The type of traffic is intermittent.

In sparse mode, each data stream goes to a relatively small number of segments in the campus network. Instead of flooding the network to determine the status of multicast members, sparse-mode PIM defines a rendezvous point. When a sender wants to send data, the sender first sends to the rendezvous point. When a receiver wants to receive data, the receiver registers with the rendezvous point. After the data stream begins to flow from sender to rendezvous point to receiver, the routers in the path will optimize the path automatically to remove any unnecessary hops. Sparse-mode PIM assumes that no hosts want the multicast traffic unless they specifically ask for it.

PIM is able to simultaneously support dense mode for some multicast groups and sparse mode for others. Cisco has implemented an alternative to choosing just dense mode or just sparse mode on a router interface. PIM sparse-dense mode allows the network to determine which IP Multicast groups should use sparse mode and which groups should use dense mode. PIM sparse mode and sparse-dense mode require the use of a rendezvous point.

IGMP and CGMP Operation

Internet Group Management Protocol (IGMP) is used to dynamically register individual hosts in a multicast group. Cisco Group Management Protocol (CGMP) allows Catalyst switches to learn about the existence of multicast clients from Cisco routers and Layer 3 switches. This topic discusses IGMP and CGMP.

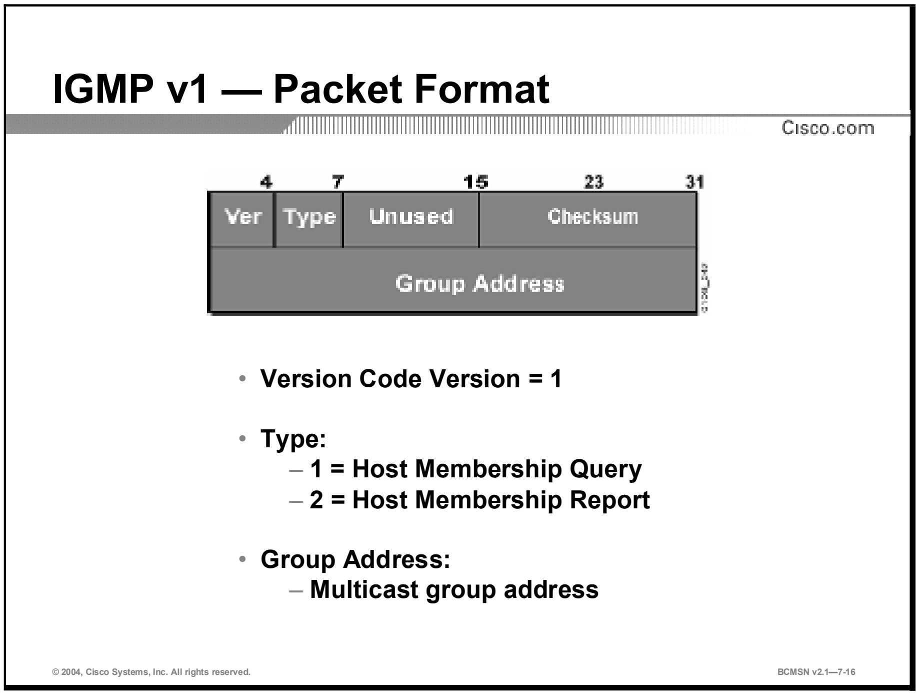

IGMP v1 — Packet Format

Cisco.com

4 7 15 23 31

Ver	Type	Unused	Checksum
Group Address			

- Version Code Version = 1
- Type:
  - 1 = Host Membership Query
  - 2 = Host Membership Report
- Group Address:
  - Multicast group address

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—7-16

IGMP is used to dynamically register individual hosts in a multicast group on a particular LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

According to the IGMP version 1 (IGMP v1) specification, one multicast router per LAN must periodically transmit Host Membership Query messages. This action determines which host groups have members on the networks to which the querier is directly attached. IGMP query messages are addressed to the all-host group (224.0.0.1) and have an IP TTL equal to 1. This TTL ensures that the query messages sourced from a router are transmitted onto the directly attached network and are not forwarded by any other multicast routers.

When the end station receives an IGMP query message, the end station responds with a host membership report for each group into which the end station belongs.

IGMP messages are specified in the IP datagram with a protocol value of 2.

The table describes the fields of the IGMP message.

Field	Definition
Type	There are two types of IGMP messages of concern to hosts: 1 = Host Membership Query 2 = Host Membership Report
Unused	Unused field, zeroed when sent, ignored when received.
Checksum	The checksum is the 16-bit one’s complement of the one's complement sum of the 8-octet IGMP message. For computing the checksum, the checksum field is zeroed.
Group Address	In a Host Membership Query message, the group address field is set to zero when sent and ignored when received. In a Host Membership Report message, the group address field holds the IP host group address of the group being reported.

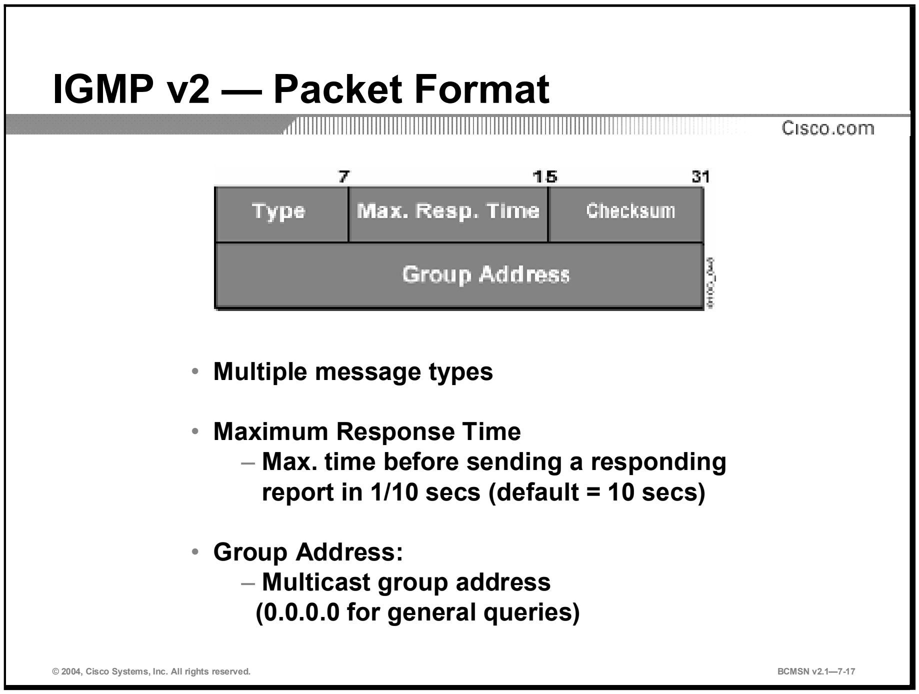

IGMP v2 — Packet Format

Cisco.com

7	15	31
Type	Max. Resp. Time	Checksum
Group Address		

- Multiple message types
- Maximum Response Time
  - Max. time before sending a responding report in 1/10 secs (default = 10 secs)
- Group Address:
  - Multicast group address (0.0.0.0 for general queries)

BCMSN v2.1—7-17

Version 2 of IGMP (IGMP v2) made some enhancements to the previous version, including the definition of a Group Specific Query. This type of message allows the router to transmit a specific query to one particular group. IGMP v2 also defines a Leave Group Message for the hosts, which results in lower leave latency.

There are four types of IGMP messages of concern to the host-router interaction:

- Membership query
- IGMP v2 membership report
- Leave report
- IGMP v1 membership report

The IGMP v1 membership report is used for backward compatibility with IGMP v1. New message types can be used by newer versions of IGMP or by multicast routing protocols. Any other or unrecognized message types are ignored.

The table describes the IGMP v2 message fields.

Field	Value
Type	0 x 11 = Membership Query 0 x 12 = Version 1 membership report 0 x 16 = Version 2 membership report 0 x 17 = Leave Report
Maximum Response Time	10 seconds = Default value. Meaningful only in a Membership Query. Specifies the maximum time allowed before sending a responding report in units of 1/10 second. 0 = All other messages
Checksum	Calculated the same as for the ICMP checksum
Group Address	0.0.0.0 in a General Query Group address queried in a Group Specific Query Multicast group address in a report

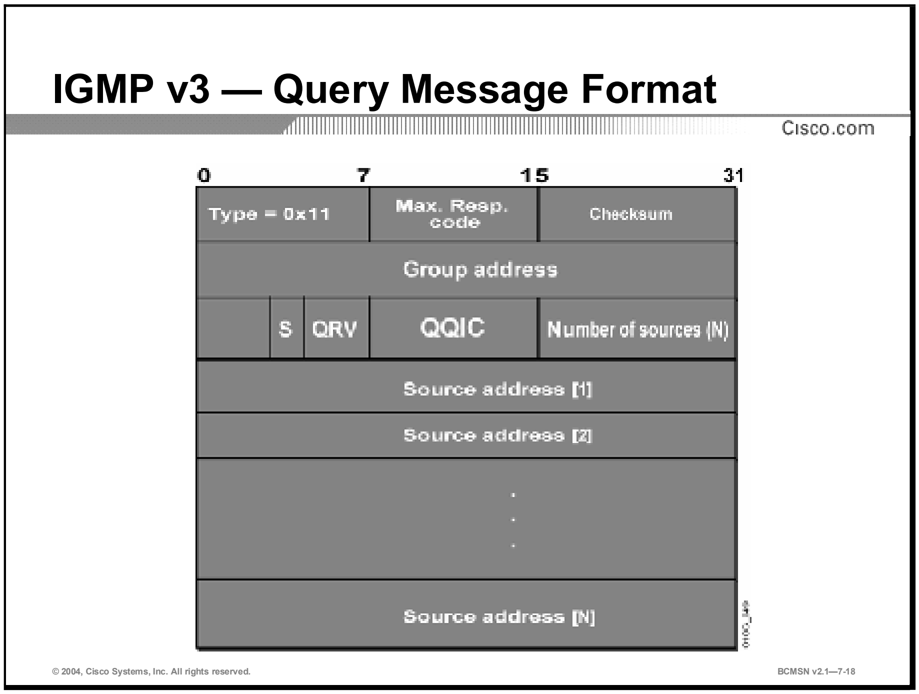

IGMP v3 — Query Message Format

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—7-18

IGMP version 3 (IGMP v3) is the next step in the evolution of IGMP. IGMP v3 adds support for source filtering. This enables a multicast receiver host to signal to a router the groups from which it wants to receive multicast traffic and from which sources this traffic is expected. This membership information enables IOS software to forward traffic from only those sources from which receivers requested the traffic.

In IGMP v3, the following types of IGMP messages exist:

- IGMP v3 membership query
- IGMP v3 membership report

The table describes the fields in the IGMP v3 query message.

Field	Definition
Type = 0x11	IGMP query.
Max resp. code	Maximum response code (in seconds). This field specifies the maximum time allowed before sending a responding report.
Group address	Multicast group address. This address is 0.0.0.0 for general queries.
S	S flag. This flag indicates that processing by routers is being suppressed.
QRV	Querier Robustness Value. This value affects timers and the number of retries.
QQIC	Query Interval Code of the Querier (in seconds). This field specifies the query interval used by the querier.
No. of sources [N]	Number of sources present in the query. This number is nonzero for a group-and-source query.
Source address [1...N]	Address of the source(s).

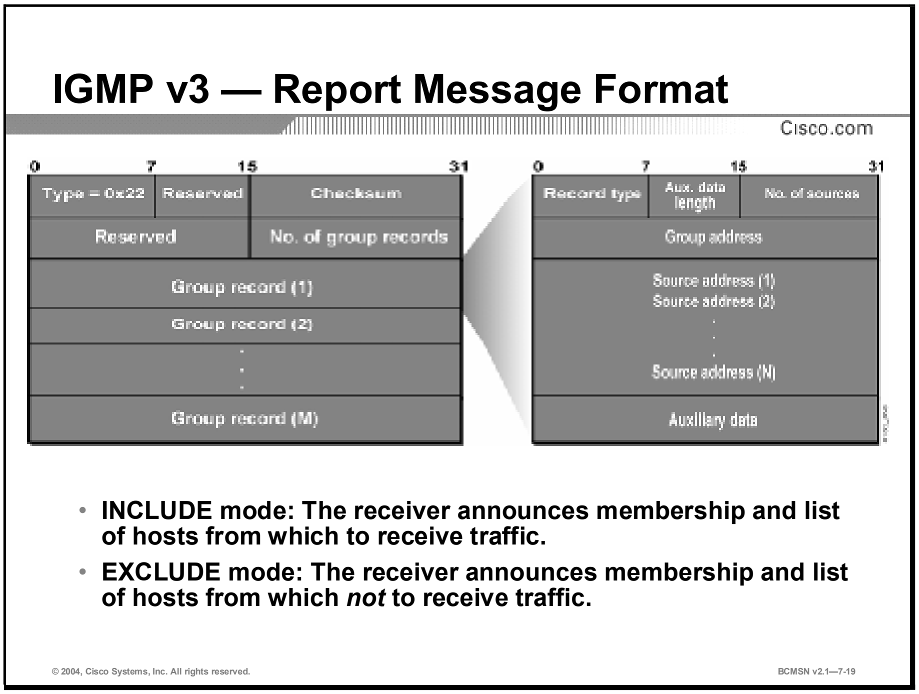

IGMP v3 — Report Message Format

Cisco.com

- **INCLUDE mode: The receiver announces membership and list of hosts from which to receive traffic.**
- **EXCLUDE mode: The receiver announces membership and list of hosts from which *not* to receive traffic.**

BCMSN v2.1—7-19

The table describes the fields in the IGMP v3 report message.

Field	Definition
No. of group records [M]	Number of group records present in the report
Group record [1...M]	Block of fields containing information regarding the sender membership with a single multicast group on the interface from which the report was sent
Record type	The group record type (e.g., MODE_IS_INCLUDE, MODE_IS_EXCLUDE)
No. of sources [N]	Number of sources present in the record
Source address [1...N]	Address of the source(s)

IGMP v3 supports applications that explicitly signal sources from which they want to receive traffic. With IGMP v3, receivers signal membership to a multicast host group in these two modes:

- **INCLUDE mode:** In this mode, the receiver announces membership to a host group and provides a list of source addresses (the INCLUDE list) from which it wants to receive traffic.
- **EXCLUDE mode:** In this mode, the receiver announces membership to a multicast group and provides a list of source addresses (the EXCLUDE list) from which it does not want to receive traffic. The host will receive traffic only from sources whose IP addresses are not listed in the EXCLUDE list. To receive traffic from all sources, which is the behavior of IGMP v2, a host uses EXCLUDE mode membership with an empty EXCLUDE list.

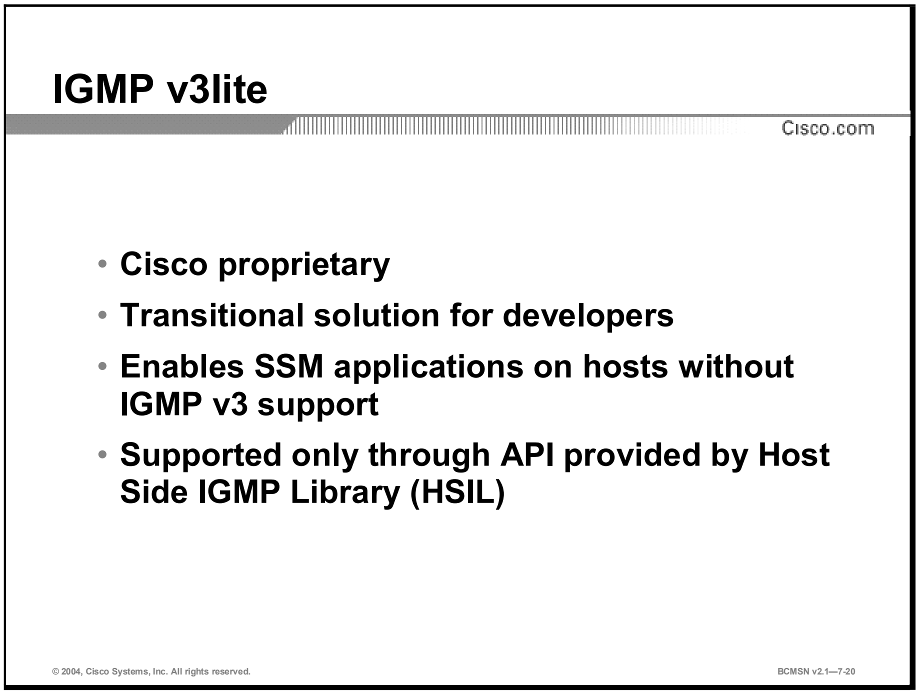

IGMP v3lite

Cisco.com

- Cisco proprietary
- Transitional solution for developers
- Enables SSM applications on hosts without IGMP v3 support
- Supported only through API provided by Host Side IGMP Library (HSIL)

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—7-20

IGMP version 3 lite (IGMP v3lite) is a Cisco transitional solution for application developers to start programming SSM applications immediately. It allows you to write and run SSM applications on hosts that do not yet support IGMP v3 in their operating system kernel.

For IGMP v3lite, applications must be compiled with the Host Side IGMP Library (HSIL). The HSIL software provides applications with a subset of the IGMP v3 application programming interface (API), which is required to write SSM applications.

One part of the HSIL is a client library linked to the SSM application. This library provides the SSM subset of the IGMP v3 API to the SSM application. If possible, the library checks whether the operating system kernel supports IGMP v3. If it does, the API calls are simply passed through to the kernel. If the kernel does not support IGMP v3, the library uses the IGMP v3lite mechanism.

When using the IGMP v3lite mechanism, the library tells the operating system kernel to join to the whole multicast group. Joining to the whole group is the only method by which the application can receive traffic for that multicast group for IGMP v1 or IGMP v2. In addition, the library signals the (S, G) channel subscriptions to an IGMP v3lite server process. A server process is needed because multiple SSM applications may be on the same host. This server process will then send IGMP v3lite-specific (S, G) channel subscriptions to the last-hop IOS router, which needs to be enabled for IGMP v3lite.

This IOS router will then see both the IGMP v1 and IGMP v2 group membership report from the operating system kernel. The router also will see the (S, G) channel subscription from the HSIL daemon. If the router sees both of these messages, the router will interpret the messages as an SSM (S, G) channel subscription and join to the channel through PIM SSM. Refer to the documentation accompanying the HSIL software for further information on how to utilize IGMP v3lite with your application.

IGMP v3lite is supported by Cisco only through the API provided by the HSIL, not as a function of the router independent of the HSIL. By default, IGMP v3lite is disabled. When

IGMP v3lite is configured through the **ip igmp v3lite** command on an interface, it will be active for IP multicast addresses in the SSM range only.

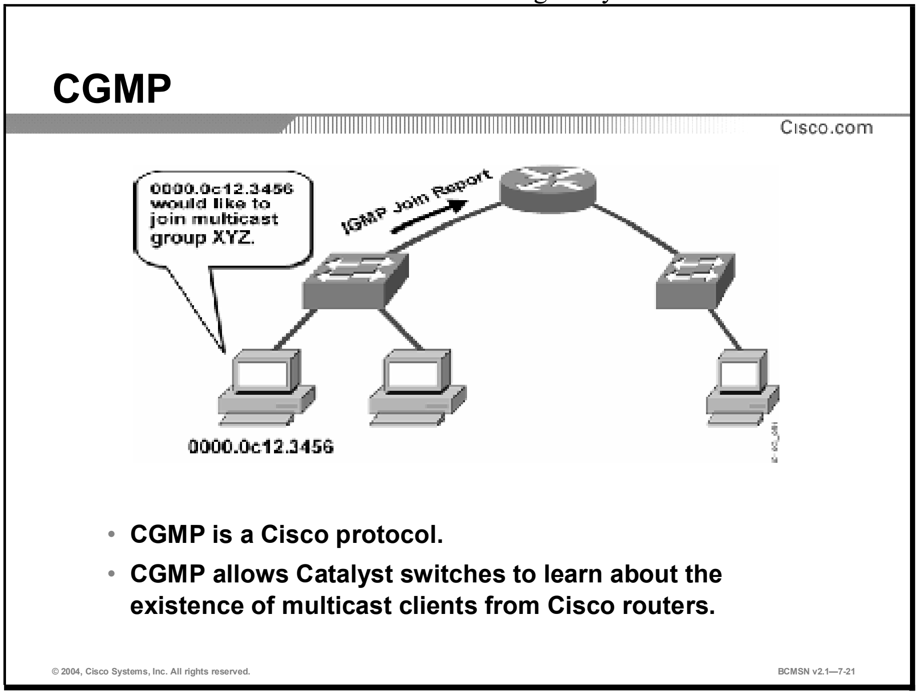

CGMP

- **CGMP is a Cisco protocol.**
- **CGMP allows Catalyst switches to learn about the existence of multicast clients from Cisco routers.**

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—7-21

For example, a video client wants to watch a 1.5-Mbps IP multicast-based video feed sent from a corporate video server. The video client sends an IGMP join message to the video server. The next-hop router for the client logs the IGMP join message. IP multicast traffic is transmitted downstream to the video client. The switch detects the incoming traffic and examines the destination MAC address to determine where the traffic should be forwarded. Because the destination MAC address is a multicast address and there are no entries in the switching table for where the traffic should go, the 1.5-Mbps video feed is simply sent to all ports.

Switches must have an architecture that allows multicast traffic to be forwarded to a large number of attached group members without unduly loading the switch fabric. This function allows the switch to provide support for the growing number of new multicast applications without having an impact on other traffic. Layer 2 switches also need some degree of multicast awareness to avoid flooding multicasts to all switch ports.

Multicast control in Layer 2 switches can be accomplished in the following ways:

- VLANs can correspond to the boundaries of the multicast group. This is a simple approach. However, this approach does not support dynamic changes to group membership, and it adds to the administrative burden of unicast VLANs.
- Layer 2 switches can examine, or "snoop," IGMP queries and reports to learn the port mappings of multicast group members. This allows the switch to dynamically track group membership. However, because snooping every multicast data and control packet consumes a lot of switch processing capacity, doing so can degrade forwarding performance and increase latency.

The traditional role of the router as a control point in the network can be maintained by defining a multicast router-to-switch protocol. The CGMP allows the router to configure the multicast forwarding table in the switch to correspond with the current group membership.

CGMP is a Cisco protocol that allows Catalyst switches to learn about the existence of multicast clients from Cisco routers and Layer 3 switches.

CGMP is based on a client-server model. The router is considered a CGMP server, with the switch taking on the client role. The basis of CGMP is that the IP multicast router sees all IGMP packets and, therefore, can inform the switch when specific hosts join or leave multicast groups. The switch then uses this information to construct a forwarding table.

When the router sees an IGMP control packet, the router creates a CGMP packet. This CGMP packet contains the request type, either a join message or a leave message, the multicast group address, and the actual MAC address of the client. The packet is sent to a well-known address to which all switches listen. Each switch then interprets the packet and creates the proper entries in a forwarding table.

Building on the previous video example, the client starts by sending an IGMP join message. However, when the next-hop router receives the IGMP join message, the router records the source MAC address of the IGMP join message and issues a CGMP join message back downstream to the Catalyst switch. The Catalyst switch uses the CGMP join message to dynamically build an entry in the switching table that maps the multicast traffic to the switch port of the client. In the example, the server delivers the 1.5-Mbps video feed only to switch ports defined in the switching table. The ports on the switch that do not support any hosts in the multicast group do not propagate the traffic.

IGMP Snooping

IGMP snooping is an IP multicast constraining mechanism that examines some Layer 3 information to maintain a Layer 2 multicast table. This topic discusses IGMP snooping.

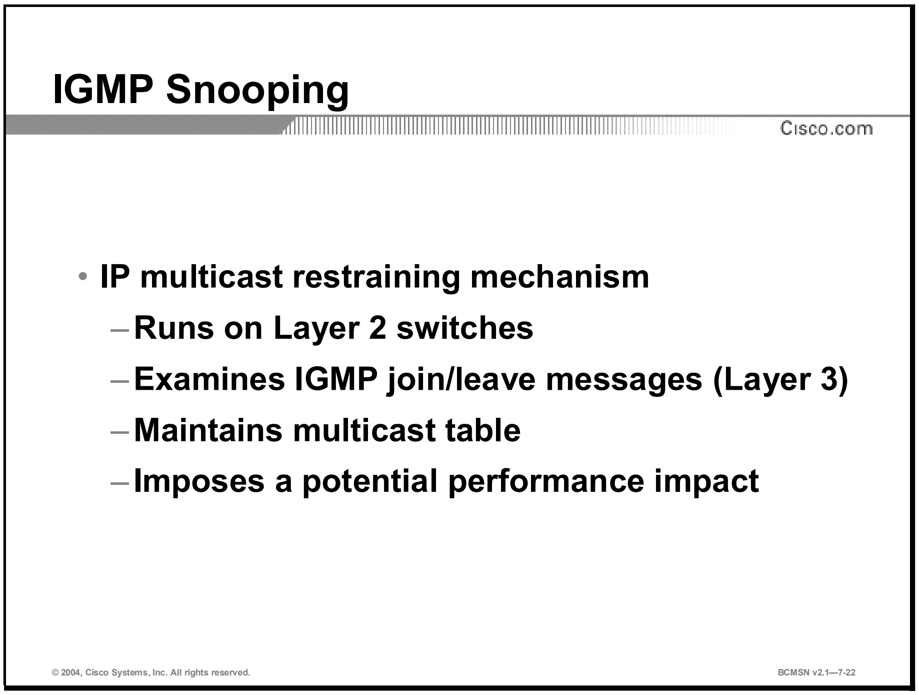

IGMP Snooping

Cisco.com

- IP multicast restraining mechanism
  - Runs on Layer 2 switches
  - Examines IGMP join/leave messages (Layer 3)
  - Maintains multicast table
  - Imposes a potential performance impact

BCMSN v2.1—7-22

The default behavior for a Layer 2 switch is to forward all multicast traffic to every port that belongs to the destination LAN on the switch. This behavior reduces the efficiency of the switch, whose purpose is to limit traffic to the ports that need to receive the data.

IGMP snooping runs on a Layer 2 LAN switch. IGMP snooping requires the LAN switch to snoop some Layer 3 information, such as IGMP join and leave messages, in the IGMP packets sent between the hosts and the router. When the switch hears the IGMP host report from a host for a particular multicast group, the switch adds the port number of the host to the associated multicast table entry. When the switch hears the IGMP leave group message from a host, the switch removes the table entry of the host.

Because IGMP control messages are transmitted as multicast packets, they are indistinguishable from multicast data at Layer 2. A switch running IGMP snooping must examine every multicast data packet to check if it contains any pertinent IGMP control information. If IGMP snooping is implemented on a low-end switch with a slow CPU, this could have a severe performance impact when data is transmitted at high rates. The solution is to implement IGMP snooping on high-end switches with special application-specific integrated circuits (ASICs) that can perform the IGMP checks in hardware. CGMP is ideal for low-end switches without special hardware.

Summary

This topic summarizes the key points discussed in this lesson.

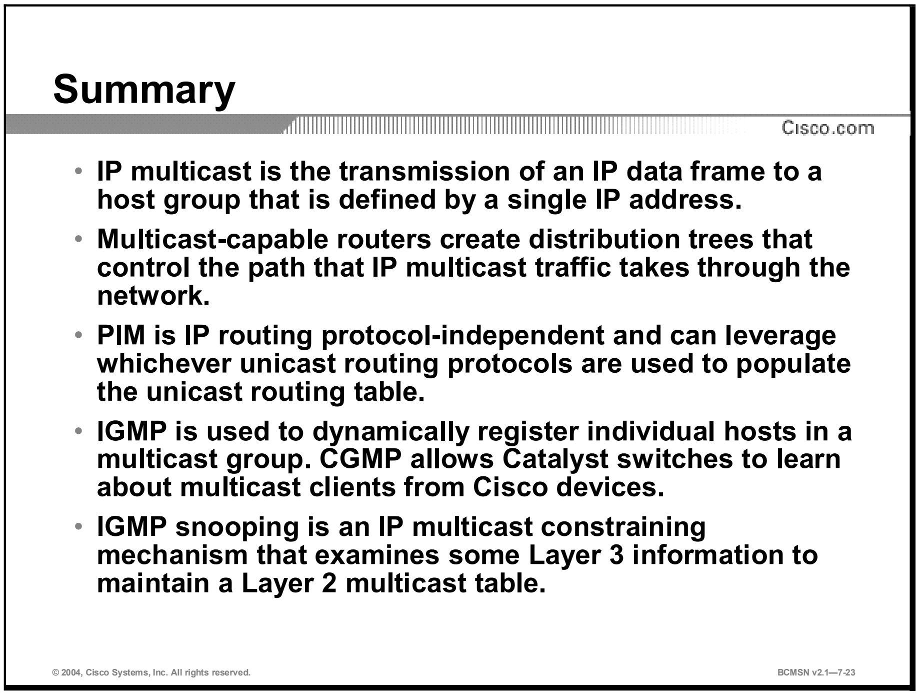

Summary

Cisco.com

- IP multicast is the transmission of an IP data frame to a host group that is defined by a single IP address.
- Multicast-capable routers create distribution trees that control the path that IP multicast traffic takes through the network.
- PIM is IP routing protocol-independent and can leverage whichever unicast routing protocols are used to populate the unicast routing table.
- IGMP is used to dynamically register individual hosts in a multicast group. CGMP allows Catalyst switches to learn about multicast clients from Cisco devices.
- IGMP snooping is an IP multicast constraining mechanism that examines some Layer 3 information to maintain a Layer 2 multicast table.

BCMSN v2.1—7-23

References

For additional information, refer to this resource:

- "Multicast Services" at http://www.cisco.com/warp/public/779/largeent/learn/technologies/multicast.html

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What differentiates IP multicast from other transmission modes?

A) IP multicast sends packets to a single host.

B) IP multicast sends packets to members of a group.

C) IP multicast sends packets to all hosts sequentially.

D) IP multicast sends packets to all hosts simultaneously.

Q2) What is one potential drawback to source distribution trees compared to shared distribution trees?

A) increased latency

B) increased memory overhead

C) suboptimal path calculations

D) increased bandwidth utilization

Q3) What is indicated by PIM sparse mode as opposed to dense mode?

A) Bandwidth is plentiful.

B) The multicast group has many members.

C) The hosts in the multicast group are widely dispersed.

D) Almost all hosts on the network belong to the multicast group.

Q4) Which version of IGMP added support for source filtering?

A) IGMP v1

B) IGMP v2

C) IGMP v3

D) IGMP v3lite

Q5) What is one potential drawback of IGMP snooping on a low-end switch without special ASICs?

A) increased latency

B) reduced bandwidth

C) degraded performance

D) increased administrative complexity

Quiz Answer Key

Q1) B

Relates to: IP Multicast Operation Overview

Q2) B

Relates to: Reverse Path Forwarding

Q3) C

Relates to: PIM Dense Mode and PIM Sparse Mode

Q4) C

Relates to: IGMP and CGMP Operation

Q5) C

Relates to: IGMP Snooping

Configuring IP Multicast in a Multilayer Switched Network

Overview

IP multicast provides bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to multiple recipients. It enables distribution of video conferencing, corporate communications, distance learning, distribution of software, and other applications. IP multicast packets are replicated in the network by Cisco Catalyst multilayer switches enabled with PIM and other supporting multicast protocols, which results in the most efficient delivery of data to multiple receivers. IP multicast is a large enough topic that it has its own course, *Implementing Cisco Multicast* (MCAST). The focus of this lesson is configuring PIM on Catalyst multilayer switches.

Relevance

With the increasing need for multicast applications in the multilayer switched network, the ability to support such applications with IP multicast is a basic requirement.

Objectives

Upon completing this lesson, you will be able to:

- Enable IP multicast on a Catalyst switch
- Configure IP multicast on a Catalyst switch
- Configure PIM version 2 on a Catalyst switch
- Monitor IP multicast on a Catalyst switch

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Enabling IP Multicast
- Configuring Auto-RP
- Configuring PIM Version 2
- Monitoring IP Multicast
- Summary
- Quiz

Enabling IP Multicast

You enable IP multicast with a single command in global configuration mode. You can then specify the PIM mode at the interface level. This topic discusses IP multicast configuration.

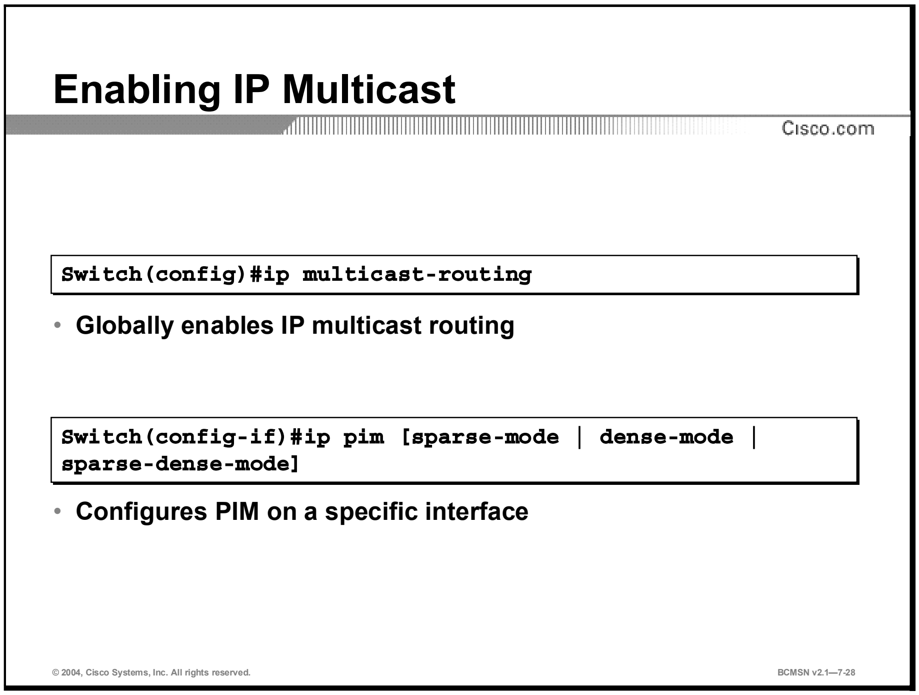

Enabling IP Multicast

Cisco.com

```
Switch(config)#ip multicast-routing
```

- Globally enables IP multicast routing

```
Switch(config-if)#ip pim [sparse-mode | dense-mode |
sparse-dense-mode]
```

- Configures PIM on a specific interface

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—7-28

Enabling IP multicast routing allows the switch to forward multicast packets. To enable IP multicast routing on the router, enter this command in global configuration mode:

```
Switch(config)#ip multicast-routing
```

Enabling PIM on an interface also enables IGMP operation on that interface. An interface can be configured to be in dense mode, sparse mode, or sparse-dense mode. The mode determines how the Layer 3 switch or router populates its multicast routing table and how the Layer 3 switch or router forwards multicast packets that it receives from its directly connected LANs. You must enable PIM in one of these modes for an interface to perform IP multicast routing.

When the switch populates the multicast routing table, dense-mode interfaces are always added to the table. Sparse-mode interfaces are added to the table only when periodic join messages are received from downstream routers, or when there is a directly connected member on the interface. Sparse-mode operation occurs if there is a RP known for the group. If so, the packets are encapsulated and sent toward the RP. When no RP is known, the packet is flooded in a dense-mode fashion. By default, multicast routing is disabled on an interface. There is no default mode setting.

To configure PIM on an interface to be in dense mode, enter this command in interface configuration mode:

```
Switch(config-if)#ip pim dense-mode
```

To configure PIM on an interface to be in sparse mode, enter this command in interface configuration mode:

```
Switch(config-if)#ip pim sparse-mode
```

When you enter either the **ip pim sparse-mode** or **ip pim dense-mode** command, that mode is applied to the interface as a whole. However, some environments might require PIM to run in a single region in sparse mode for some groups and in dense mode for other groups.

An alternative to enabling only dense mode or only sparse mode is to enable sparse-dense mode. In this case, the interface is treated as dense mode if the group is in dense mode; the interface is treated in sparse mode if the group is in sparse mode. If you want to treat the group as a sparse group, and the interface is in sparse-dense mode, you must have an RP.

If you configure sparse-dense mode, the idea of sparseness or denseness is applied to the group on the switch, and the network manager should apply the same concept throughout the network.

Another benefit of sparse-dense mode is that Auto-RP information can be distributed in a dense-mode manner, and multicast groups for user groups can be used in a sparse-mode manner. Thus, there is no need to configure a default RP at the leaf routers if Auto-RP is enabled.

To enable PIM to operate in the same mode as the group, enter this command in interface configuration mode:

```
Switch(config-if)#ip pim sparse-dense-mode
```

RPs are used by senders to a multicast group to announce their existence and used by receivers of multicast packets to learn about new senders. The IOS software can be configured so that packets for a single multicast group can use one or more RPs.

The RP address is used by first-hop routers to send PIM register messages on behalf of a host sending a packet to the group. The RP address is also used by last-hop routers to send PIM join and prune messages to the RP to inform it about group membership. You must configure the RP address on all routers (including the RP router).

To configure the address of the RP, enter this command in global configuration mode:

```
Switch(config)#ip pim rp-address ip-address [access-list-
number] [override]
```

Configuring Auto-RP

Auto-RP is a feature that automates the distribution of group-to-RP mappings in a network supporting sparse-mode PIM. This topic explains how to configure Auto-RP.

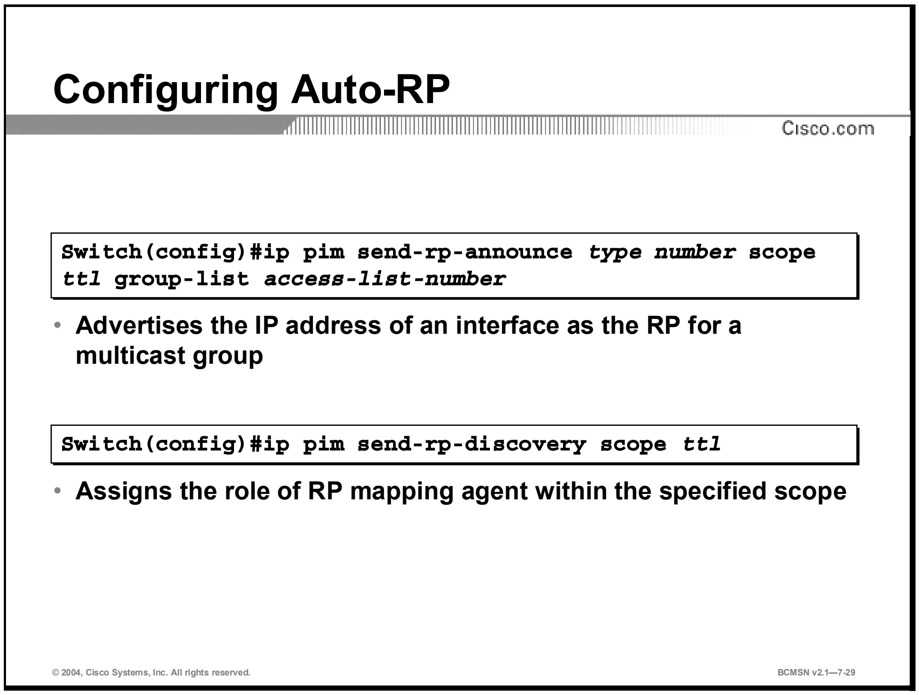

Configuring Auto-RP

Cisco.com

```
Switch(config)#ip pim send-rp-announce type number scope
ttl group-list access-list-number
```

- **Advertises the IP address of an interface as the RP for a multicast group**

```
Switch(config)#ip pim send-rp-discovery scope ttl
```

- **Assigns the role of RP mapping agent within the specified scope**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—7-29

Auto-RP is a feature that automates the distribution of group-to-RP mappings in a PIM network. This feature has these benefits:

- It is easy to use multiple RPs within a network to serve different group ranges.
- Auto-RP allows load splitting among different RPs and arrangement of RPs according to the location of group participants.
- Auto-RP avoids inconsistent, manual RP configurations that can cause connectivity problems.

Note If you configure PIM in sparse mode or sparse-dense mode and do not configure Auto-RP or any other RP learning component, you must statically configure an RP.

Sparse-mode environments need a default RP; sparse-dense-mode environments do not. If you have sparse-dense mode configured everywhere, you do not choose a default RP. If you are setting up Auto-RP in a new internetwork, you do not need a default RP because you configure all the interfaces for sparse-dense mode.

Note You only designate a default RP if you are configuring Auto-RP in an existing environment with sparse mode already configured.

Adding Auto-RP to a sparse-mode network requires a default RP. In an existing PIM sparse-mode region, at least one RP is defined across the network that has good connectivity and availability. That is, the **ip pim rp-address** command is already configured on all routers in this network.

Use that RP for the global groups (for example, 224.x.x.x and other global groups). There is no need to reconfigure the group address range that RP serves. RPs discovered dynamically through Auto-RP take precedence over statically configured RPs. Typically, you would use a second RP for the local groups.

Find another router to serve as the RP for the local groups. The RP mapping agent can double as an RP itself. Assign the whole range of 239.x.x.x to that RP, or assign a subrange of that range (for example, 239.2.x.x).

To designate a router as the RP, enter this command in global configuration mode:

```
Switch(config)#ip pim send-rp-announce type number scope ttl
group-list access-list-number
```

To change the group ranges that this RP optimally serves in the future, change the announcement setting on the RP. If the change is valid, all other routers automatically adopt the new group-to-RP mapping.

The RP mapping agent a discovery message notifying other routers which group-to-RP mapping to use. Such a role is necessary in the event of conflicts (such as overlapping group-to-RP ranges).

Find a router for which connectivity is not likely to be interrupted and assign it the role of RP mapping agent. All routers within the TTL number of hops from the source router receive the Auto-RP discovery messages. To assign the role of RP mapping agent in that router, enter this command in global configuration mode:

```
Switch(config)#ip pim send-rp-discovery scope ttl
```

Example: Designating an RP

The following example advertises the IP address of Ethernet 0 as the RP for the administratively scoped groups:

```
ip pim send-rp-announce ethernet0 scope 16 group-list 1
access-list 1 permit 239.0.0.0 0.255.255.255
```

Configuring PIM Version 2

PIM version 2 provides standards compliance and additional features over PIM version 1. This topic explains how to configure PIM version 2.

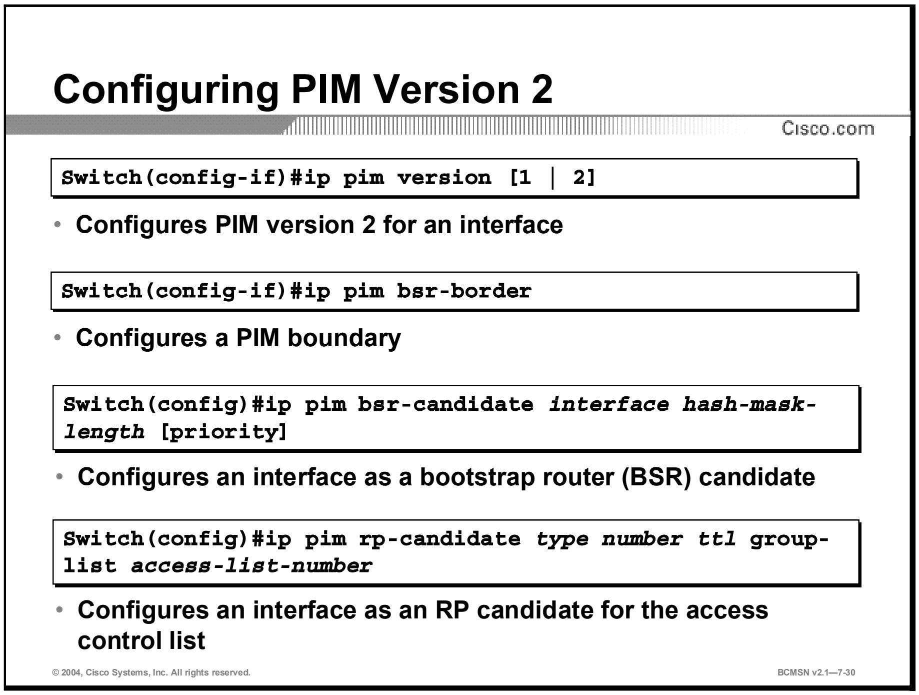

Configuring PIM Version 2

Cisco.com

```
Switch(config-if)#ip pim version [1 | 2]
```

- **Configures PIM version 2 for an interface**

```
Switch(config-if)#ip pim bsr-border
```

- **Configures a PIM boundary**

```
Switch(config)#ip pim bsr-candidate interface hash-mask-
length [priority]
```

- **Configures an interface as a bootstrap router (BSR) candidate**

```
Switch(config)#ip pim rp-candidate type number ttl group-
list access-list-number
```

- **Configures an interface as an RP candidate for the access control list**

BCMSN v2.1—7-30

PIM version 2 includes these improvements over PIM version 1:

- A single, active RP exists per multicast group, with multiple backup RPs. This single RP compares to multiple active RPs for the same group in PIM version 1.
- A bootstrap router (BSR) provides a fault-tolerant, automated RP discovery and distribution mechanism. Thus, routers dynamically learn the group-to-RP mappings.
- Sparse mode and dense mode are properties of a group, as opposed to an interface. Cisco strongly recommends sparse-dense mode.
- PIM join and prune messages have more flexible encodings for multiple address families.
- A more flexible hello packet format replaces the query packet to encode current and future capability options.
- Register messages to an RP indicate whether they were sent by a border router or a designated router.
- PIM packets are no longer inside IGMP packets; they are standalone packets.

PIM version 1, used with the Auto-RP feature, can perform the same tasks as the PIM version 2 BSR. However, Auto-RP is a standalone protocol, separate from PIM version 1, and it is Cisco proprietary. PIM version 2 is a standards track protocol in the Internet Engineering Task Force (IETF). Cisco recommends that you use PIM version 2.

Choose either the BSR or Auto-RP for a given range of multicast groups. If there are PIM version 1 routers in the network, do not use the BSR.

All systems using IOS Release 11.3(2)T or later start in PIM version 2 mode by default. In case you need to re-enable PIM version 2 or specify PIM version 1, you can control the PIM version by entering this command in interface configuration mode:

```
Switch(config-if)#ip pim version [1 | 2]
```

To configure PIM version 2 exclusively, perform the tasks in this topic. It is assumed that no PIM version 1 system exists in the PIM domain. The first task is to configure sparse-dense mode on all interfaces. If you configure Auto-RP, none of the other tasks are required to run PIM version 2.

Configure a border for the PIM domain, so that bootstrap messages do not cross this border in either direction. Therefore, different BSRs will be elected on the two sides of the PIM border. Use the following command on the interface of a border router peering with one or more neighbors outside the PIM domain. To configure a PIM domain boundary, enter this command in interface configuration mode:

```
Switch(config-if)#ip pim bsr-border
```

To avoid a single point of failure, you can configure several candidate BSRs in a PIM domain. A BSR is elected among the candidate BSRs automatically; they use bootstrap messages to discover which BSR has the highest priority. This router then announces to all PIM routers in the PIM domain that it is the BSR.

Routers that are configured as candidate RPs send the group range for which they are responsible to the BSR. The BSR includes this information in its bootstrap messages and disseminates it to all PIM routers in the domain. Based on this information, all routers will be able to map multicast groups to specific RPs. As long as a router is receiving the bootstrap message, the router has a current RP map.

You should configure one or more candidate BSRs. The routers that serve as candidate BSRs should be well-connected and be in the backbone portion of the network, rather than the dialup portion of the network. On the candidate BSRs, enter this command in global configuration mode:

```
Switch(config)#ip pim bsr-candidate interface hash-mask-length
[priority]
```

Configure one or more candidate RPs. Similar to BSRs, the RPs should also be well-connected and in the backbone portion of the network. An RP can serve the entire IP multicast address space or a portion of it. Candidate RPs send candidate RP advertisements to the BSR.

On the candidate RPs, enter this command in global configuration mode:

```
Switch(config)#ip pim rp-candidate type number ttl group-list
access-list-number
```

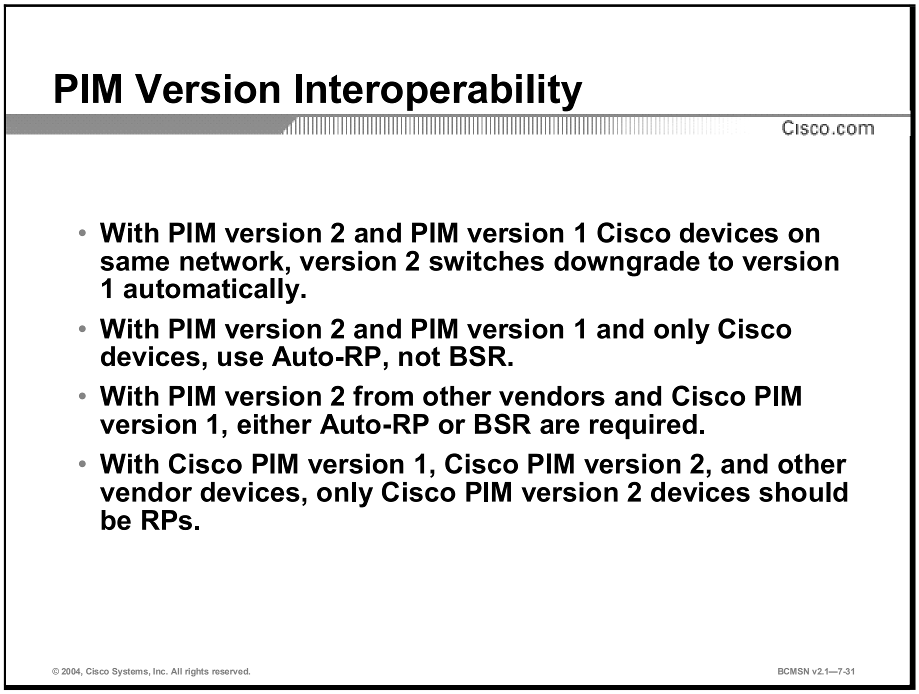

PIM Version Interoperability

- With PIM version 2 and PIM version 1 Cisco devices on same network, version 2 switches downgrade to version 1 automatically.
- With PIM version 2 and PIM version 1 and only Cisco devices, use Auto-RP, not BSR.
- With PIM version 2 from other vendors and Cisco PIM version 1, either Auto-RP or BSR are required.
- With Cisco PIM version 1, Cisco PIM version 2, and other vendor devices, only Cisco PIM version 2 devices should be RPs.

BCMSN v2.1—7-31

Cisco PIM version 2 implementation allows interoperability and transition between version 1 and version 2, although there might be some minor problems. You can upgrade to PIM version 2 incrementally. PIM versions 1 and 2 can be configured on different routers within one network. Internally, all routers on a shared media network must run the same PIM version. Therefore, if a PIM version 2 router detects a PIM version 1 router, the version 2 router downgrades itself to version 1 until all version 1 routers have been shut down or upgraded.

PIM uses the BSR to discover and announce RP-set information for each group prefix to all the routers in a PIM domain. This is the same function accomplished by Auto-RP, but the BSR is part of the PIM version 2 specification. The BSR mechanism interoperates with Auto-RP on Cisco routers.

When PIM version 2 routers interoperate with PIM version 1 routers, Auto-RP should have already been deployed. A PIM version 2 BSR that is also an Auto-RP mapping agent will automatically advertise the RP elected by Auto-RP. That is, Auto-RP prevails in its imposition of a single RP on every router in the group. All routers in the domain refrain from trying to use the PIM version 2 hash function to select multiple RPs.

Because bootstrap messages are sent hop by hop, a PIM version 1 router will prevent these messages from reaching all routers in your network. It is best to use Auto-RP rather than the bootstrap mechanism if your network contains Cisco routers and a PIM version 1 router. If you have a network that includes routers from other vendors, configure the Auto-RP mapping agent and the BSR on a Cisco PIM version 2 router. You should also ensure that no PIM version 1 router is located on the path between the BSR and a PIM version 2 router from another vendor.

If there are only Cisco routers in your network (no routers from other vendors), there is no need to configure a bootstrap router. Configure Auto-RP in the mixed PIM version 1 and version 2 environment.

Both Auto-RP and a BSR are required if you have other vendor PIM version 2 routers that need to interoperate with Cisco routers running PIM version 1. Cisco recommends that a Cisco PIM version 2 router be both the Auto-RP mapping agent and the BSR.

Note Dense-mode groups in a mixed PIM version 1 and version 2 region need no special configuration; they will interoperate automatically.

Sparse-mode groups in a mixed PIM version 1 and version 2 region are possible because the Auto-RP feature in version 1 interoperates with the RP feature of version 2. Although all PIM version 2 routers are also capable of using version 1, Cisco recommends that the RPs be upgraded to version 2 (or at least upgraded to PIM version 1 in the Cisco IOS Release 11.3 software).

To ease the transition to PIM version 2, Cisco also recommends that you do the following:

- Use Auto-RP throughout the region.
- Configure sparse-dense mode throughout the region.
- If Auto-RP was not already configured in the PIM version 1 regions, configure Auto-RP.

There are two approaches to using PIM version 2. You can use version 2 exclusively in your network, or migrate to version 2 by employing a mixed PIM version environment.

Consider the following when deciding which routers should be RPs:

- Any router can be configured as an RP in a network of Cisco routers using only Auto-RP.
- In a network of routers that includes only Cisco PIM version 2 routers and routers from other vendors, any router can be used as an RP.
- In a network of Cisco PIM version 1 routers, Cisco PIM version 2 routers, and routers from other vendors, only Cisco PIM version 2 routers should be configured as RPs.

Monitoring IP Multicast

You can use a variety of **show** commands to verify and monitor IP multicast. This topic explains how to verify and monitor IP multicast.

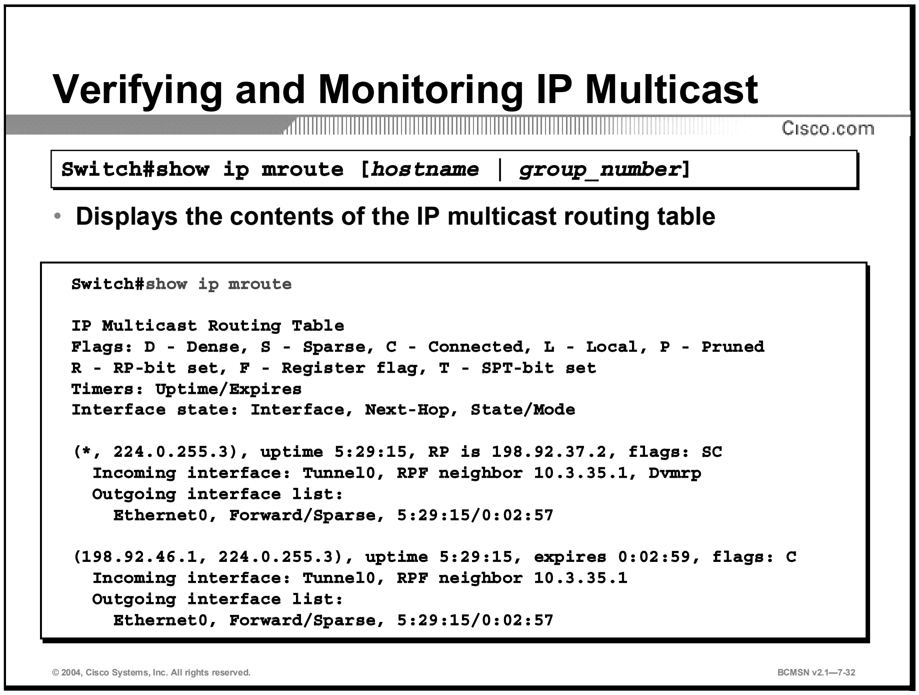

Verifying and Monitoring IP Multicast

Cisco.com

```
Switch#show ip mroute [hostname | group_number]
```

- **Displays the contents of the IP multicast routing table**

```
Switch#show ip mroute

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
R - RP-bit set, F - Register flag, T - SPT-bit set
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.0.255.3), uptime 5:29:15, RP is 198.92.37.2, flags: SC
  Incoming interface: Tunnel0, RPF neighbor 10.3.35.1, Dvmrp
  Outgoing interface list:
    Ethernet0, Forward/Sparse, 5:29:15/0:02:57

(198.92.46.1, 224.0.255.3), uptime 5:29:15, expires 0:02:59, flags: C
  Incoming interface: Tunnel0, RPF neighbor 10.3.35.1
  Outgoing interface list:
    Ethernet0, Forward/Sparse, 5:29:15/0:02:57
```

BCMSN v2.1—7-32

You can use a variety of **show** commands to verify and monitor IP multicast. You can display specific statistics, such as the contents of IP routing tables and databases. The information provided can be used to determine resource utilization and solve network problems. You can also display information about node reachability and discover the routing path that your device packets are taking through the network.

To display various routing statistics, you can enter any of these commands in EXEC mode:

Command	Description
ping [*group-name* | *group-address*]	Sends an ICMP echo request to a multicast group address.
show ip mroute [*hostname* | *group_number*]	Displays the contents of the IP multicast routing table.
show ip pim interface [*type number*] [**count**]	Displays information about interfaces configured for PIM.
show ip interface	Displays PIM information for all interfaces.

Example: show ip mroute Command for Sparse Mode

This is sample output from the **show ip mroute** command for a router operating in sparse mode:

```
Switch#show ip mroute

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P -
Pruned
R - RP-bit set, F - Register flag, T - SPT-bit set
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.0.255.3), uptime 5:29:15, RP is 198.92.37.2, flags: SC
  Incoming interface: Tunnel0, RPF neighbor 10.3.35.1, Dvmrp
  Outgoing interface list:
    Ethernet0, Forward/Sparse, 5:29:15/0:02:57

(198.92.46.1, 224.0.255.3), uptime 5:29:15, expires 0:02:59,
flags: C
  Incoming interface: Tunnel0, RPF neighbor 10.3.35.1
  Outgoing interface list:
    Ethernet0, Forward/Sparse, 5:29:15/0:02:57
```

Note Output interface timers are not updated for hardware-forwarded packets. Entry timers are updated approximately every 5 seconds.

Example: show ip mroute summary Command

This is sample output from the **show ip mroute** command with the **summary** keyword:

```
Switch#show ip mroute summary

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P -
Pruned
       R - RP-bit set, F - Register flag, T - SPT-bit set, J -
Join SPT
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.255.255.255), 2d16h/00:02:30, RP 171.69.10.13, flags:
SJPC

(*, 224.2.127.253), 00:58:18/00:02:00, RP 171.69.10.13, flags:
SJC

(*, 224.1.127.255), 00:58:21/00:02:03, RP 171.69.10.13, flags:
SJC

(*, 224.2.127.254), 2d16h/00:00:00, RP 171.69.10.13, flags:
SJCL
  (128.9.160.67/32, 224.2.127.254), 00:02:46/00:00:12, flags:
CLJT
  (129.48.244.217/32, 224.2.127.254), 00:02:15/00:00:40,
flags: CLJT
  (130.207.8.33/32, 224.2.127.254), 00:00:25/00:02:32, flags:
```

```
CLJT
  (131.243.2.62/32, 224.2.127.254), 00:00:51/00:02:03, flags:
CLJT
  (140.173.8.3/32, 224.2.127.254), 00:00:26/00:02:33, flags:
CLJT
  (171.69.60.189/32, 224.2.127.254), 00:03:47/00:00:46, flags:
CLJT
```

Example: show ip mroute active Command

This is sample output from the **show ip mroute** command with the **active** keyword:

```
Switch#show ip mroute active

Active IP Multicast Sources - sending >= 4 kbps

Group: 224.2.127.254, (sdr.cisco.com)
   Source: 146.137.28.69 (mbone.ipd.anl.gov)
     Rate: 1 pps/4 kbps(1sec), 4 kbps(last 1 secs), 4
kbps(life avg)

Group: 224.2.201.241, ACM 97
   Source: 130.129.52.160 (webcast3-e1.acm97.interop.net)
     Rate: 9 pps/93 kbps(1sec), 145 kbps(last 20 secs), 85
kbps(life avg)

Group: 224.2.207.215, ACM 97
   Source: 130.129.52.160 (webcast3-e1.acm97.interop.net)
     Rate: 3 pps/31 kbps(1sec), 63 kbps(last 19 secs), 65
kbps(life avg)
```

Example: show ip mroute count Command

This is sample output from the **show ip mroute** command with the **count** keyword:

```
Switch#show ip mroute count

IP Multicast Statistics - Group count: 8, Average sources per
group: 9.87
Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per
second

Group: 224.255.255.255, Source count: 0, Group pkt count: 0
  RP-tree: 0/0/0/0

Group: 224.2.127.253, Source count: 0, Group pkt count: 0
  RP-tree: 0/0/0/0

Group: 224.1.127.255, Source count: 0, Group pkt count: 0
  RP-tree: 0/0/0/0

Group: 224.2.127.254, Source count: 9, Group pkt count: 14
  RP-tree: 0/0/0/0
  Source: 128.2.6.9/32, 2/0/796/0
  Source: 128.32.131.87/32, 1/0/616/0
  Source: 128.125.51.58/32, 1/0/412/0
  Source: 130.207.8.33/32, 1/0/936/0
  Source: 131.243.2.62/32, 1/0/750/0
  Source: 140.173.8.3/32, 1/0/660/0
  Source: 146.137.28.69/32, 1/0/584/0
  Source: 171.69.60.189/32, 4/0/447/0
  Source: 204.162.119.8/32, 2/0/834/0

Group: 224.0.1.40, Source count: 1, Group pkt count: 3606
  RP-tree: 0/0/0/0
  Source: 171.69.214.50/32, 3606/0/48/0, RPF Failed: 1203

Group: 224.2.201.241, Source count: 36, Group pkt count: 54152
  RP-tree: 7/0/108/0
  Source: 13.242.36.83/32, 99/0/123/0
  Source: 36.29.1.3/32, 71/0/110/0
  Source: 128.9.160.96/32, 505/1/106/0
  Source: 128.32.163.170/32, 661/1/88/0
  Source: 128.115.31.26/32, 192/0/118/0
  Source: 128.146.111.45/32, 500/0/87/0
  Source: 128.183.33.134/32, 248/0/119/0
  Source: 128.195.7.62/32, 527/0/118/0
  Source: 128.223.32.25/32, 554/0/105/0
  Source: 128.223.32.151/32, 551/1/125/0
  Source: 128.223.156.117/32, 535/1/114/0
  Source: 128.223.225.21/32, 582/0/114/0
  Source: 129.89.142.50/32, 78/0/127/0
  Source: 129.99.50.14/32, 526/0/118/0
  Source: 130.129.0.13/32, 522/0/95/0
  Source: 130.129.52.160/32, 40839/16/920/161
  Source: 130.129.52.161/32, 476/0/97/0
  Source: 130.221.224.10/32, 456/0/113/0
  Source: 132.146.32.108/32, 9/1/112/0
```

Note Multicast route byte and packet statistics are supported only for the first 1024 multicast routes. Output interface statistics are not maintained.

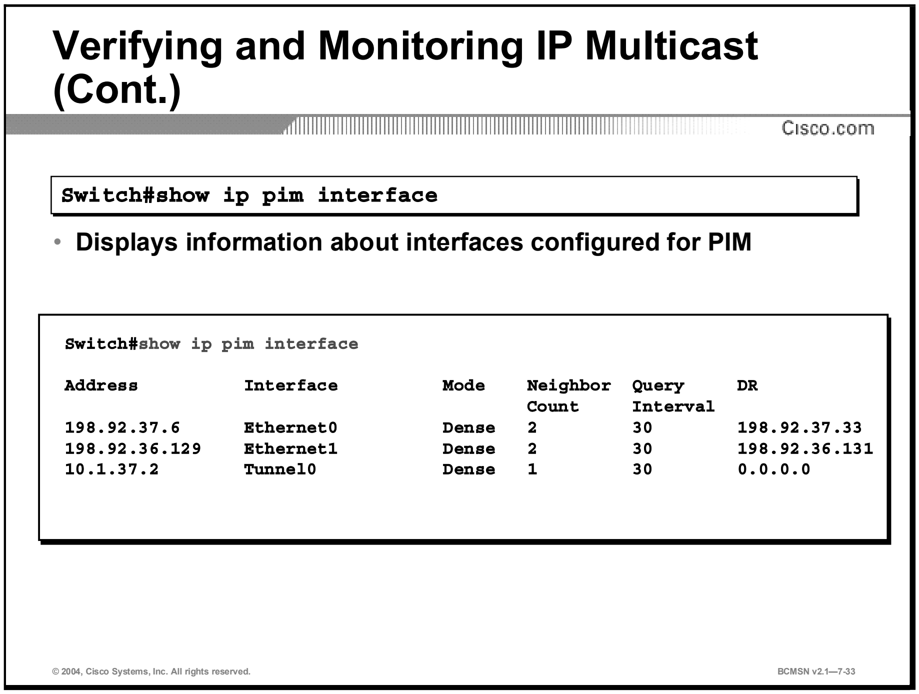

Verifying and Monitoring IP Multicast (Cont.)

```
Switch#show ip pim interface
```

- **Displays information about interfaces configured for PIM**

```
Switch#show ip pim interface

Address          Interface          Mode    Neighbor  Query     DR
                                            Count     Interval
198.92.37.6      Ethernet0          Dense   2         30        198.92.37.33
198.92.36.129    Ethernet1          Dense   2         30        198.92.36.131
10.1.37.2        Tunnel0            Dense   1         30        0.0.0.0
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—7-33

This is sample output from the **show ip pim interface** command:

```
Switch#show ip pim interface

Address          Interface          Mode    Neighbor  Query
designated router
                                            Count     Interval
198.92.37.6      Ethernet0          Dense   2         30
198.92.37.33
198.92.36.129    Ethernet1          Dense   2         30
198.92.36.131
10.1.37.2        Tunnel0            Dense   1         30
0.0.0.0
```

This is sample output from the **show ip pim interface** command with a **count**:

```
Switch#show ip pim interface count

Address          Interface          FS   Mpackets In/Out
171.69.121.35    Ethernet0          *    548305239/13744856
171.69.121.35    Serial0.33         *    8256/67052912
198.92.12.73     Serial0.1719       *    219444/862191
```

The following is sample output from the **show ip pim interface** command with a count when IP multicast is enabled. The example lists the PIM interfaces that are fast switched and process switched, and lists the packet counts for these. The H is added to interfaces on which IP multicast is enabled.

```
Switch#show ip pim interface count

States: FS - Fast Switched, H - Hardware Switched
Address          Interface          FS  Mpackets In/Out
192.1.10.2       Vlan10             * H 40886/0
192.1.11.2       Vlan11             * H 0/40554
192.1.12.2       Vlan12             * H 0/40554
192.1.23.2       Vlan23             *   0/0
192.1.24.2       Vlan24             *   0/0
```

To monitor the RP mapping information, you can enter these commands in EXEC mode:

Command	Description
show ip pim bsr	Displays information about the currently elected BSR.
show ip pim rp-hash *group*	Displays the RP that was selected for the specified group.
show ip pim rp [*group-name* | *group-address* | **mapping**]	Displays how the router learns of the RP (via bootstrap or Auto-RP mechanism).

When debugging interoperability problems between PIM version 1 and version 2, perform these tasks:

Step 1 Verify RP mapping with the **show ip pim rp-hash** command, making sure that all systems agree on the same RP for the same group.

Step 2 Verify interoperability between different versions of designated routers and RPs. Ensure that RPs are interacting with the designated routers properly (by responding with register-stop packets and forwarding de-encapsulated data packets from registers).

You can clear all contents of a particular cache, table, or database. This action is necessary if the contents of the particular structure have become, or are suspected to be, invalid.

To clear IP multicast caches, tables, and databases, enter these commands in EXEC mode:

Command	Description
clear ip mroute	Deletes entries from the IP routing table.
clear ip mfib counters	Deletes all per-route and global Multicast Forwarding Information Base (MFIB) counters.
clear ip mfib fastdrop	Deletes all fast-drop entries.

Note IP multicast routes can be regenerated in response to protocol events and as data packets arrive.

Summary

This topic summarizes the key points discussed in this lesson.

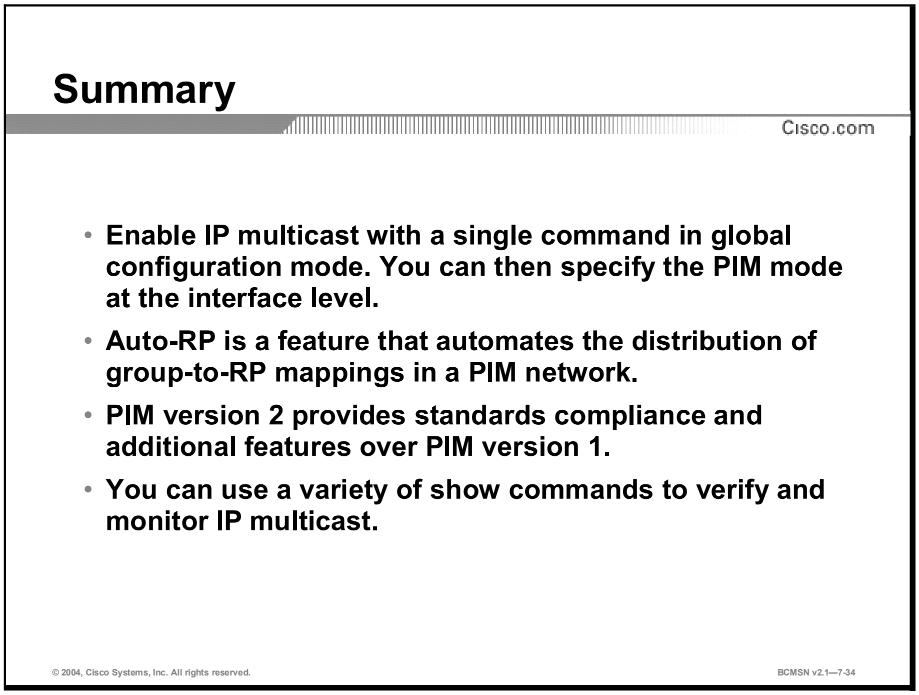

Summary

- **Enable IP multicast with a single command in global configuration mode. You can then specify the PIM mode at the interface level.**
- **Auto-RP is a feature that automates the distribution of group-to-RP mappings in a PIM network.**
- **PIM version 2 provides standards compliance and additional features over PIM version 1.**
- **You can use a variety of show commands to verify and monitor IP multicast.**

BCMSN v2.1—7-34

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which command correctly enables the preferred mode of PIM for IP multicast?

A) **ip pim dense-mode**

B) **ip pim sparse-mode**

C) **ip pim sparse-dense-mode**

D) **ip pim dense-sparse-mode**

Q2) Which command correctly advertises the IP address of Ethernet 0 as the RP for the administratively scoped groups?

A) **ip pim send-rp-announce ethernet0 scope 16 group-list 1 access-list 1 permit 239.0.0.0 0.255.255.255**

B) **ip pim send-rp-announce ethernet0 scope 16 group-list 1 access-list 1 permit 224.0.0.0 0.255.255.255**

C) **ip pim send-rp-discovery ethernet0 scope 16 group-list 1 access-list 1 permit 239.0.0.0 0.255.255.255**

D) **ip pim send-rp-discovery ethernet0 scope 16 group-list 1 access-list 1 permit 224.0.0.0 0.255.255.255**

Q3) What is the purpose of a bootstrap router in PIM version 2?

A) to ensure RP availability

B) to provide redundant RPs

C) to provide RP discovery and distribution

D) to configure a single RP for multiple groups

Q4) Which command correctly displays the contents of the IP multicast routing table?

A) **show ip mfib**

B) **show ip mroute**

C) **show mroute table**

D) **show ip pim interface**

Quiz Answer Key

Q1) C

Relates to: Enabling IP Multicast

Q2) A

Relates to: Configuring Auto-RP

Q3) C

Relates to: Configuring PIM Version 2

Q4) B

Relates to: Monitoring IP Multicast

Introducing Cisco IP Telephony

Overview

The flexibility and functionality of the Cisco Architecture for Voice, Video and Integrated Data (AVVID) network infrastructure provides a framework that permits rapid deployment of IP telephony applications.

Relevance

The Cisco AVVID framework, combined with multicast services and the Cisco IP telephony solution, provides universal transport for data, voice, and video applications today.

Objectives

Upon completing this lesson, you will be able to:

- Identify the network and device design considerations to support voice traffic
- Implement IP telephony on a switched network with auxiliary VLANs
- Explain how to implement voice in the campus network

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Network Design Issues for Voice
- Reviewing Auxiliary VLANs
- Best Practices for IP Telephony in the Campus
- Summary
- Quiz

Network Design Issues for Voice

IP telephony places strict requirements on the network infrastructure. The network must provide sufficient bandwidth and quick convergence after network failures or changes. This topic identifies the network and device design considerations to support voice traffic.

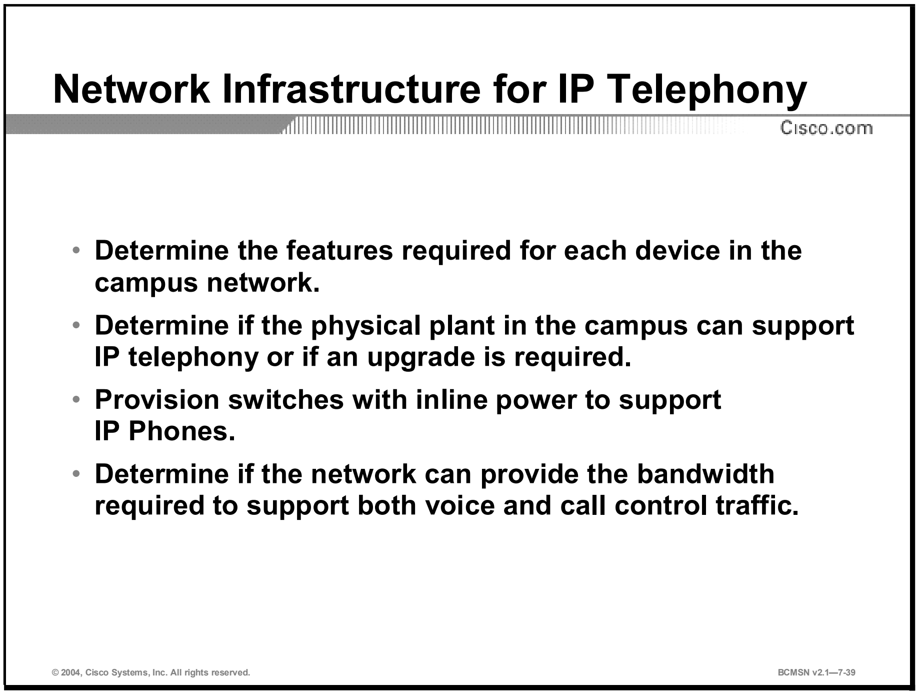

Network Infrastructure for IP Telephony

- **Determine the features required for each device in the campus network.**
- **Determine if the physical plant in the campus can support IP telephony or if an upgrade is required.**
- **Provision switches with inline power to support IP Phones.**
- **Determine if the network can provide the bandwidth required to support both voice and call control traffic.**

BCMSN v2.1—7-39

Most IP telephony installations are built on an existing network infrastructure, but the infrastructure may require enhancements. In the voice solution, it is critical that you prepare to allow voice traffic to have priority over all other traffic in the network. To design the infrastructure to support voice, complete these tasks:

- **Determine the features required for each device in the campus network.** IP phones require power. Most enterprises put IP telephony applications on a separate VLAN with prioritization. The infrastructure must support voice to ensure success.

- **Determine if the physical plant in the campus can support IP telephony or if an upgrade is required.** The wiring and cabling plant are critical for IP telephony. At a minimum, Category 5 cabling is critical for IP telephony.

- **Provision switches with inline power to support IP phones.** Many Cisco Catalyst switches provide inline power. Within a wiring closet, you can deploy a Catalyst Inline Power Patch Panel, which provides Fast Ethernet enhancements needed for multiservice networking while preserving customer investments in existing Catalyst switch equipment.

- **Determine if the network can provide the bandwidth required to support voice and call control traffic.** Bandwidth must consider both voice traffic and call control traffic. Consider both in your traffic engineering efforts. In the campus network, bandwidth provisioning requires careful planning of the LAN infrastructure so that the available bandwidth is always considerably higher than the load and so that there is no steady-state congestion over the LAN links. This ensures that the network is responsive to the offered traffic.

Note You should plan to work with a voice specialist to complete a traffic engineering analysis for the network.

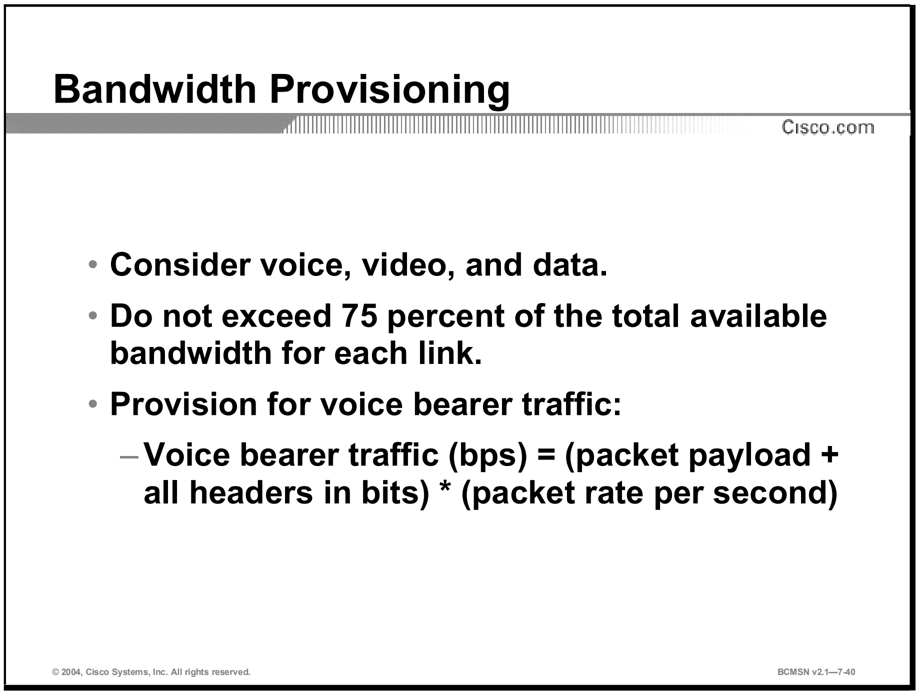

Bandwidth Provisioning

Cisco.com

- Consider voice, video, and data.
- Do not exceed 75 percent of the total available bandwidth for each link.
- Provision for voice bearer traffic:
  - Voice bearer traffic (bps) = (packet payload + all headers in bits) * (packet rate per second)

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—7-40

Properly provisioning the network bandwidth is a major component of designing a successful IP telephony network. You can calculate the required bandwidth by adding the bandwidth requirements for each major application, including voice, video, and data. This sum then represents the minimum bandwidth requirement for any given link, and it should not exceed approximately 75 percent of the total available bandwidth for the link.

From a traffic standpoint, an IP telephony call consists of these two parts:

- **Voice carrier stream:** This consists of Real-Time Transport Protocol (RTP) packets that contain the actual voice samples.
- **Call control signaling:** This consists of packets belonging to one of several protocols, according to the endpoints involved in the call; for example, H.323 or Media Gateway Control Protocol (MGCP). Call control functions are, for instance, those used to set up, maintain, tear down, or redirect a call.

Bandwidth provisioning must include not only the voice stream traffic but also the call control traffic.

A Voice over IP (VoIP) packet consists of the voice payload, IP header, UDP header, RTP header, and Layer 2 link header. Coder-decoder (codec) type (G.711, G.729, etc.) is configurable by device. However, G.729 does not support fax or modem traffic. The IP header is 20 bytes, the UDP header is 8 bytes, and the RTP header is 12 bytes. The link header varies in size according to the Layer 2 media used; Ethernet requires 14 bytes of header. The voice payload size and packetization period are device-dependent.

Use this formula to calculate the bandwidth that voice streams consume:

- (Packet payload + all headers in bits) * Packet rate per second; for example, 50 packets per second (pps) when using a 20-ms packet period

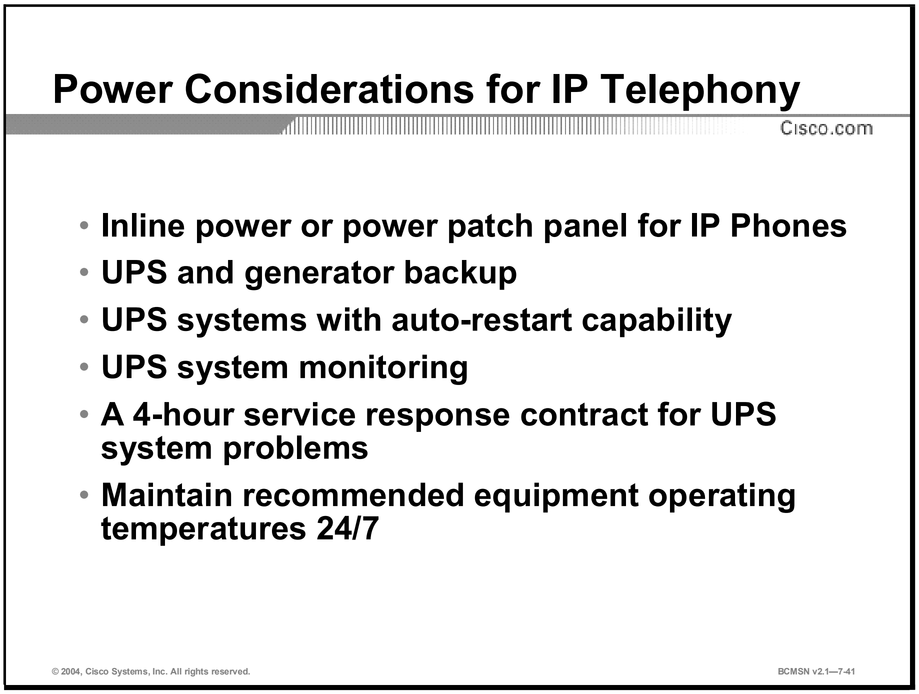

Power Considerations for IP Telephony

- Inline power or power patch panel for IP Phones
- UPS and generator backup
- UPS systems with auto-restart capability
- UPS system monitoring
- A 4-hour service response contract for UPS system problems
- Maintain recommended equipment operating temperatures 24/7

BCMSN v2.1—7-41

Power and environment can have an impact on availability that affects the IP telephony solution. Power can be supplied to the IP phones either directly from Catalyst switches with inline power capabilities or by inserting a Catalyst Inline Power Patch Panel for that purpose. Power is unique in that can affect an entire building or multiple buildings. This affect can have an impact on all devices in the IP telephony availability definition, including Building Distribution and Campus Backbone submodules, gateway, and Cisco CallManager all at once. The calculations, therefore, change from device-based to network-based. This can have a significant impact on theoretical availability, depending on the power protection strategy used.

Providing high availability based on the power protection strategy requires an uninterruptible power supply (UPS) system with a minimum battery life of 1 hour and a 4-hour response for UPS system failures, or a generator with an onsite service contract. That is, the high-availability IP telephony solution must include UPS and generator backup for all telephony devices. In addition, the organization should have UPS systems that have auto-restart capability and a service contract for 4-hour response to support the UPS system or generator.

IP telephony high-availability power and environment recommendations include the following:

- UPS and generator backup
- UPS systems with auto-restart capability
- UPS system monitoring
- A 4-hour service response contract for UPS system problems
- Maintain recommended equipment operating temperatures 24/7

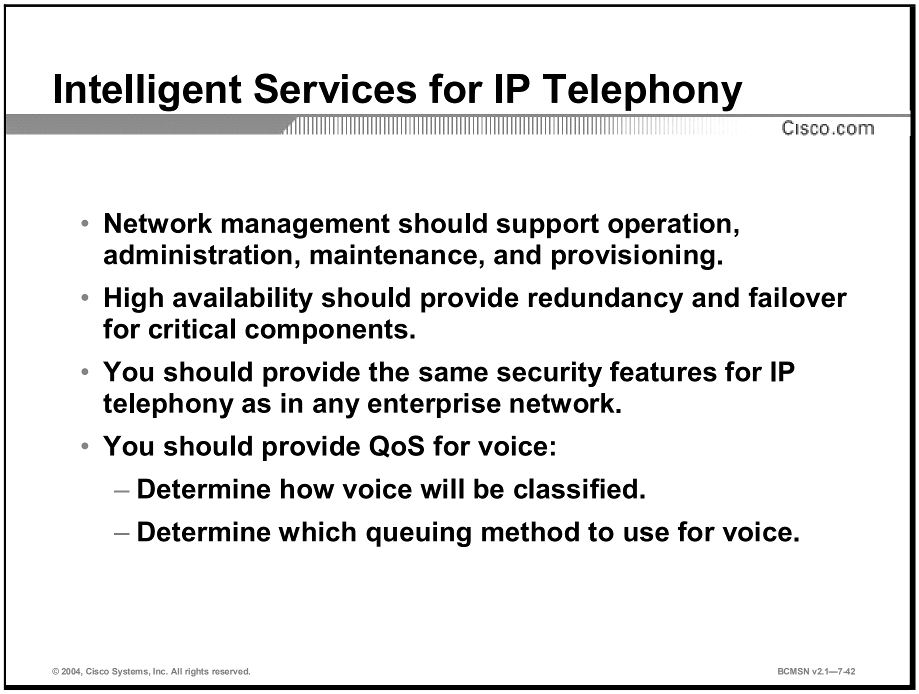

Intelligent Services for IP Telephony

Cisco.com

- Network management should support operation, administration, maintenance, and provisioning.
- High availability should provide redundancy and failover for critical components.
- You should provide the same security features for IP telephony as in any enterprise network.
- You should provide QoS for voice:
  - Determine how voice will be classified.
  - Determine which queuing method to use for voice.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—7-42

The network management, high availability, security, and quality of service (QoS) intelligent network services extend to incorporate voice-specific attributes.

Network Management

In traditional voice networks, there is a distinct set of voice management concepts and processes. The convergence of voice and data has brought about a similar merge of data network and voice-only management.

In fact, this merging of management tasks and processes is one of the key benefits of using a converged network as opposed to a dedicated voice-only network. However, it is still necessary to understand the traditional voice-only management concepts to relate the features available in that technology to the converged network management techniques.

High Availability

Cisco AVVID IP telephony is based on a distributed model for high availability. Cisco CallManager clusters support Cisco CallManager redundancy. The gateways must support the ability to “re-home” to a secondary Cisco CallManager in the event that a primary Cisco CallManager fails, thereby providing Cisco CallManager redundancy. This differs from call survivability in the event of a Cisco CallManager or network failure, when the call is routed to an alternate gateway, such as an MGCP gateway.

As with any network capability, you need to plan redundancy for critical components such as the Cisco CallManager and the associated gateway and infrastructure devices that support the voice network.

Security

The subject of securing voice communications has received even more visibility recently as network convergence becomes the accepted design model. With the advent of IP telephony, which uses IP data network devices for voice communication, the potential exists for malicious attacks on call-processing components and telephony applications.

To help safeguard against attacks, you should implement the same security precautions as in the rest of the enterprise network.

Securing the voice call-processing platform and installed applications is perhaps the most vital step in securing Cisco AVVID networks.

Every enterprise should have a predefined security policy for all devices, applications, and users to follow. The strictness of the security policy depends upon the level of caution required.

QoS

Voice, as a class of IP network traffic, has strict requirements concerning delay and delay variation (also known as “jitter”). Compared to most data, voice is relatively tolerant of loss. To meet the requirements for voice traffic, the Cisco AVVID IP telephony solution uses a wide range of IP QoS features, such as classification, queuing, congestion detection, traffic shaping, and compression.

The overall goal of QoS in the network is to be able to manage applications to determine which are less likely to be affected by loss, delay, and jitter. When a network becomes congested, some traffic will be delayed or even, at times, lost. The goal is to give critical applications a higher priority for service so that they are least likely to be delayed or dropped in times of congestion. In many converged networks, voice is the most critical application. In others, voice may opportunistically use bandwidth not required for data and fall back to the Public Switched Telephone Network (PSTN) in times of congestion.

Reviewing Auxiliary VLANs

The auxiliary VLAN feature provides automatic VLAN configuration for IP phones. This topic explains how to implement IP telephony on a switched network with auxiliary VLANs.

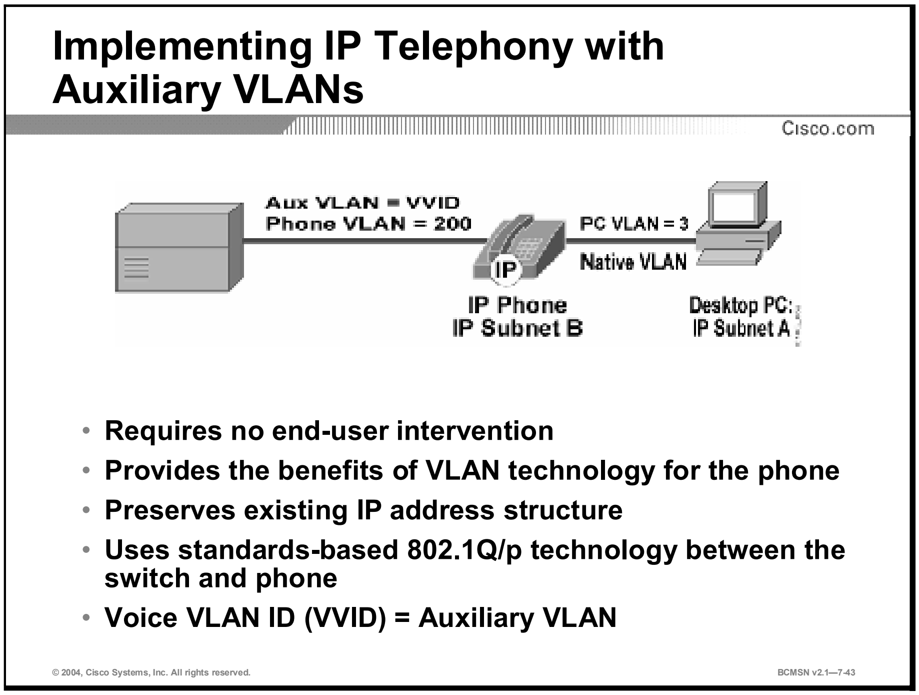

Implementing IP Telephony with Auxiliary VLANs

Cisco.com

- Requires no end-user intervention
- Provides the benefits of VLAN technology for the phone
- Preserves existing IP address structure
- Uses standards-based 802.1Q/p technology between the switch and phone
- Voice VLAN ID (VVID) = Auxiliary VLAN

BCMSN v2.1—7-43

Some Cisco Catalyst switches offer a unique feature called “auxiliary VLAN.” The auxiliary VLAN feature allows you to overlay a voice topology onto a data network. You can segment phones into separate logical networks, even though the data and voice infrastructure are physically the same.

The auxiliary VLAN feature places the phones into their own VLANs without any end-user intervention. Furthermore, these VLAN assignments can be seamlessly maintained, even if the phone is moved to a new location. The user simply plugs the phone into the switch, and the switch will provide the phone with the necessary VLAN information. By placing phones into their own VLANs, network administrators gain the advantages of network segmentation and control. Furthermore, network administrators can preserve their existing IP topology for the data end stations. IP phones can be easily assigned to different IP subnets using standards-based Dynamic Host Configuration Protocol (DHCP) operation.

With the phones in their own IP subnets and VLANs, network administrators can more easily identify and troubleshoot network problems. Additionally, network administrators can create and enforce QoS or security policies. With the auxiliary VLAN feature, Cisco enables network administrators to gain all the advantages of physical infrastructure convergence while maintaining separate logical topologies for voice and data terminals. This creates the most effective way to manage a multiservice network.

Auxiliary VLAN Implementation Guidelines

Follow these guidelines when implementing auxiliary VLANs:

- The IP phone and a device attached to the phone are in the same VLAN and must be in the same IP subnet if the following is true:
  - They use the same frame type
  - The phone uses 802.1p frames and the device uses untagged frames
  - The phone uses untagged frames and the device uses 802.1p frames
  - The phone uses 802.1Q frames and the auxiliary VLAN equals the native VLAN
- The IP phone and a device attached to the phone cannot communicate if they are in the same VLAN and subnet but use different frame types. This is because traffic between devices in the same subnet is not routed (routing would eliminate the frame type difference).
- You cannot use switch commands to configure the frame type used by traffic received from a device attached to the phone access port.

Best Practices for IP Telephony in the Campus

Deploying IP telephony in the enterprise campus requires the implementation of various features particular to each submodule. This topic discusses implementing IP telephony in the campus.

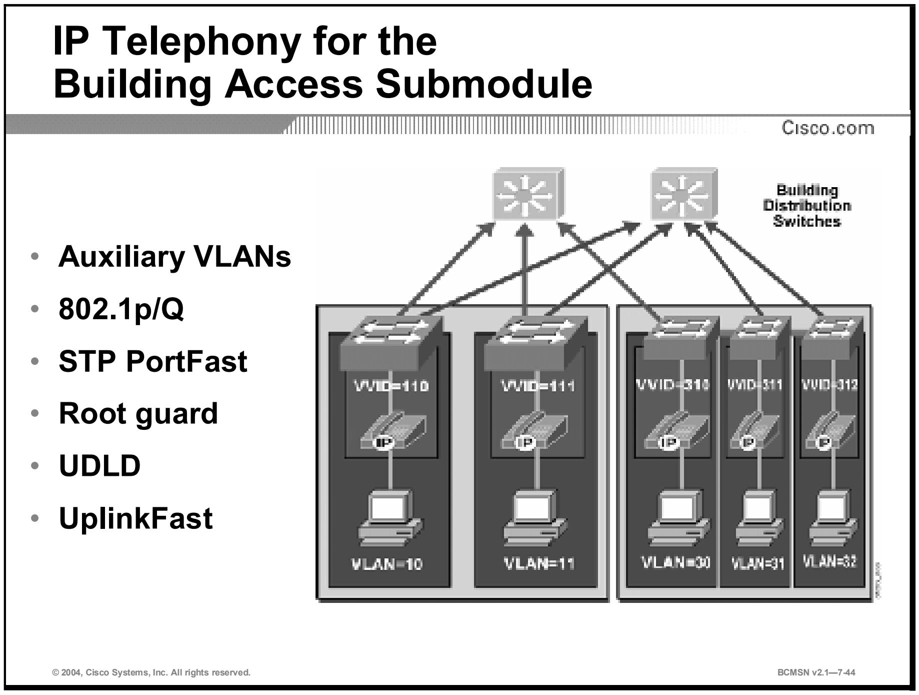

IP Telephony for the Building Access Submodule

Cisco.com

- Auxiliary VLANs
- 802.1p/Q
- STP PortFast
- Root guard
- UDLD
- UplinkFast

BCMSN v2.1—7-44

Within the Building Access submodule, implement these features to support IP telephony:

- Auxiliary VLANs
- 802.1p/Q
- Hardware support for multiple output queues
- Hardware support for in-line power to IP phones
- STP PortFast
- Root Guard
- Unidirectional Link Detection (UDLD)
- UplinkFast

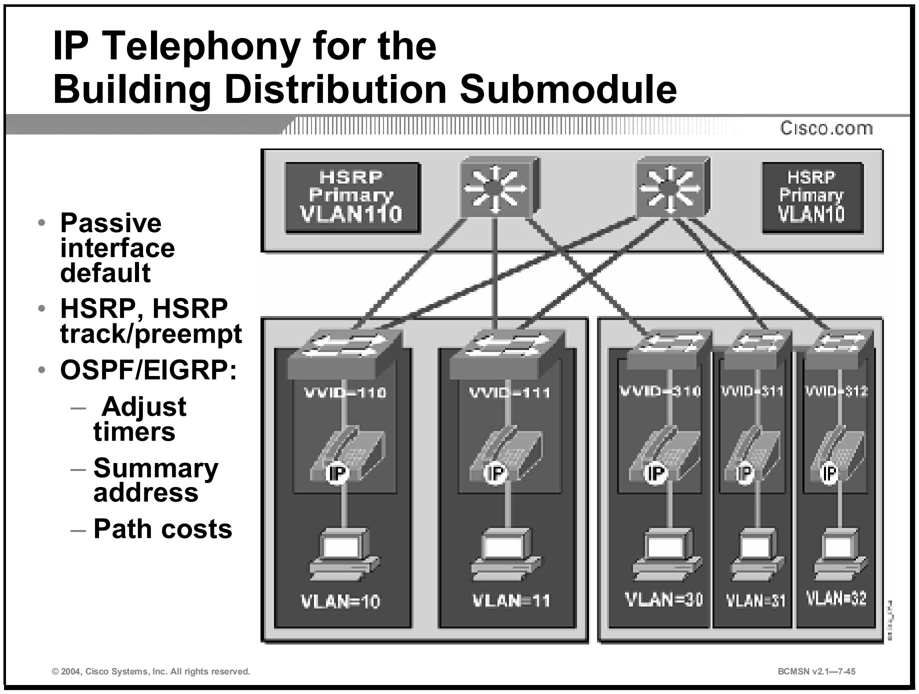

IP Telephony for the Building Distribution Submodule

Cisco.com

- Passive interface default
- HSRP, HSRP track/preempt
- OSPF/EIGRP:
  - Adjust timers
  - Summary address
  - Path costs

BCMSN v2.1—7-45

Within the Building Distribution submodule, implement the following features to support IP telephony:

- Passive interfaces as the default
- Layer 3 redundancy with Hot Standby Router Protocol (HSRP), HSRP track, and HSRP preempt
- OSPF or Enhanced Interior Gateway Routing Protocol (EIGRP) routing with adjusted timers, summary addresses, and path costs

Summary

This topic summarizes the key points discussed in this lesson.

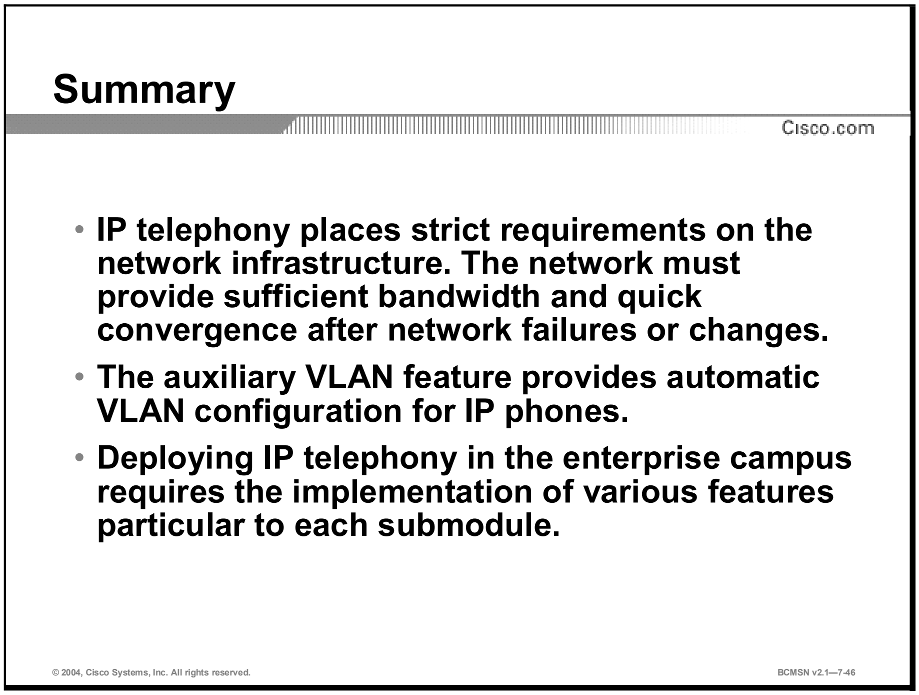

Summary

Cisco.com

- IP telephony places strict requirements on the network infrastructure. The network must provide sufficient bandwidth and quick convergence after network failures or changes.
- The auxiliary VLAN feature provides automatic VLAN configuration for IP phones.
- Deploying IP telephony in the enterprise campus requires the implementation of various features particular to each submodule.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—7-46

References

For additional information, refer to this resource:

- "Cisco IP Telephony Solution" at http://www.cisco.com/en/US/netsol/ns110/ns163/ns165/ns268/networking_solutions_package.html

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What is the formula for calculating voice bearer traffic?

A) packet payload * packet rate

B) (packet payload * packet rate) + headers

C) (packet payload + headers) * packet rate

D) (packet payload + headers) * packet rate * 75%

Q2) What are two benefits of the auxiliary VLAN feature of Catalyst switches? (Choose two.)

A) increased availability

B) easier network management

C) reduced bandwidth utilization

D) easier network troubleshooting

E) network segmentation and control

Q3) In which module or submodule should you implement Layer 3 redundancy to support IP telephony?

A) Server Farm

B) Building Access

C) Campus Backbone

D) Building Distribution

Quiz Answer Key

Q1) C

Relates to: Network Design Issues for Voice

Q2) D , E

Relates to: Reviewing Auxiliary VLANs

Q3) D

Relates to: Best Practices for IP Telephony in the Campus

Lesson Assessments

Overview

Use the lesson assessments here to test what you learned in this module. The correct answers and solutions are found in the Lesson Assessment Answer Key.

Outline

This section includes these assessments:

- Quiz 7-1: Examining IP Multicast in a Multilayer Switched Network
- Quiz 7-2: Configuring IP Multicast in a Multilayer Switched Network
- Quiz 7-3: Introducing Cisco IP Telephony

Quiz 7-1: Examining IP Multicast in a Multilayer Switched Network

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Explain how IP multicast operates on a multilayer switched network
- Describe the operation of Reverse Path Forwarding RPF
- Describe the operation and issues of PIM dense mode and PIM sparse mode
- Describe the operation of CGMP and of IGMP versions 1, 2, 3, and IGMPv3lite
- Explain the impact of CGMP and of IGMP versions 1, 2, 3, and IGMPv3lite on a Catalyst switch
- Define IGMP snooping

Quiz

Answer these questions:

Q1) Which type of multicast addresses are only to be used within a local group or organization?

A) GLOP addresses

B) limited scope addresses

C) globally scoped addresses

D) source specific multicast addresses

Q2) In which situation does an RPF check succeed?

A) if a multicast packet arrives on the interface leading back to the source

B) if a multicast packet arrives on an interface not leading back to the source

C) if a multicast packet arrives on an interface leading to the multicast group address

D) if a multicast packet arrives on an interface not leading to the multicast group address

Q3) In which environment is PIM sparse mode most useful, as opposed to PIM dense mode?

A) The volume of multicast traffic is high.

B) The stream of multicast traffic is constant.

C) There are few senders and many receivers.

D) There are few receivers in a multicast group.

Q4) What is the default value for maximum response time in an IGMP v2 membership query?

A) 1 second

B) 10 seconds

C) 1/10 second

D) 1/100 second

Q5) What does IGMP snooping look for in Layer 3 information?

A) IGMP join and leave messages

B) IGMP membership reports

C) IGMP membership queries

D) IGMP INCLUDE messages

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Quiz 7-2: Configuring IP Multicast in a Multilayer Switched Network

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Enable IP multicast on a Catalyst switch
- Configure IP multicast on a Catalyst switch
- Configure PIM version 2 on a Catalyst switch
- Monitor IP multicast on a Catalyst switch

Quiz

Answer these questions:

Q1) In which situation must you identify an RP for IP multicast?

A) PIM dense mode

B) PIM sparse mode

C) all PIM configurations

D) PIM sparse-dense mode

Q2) Which command correctly specifies a router as an RP mapping agent?

A) Switch#**ip pim send-rp-announce scope 16**

B) Switch#**ip pim send-rp-discovery scope 16**

C) Switch(config#**ip pim send-rp-announce scope 16**

D) Switch(config#**ip pim send-rp-discovery scope 16**

Q3) What is the purpose of the command **ip pim bsr-border** in configuring PIM version 2?

A) to identify the domain for a single RP

B) to identify the domain for multiple BSRs

C) to specify a domain to contain Auto-RP messages

D) to specify a boundary that bootstrap messages will not cross

Q4) Which command correctly displays information about interfaces configured for PIM?

A) **show interface pim**

B) **show pim interface**

C) **show ip pim interface**

D) **show ip interface pim**

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 75 percent or better.

Quiz 7-3: Introducing Cisco IP Telephony

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Identify the network and device design considerations to support voice traffic
- Implement IP telephony on a switched network with auxiliary VLANs
- Explain how to implement voice in the campus network

Quiz

Answer these questions:

Q1) Which network infrastructure feature should you provision on Catalyst switches specifically to support IP phones?

A) VLANs

B) inline power

C) multilayer switching

D) Spanning Tree Protocol

Q2) In what situation can an IP phone and a device attached to the phone not communicate to each other in an auxiliary VLAN implementation?

A) if they are in the same VLAN and subnet and use the same frame type

B) if they are in the same VLAN and subnet but use different frame types

C) if the phone uses 802.1Q frames and the auxiliary VLAN equals the native VLAN

D) if the phone uses 802.1p frames and the auxiliary VLAN equals the native VLAN

Q3) In which module or submodule should you implement hardware support for multiple output queues?

A) Server Farm

B) Building Access

C) Campus Backbone

D) Building Distribution

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 67 percent or better.

Lesson Assessment Answer Key

Quiz 7-1: Examining IP Multicast in a Multilayer Switched Network

Q1) B

Q2) A

Q3) D

Q4) B

Q5) A

Quiz 7-2: Configuring IP Multicast in a Multilayer Switched Network

Q1) B

Q2) D

Q3) D

Q4) C

Quiz 7-3: Introducing Cisco IP Telephony

Q1) B

Q2) B

Q3) B

Module 8

Implementing QoS in Multilayer Switched Networks

Overview

Enterprise and service provider networks transport varied applications and data, including high-quality video and critical application data. Bandwidth-intensive applications may stretch network capabilities and resources. Networks must provide secure, predictable, measurable, and sometimes guaranteed levels of service.

It is critical for the end-to-end network to achieve the required quality of service (QoS) by managing delay, delay variation (jitter), bandwidth, and packet loss parameters. QoS tools give network administrators the techniques to manage the network resources.

Upon completing this module, you will be able to:

- Describe each component of a Cisco QoS solution and explain how QoS solves quality issues on a multilayer switched network
- Configure QoS features on multilayer switched networks to provide optimal quality and bandwidth utilization for applications and data

Outline

The module contains these components:

- Examining the Cisco QoS Solution
- Configuring QoS in Multilayer Switched Networks
- Lesson Assessments

Examining the Cisco QoS Solution

Overview

Cisco IOS software provides a range of quality of service (QoS) tools that address the needs of voice, video, and data applications. IOS QoS technology lets you implement complex networks that predictably manage services to a variety of networked applications and traffic types.

Using the QoS tools and services in IOS software, you can design and implement networks that conform to either the Internet Engineering Task Force (IETF) Integrated Services (IntServ) model or the Differentiated Services (DiffServ) model. IOS QoS tools provide additional functionality such as classification and marking, congestion avoidance, and congestion management.

Relevance

Network administrators can enhance the performance and bandwidth utilization on a campus or WAN using QoS features. QoS features lead to efficient, predictable services for applications running on the network.

Objectives

Upon completing this lesson, you will be able to:

- Identify network requirements for QoS
- Describe the IntServ and DiffServ QoS architectures, and explain when to use each one
- Identify the classification and marking components of a Cisco QoS solution
- Identify the queuing and congestion management components of a Cisco QoS solution
- Identify the congestion avoidance components of a Cisco QoS solution
- Identify the traffic conditioning components of a Cisco QoS solution
- Identify the link efficiency components of a Cisco QoS solution
- Select QoS solutions for the campus network
- Summarize the key IOS software QoS features and explain when to use each feature

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Quality Issues on a Multilayer Switched Network
- QoS Service Models
- Classification and Marking
- Queuing and Congestion Management
- Congestion Avoidance
- Traffic Conditioning
- Link Efficiency Mechanisms
- QoS Within the Campus Network
- Summary of Key Cisco IOS Software QoS Categories and Features
- Summary
- Quiz

Quality Issues on a Multilayer Switched Network

Network managers know that queuing or “buffering,” not bandwidth, is the primary issue in the campus network. Almost any network can take advantage of QoS for optimum efficiency, whether it is a small corporate network, an Internet service provider (ISP), or an enterprise network. This topic identifies network requirements for QoS.

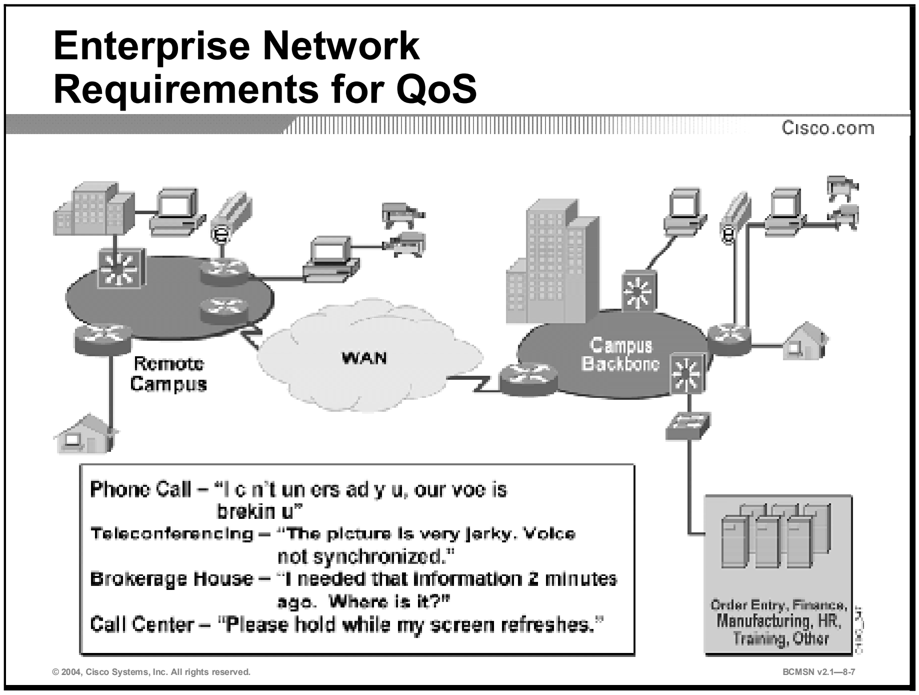

Enterprise Network Requirements for QoS

Phone Call – “I c n’t un ers ad y u, our voc is brekin u”
Teleconferencing – “The picture is very jerky. Voice not synchronized.”
Brokerage House – “I needed that information 2 minutes ago. Where is it?”
Call Center – “Please hold while my screen refreshes.”

BCMSN v2.1—8-7

QoS is defined as the application of features and functionality to actively manage and satisfy networking requirements of applications sensitive to loss, delay, and delay variation (jitter). QoS also guarantees the availability of bandwidth for critical application flows. QoS tools enable manageability and predictable service for a variety of networked applications and traffic types in a complex network.

The IOS QoS software provides these benefits:

- **Control over resources:** You have control over which network resources (bandwidth, equipment, wide-area facilities, and so on) are being used. For example, critical traffic such as voice or video may consume a link. QoS helps control the use of the resource (the link) by dropping low-priority packets, thereby denying them any use of the link.
- **More efficient use of network resources:** Using network analysis management and accounting tools, you can determine how traffic is handled, and which traffic experiences latency, jitter, and packet loss. If traffic is not handled optimally, you can use QoS tools to adjust the handling of that traffic.
- **Tailored services:** The control and visibility provided by QoS enables ISPs to offer carefully tailored grades of service differentiation to their customers. For example, a service provider can offer different service level agreements (SLAs) for a customer website that receives 3000 to 4000 hits per day, compared to another customer site that receives only 200 to 300 hits per day.

- **Coexistence of mission-critical applications:** QoS technologies make certain that mission-critical applications that are most important to a business receive the most efficient use of the network. Time-sensitive multimedia and voice applications require bandwidth and minimized delays, while other applications on a link receive fair service without interfering with mission-critical traffic.

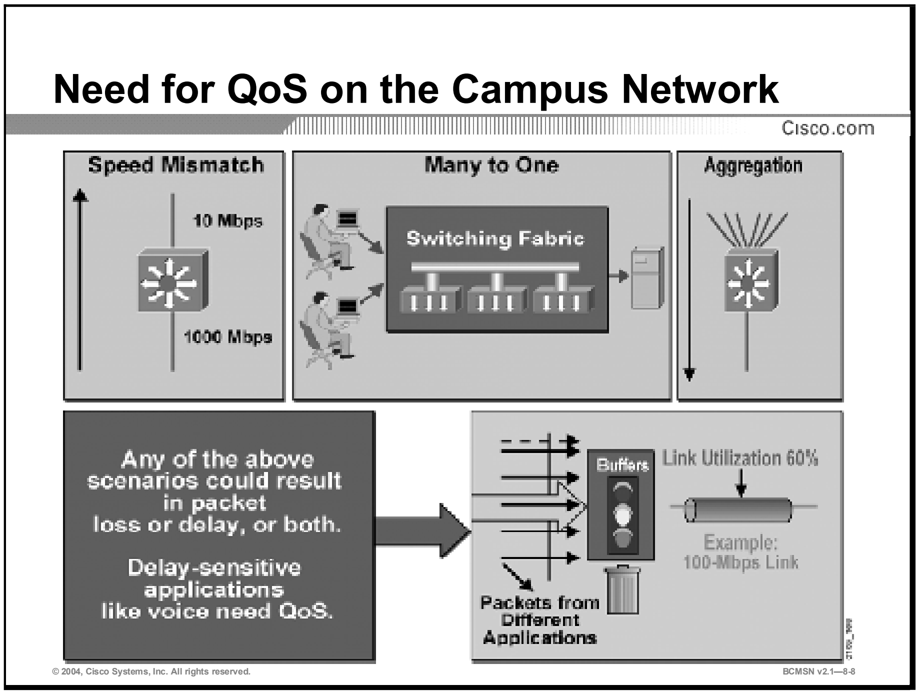

Need for QoS on the Campus Network

Speed Mismatch

10 Mbps

1000 Mbps

Many to One

Switching Fabric

Aggregation

Any of the above scenarios could result in packet loss or delay, or both.

Delay-sensitive applications like voice need QoS.

Buffers

Link Utilization 60%

Example: 100-Mbps Link

Packets from Different Applications

BCMSN v2.1—8-8

Switches today have many features, which may include high-performance backplanes, transfer of millions of switched packets per second (pps), and nonblocking characteristics. QoS provides guaranteed delivery of critical applications such as voice and video, even when a network runs suboptimally.

A switch may be the fastest switch in the world, but if you have speed mismatches, many-to-one switching fabrics, and aggregation, that switch will experience congestion. At times of congestion, if congestion management features are not in place, packets will be dropped. When packets are dropped, retransmissions typically occur. When retransmissions occur, the network load can increase. In networks that are already congested, this load increase can add to existing performance issues and potentially further degrade performance.

With converged networks, QoS is even more critical. Simply adding more buffers to a switch will not necessarily alleviate quality problems. First, you need to identify important traffic through classification techniques such as network-based application recognition (NBAR), and then implement techniques such as queuing to ensure delivery of high-priority traffic. Finally, you need to incorporate scheduling techniques to switch priority packets from queues to expedite their delivery.

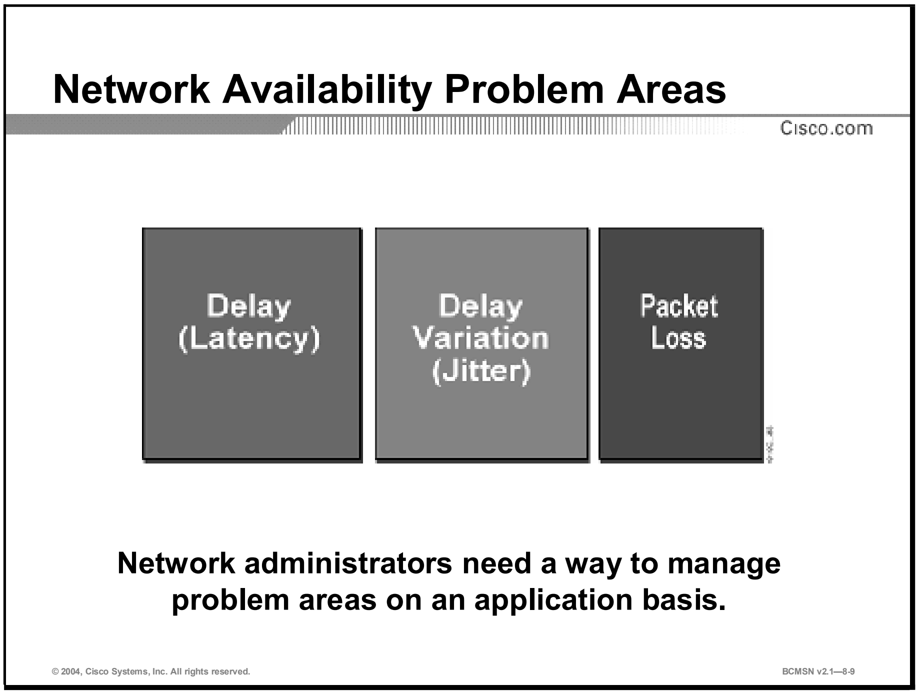

Network Availability Problem Areas

Cisco.com

Delay (Latency)

Delay Variation (Jitter)

Packet Loss

Network administrators need a way to manage problem areas on an application basis.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—8-9

An enterprise network may experience any of these network availability problems:

- **Delay:** Delay (or latency) is the amount of time that it takes a packet to reach the receiving endpoint after being transmitted from the sending endpoint. This time period is termed the "end-to-end delay," and can be broken into two areas: fixed network delay and variable network delay. Fixed network delay includes encoding and decoding time (for voice and video), as well as the amount of time required for the electrical and optical pulses to traverse the media en route to their destination. Variable network delay generally refers to network conditions, such as congestion, that may affect the overall time required for transit. In data networks, for example, these types of delay occur:
  - **Packetization delay:** The amount of time that it takes to segment data (if necessary), sample and encode signals (if necessary), process data, and turn the data into packets
  - **Serialization delay:** The amount of time that it takes to place the bits of a packet, encapsulated in a frame, onto the physical media
  - **Propagation delay:** The amount of time that it takes to transmit the bits of a frame across the physical wire
  - **Processing delay:** The time that it takes for a network device to take the frame from an input interface, place it into a receive queue, and then place it into the output queue of the output interface
  - **Queuing delay:** The amount of time that a packet resides in the output queue of an interface

- **Delay variation:** Delay variation (or jitter) is the difference in the end-to-end delay between packets. For example, if one packet requires 100 ms to traverse the network from the source endpoint to the destination endpoint, and the following packet requires 125 ms to make the same trip, then the delay variation is calculated as 25 ms.

  Each end station and Cisco network device in a voice or video conversation has a jitter buffer. Jitter buffers are used to smooth out changes in arrival times of data packets containing voice and video. A jitter buffer is dynamic and can adjust for changes in arrival times of packets. If you have instantaneous changes in arrival times of packets that are outside of the capabilities of a jitter buffer to compensate, you will have one of these situations:

  — A jitter buffer underrun, when arrival times between packets containing voice or video increase to the point where the jitter buffer has been exhausted and contains no packets to process the signal for the next piece of voice or video.

  — A jitter buffer overrun, when arrival times between packets containing voice or video decrease to the point where the jitter buffer cannot dynamically resize itself quickly enough to accommodate. When an overrun occurs, packets are dropped and voice quality is degraded.

- **Packet loss:** Packet loss is a measure of packets transmitted and received compared to the total number that were transmitted. Loss is expressed as the percentage of packets that were dropped. Tail drops occur when the output queue is full. These are the most common drops that can occur when a link is congested. Other types of drops (input, ignore, overrun, no buffer) are not as common but may require a hardware upgrade because they are usually a result of network device congestion.

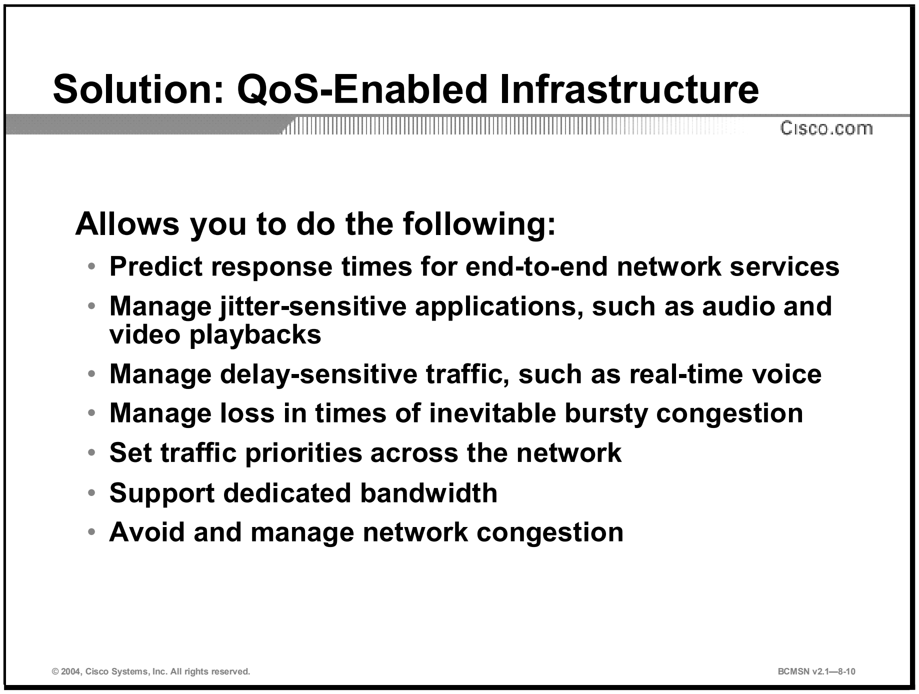

Solution: QoS-Enabled Infrastructure

Cisco.com

Allows you to do the following:

- **Predict response times for end-to-end network services**
- **Manage jitter-sensitive applications, such as audio and video playbacks**
- **Manage delay-sensitive traffic, such as real-time voice**
- **Manage loss in times of inevitable bursty congestion**
- **Set traffic priorities across the network**
- **Support dedicated bandwidth**
- **Avoid and manage network congestion**

BCMSN v2.1—8-10

Managing quality on the network is difficult because many applications deliver unpredictable bursts of traffic. For example, usage patterns for web, e-mail, and file transfer applications are virtually impossible to predict, yet network managers need to be able to support mission-critical applications even during peak periods.

QoS technology allows you to do the following:

- Predict response times for end-to-end network services
- Manage jitter-sensitive applications, such as audio and video playbacks
- Manage delay-sensitive traffic, such as real-time voice
- Manage loss in times of inevitable bursty congestion
- Set traffic priorities across the network
- Support dedicated bandwidth
- Avoid and manage network congestion

QoS includes these processes:

- **Classifying traffic:** Classification entails using a descriptor value to categorize a packet or frame within a specific group, to make it accessible for QoS handling on the network or within a network device such as a multilayer switch. Using classification, you can partition network traffic into multiple priority levels or classes of service. QoS signaling is another way in which classification is accomplished, but on an application-flow basis rather than a packet or frame basis.
- **Traffic shaping or conditioning:** Traffic shaping is used to create a traffic flow that limits the bandwidth potential of the flow. Shaping determines whether a packet is in or out of profile by comparing the traffic rate to the configured policer, which limits the bandwidth consumed by a flow of traffic. Traffic that exceeds the configured rate is buffered first in an attempt to minimize loss.

- **Policing:** Policing determines whether a packet is in or out of profile by comparing the traffic rate to the configured policer, which limits the bandwidth consumed by a flow of traffic. Policing is similar to shaping except that traffic that exceeds the configured rate is not buffered and normally is discarded.
- **Marking:** Marking is the process of setting or changing a prioritization value of a frame or packet.
- **Queuing:** Queuing evaluates the prioritization value and the configured policer and determines in which of the egress queues to place the packet.
- **Scheduling:** Scheduling services the egress (transmit) queues based on the sharing and shaping configuration of the egress (transmit) port.
- **Dropping:** Dropping is used to drop packets in order to take advantage of TCP transmission control and retransmission (windowing), which selectively drops some packets in an attempt to avoid dropping large numbers of packets. Note that this process works only with connection-oriented services that maintain sessions. User Datagram Protocol (UDP) or other connectionless traffic does not react properly to dropped packets, and may even increase the amount of traffic being transmitted as a result of packet loss.
- **Forwarding:** There are several supported forwarding mechanisms, such as process switching, fast switching, Cisco Express Forwarding (CEF) switching, and so on.

You can configure QoS features throughout a network to provide for end-to-end QoS delivery. The following three components are needed to deliver QoS across a heterogeneous network:

- QoS within a single network element, which includes queuing, scheduling, and traffic-shaping features
- QoS signaling techniques for coordinating QoS from end to end between network elements
- QoS policing and management functions to control and administer end-to-end traffic across a network

QoS Service Models

The two QoS architectures that you use in IP networks when designing a QoS solution are the Integrated Services (IntServ) and Differentiated Services (DiffServ) models. This topic describes the features of the IntServ and DiffServ models, and when to use each.

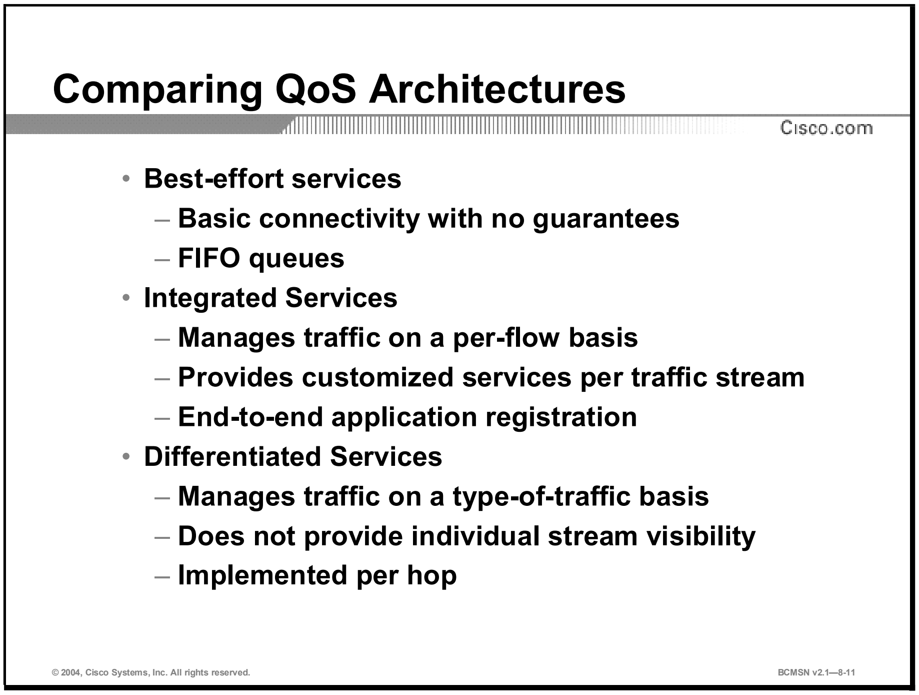

Comparing QoS Architectures

Cisco.com

- Best-effort services
  - Basic connectivity with no guarantees
  - FIFO queues
- Integrated Services
  - Manages traffic on a per-flow basis
  - Provides customized services per traffic stream
  - End-to-end application registration
- Differentiated Services
  - Manages traffic on a type-of-traffic basis
  - Does not provide individual stream visibility
  - Implemented per hop

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—8-11

A service model, also called a level of service, describes a set of end-to-end QoS capabilities. End-to-end QoS is the ability of the network to deliver the level of service required by specific applications from the source to the destination.

QoS service models differ from each other in how they enable applications to send data and in how the network attempts to deliver that data within the specified level of service. It is important to note that, regardless of the implementation, QoS does not create bandwidth. QoS tools allow you to use and manage the available bandwidth with respect to a given priority, classification, or type of traffic.

The services differ in their level of QoS strictness, which describes how tightly the service can be bound by specific bandwidth, delay, jitter, and loss characteristics.

Three basic levels of end-to-end QoS can be provided across a heterogeneous network:

- **Best-effort services:** Also known as lack of QoS, best-effort service is basic connectivity with no guarantees. This level of service is often characterized by FIFO queues, which have no differentiation between flows.
- **IntServ (also called "hard QoS"):** This level of service involves an absolute reservation of network resources for specific traffic from end to end. It is provided through the QoS tools called Resource Reservation Protocol (RSVP) and queuing.
- **DiffServ (also called "soft QoS"):** Some traffic classes are treated better than the rest (faster handling, more average bandwidth, and lower average loss rate). This level of

service represents a statistical preference, not a hard and fast guarantee of any particular traffic flow. DiffServ is provided by queuing and congestion avoidance mechanisms.

Deciding which type of service is appropriate to deploy in the network depends on the following:

- **Application supported or problem being solved:** Each of the three types of service is appropriate for certain applications. The hierarchy of service levels does not imply that an enterprise must migrate to differentiated and then to guaranteed service. A differentiated service, or even a best-effort service, may be appropriate, depending on the application requirements.
- **Speed to upgrade the infrastructure:** There is a natural upgrade path from the technology needed to provide differentiated services to the technology needed to provide guaranteed services.
- **Cost:** The cost of implementing and deploying IntServ is likely to be more than that for a differentiated service.

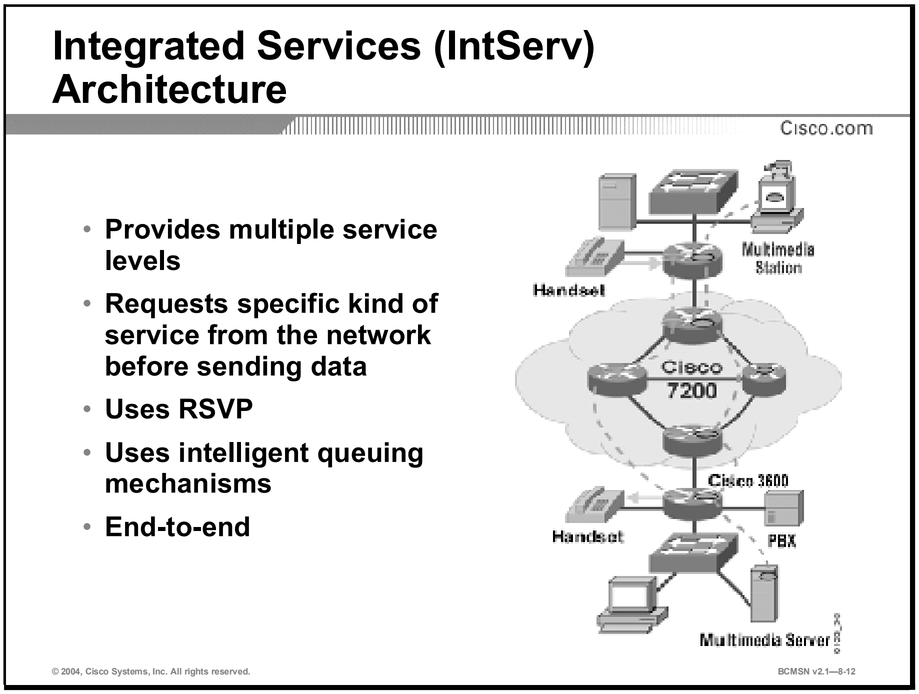

Integrated Services (IntServ) Architecture

Cisco.com

- Provides multiple service levels
- Requests specific kind of service from the network before sending data
- Uses RSVP
- Uses intelligent queuing mechanisms
- End-to-end

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—8-12

The IntServ model (defined in RFC 1633) was introduced to guarantee a predictable behavior of the network for applications. IntServ provides multiple services that can accommodate multiple QoS requirements. IntServ is implemented through the use of RSVP, enabled at both the endpoints and the network devices in between. The RSVP-enabled application requests a specific kind of service from the RSVP-enabled network before it sends data. Explicit signaling via RSVP facilitates the request. Then, the application informs the network of its traffic profile and requests a particular kind of service that can encompass its bandwidth and delay requirements. The application is expected to send data only after it gets a confirmation from the network. It is also expected to send data that lies within its described traffic profile.

The RSVP-enabled network performs admission control based on information from the requesting host application and available network resources. It either admits or rejects the application request for bandwidth. The network commits to meeting the QoS requirements of the application as long as the specific traffic flow for which the request was made remains within the profile specifications. Each flow of traffic must go through the admission control process.

You can centralize admission control using Common Open Policy Service (COPS) at a policy decision point (PDP). COPS provides these benefits when used with RSVP:

- Centralized management of services
- Centralized admission control and authorization of RSVP flows
- Increased scalability of RSVP-based QoS solutions

RSVP and IntServ offer these benefits:

- Explicit resource admission control (end-to-end)
- Per-request policy admission control (authorization object, policy object)
- Signaling of dynamic port numbers

RSVP and IntServ have these drawbacks:

- Continuous signaling due to stateless architecture
- Lack of scalability

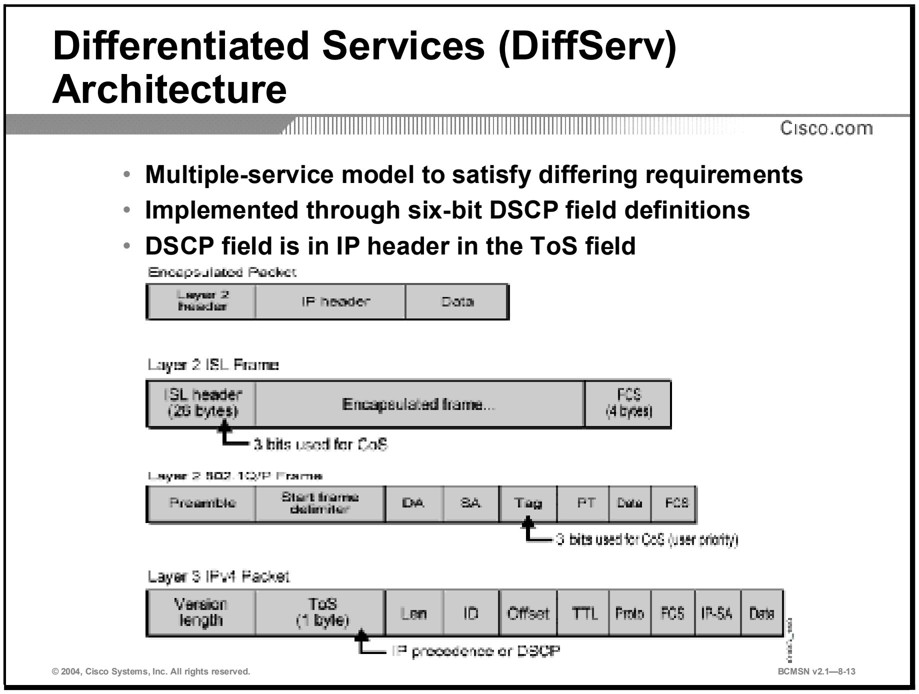

Differentiated Services (DiffServ) Architecture

Cisco.com

- **Multiple-service model to satisfy differing requirements**
- **Implemented through six-bit DSCP field definitions**
- **DSCP field is in IP header in the ToS field**

BCMSN v2.1—8-13

The DiffServ model is a multiple service-level model that can satisfy differing QoS requirements. However, unlike in the IntServ model, an application using DiffServ does not explicitly signal the network devices before sending data. DiffServ is a QoS implementation technique that is tailored for modern networks and the solutions that they depend on. DiffServ reassigns bits in the type of service (ToS) field of an IP packet. DiffServ uses differentiated services code points (DSCPs) as the QoS priority descriptor value and supports 64 classifications.

For a differentiated service, the network tries to deliver a particular kind of service based on the QoS descriptor specified in each IP packet header on a per-hop basis, rather than per-traffic flow as in IntServ.

You can use the DiffServ model for the same mission-critical applications as IntServ and to provide end-to-end QoS. Each packet is forwarded according to priorities designated within the packet on a per-device, per-hop basis. Typically, this service model is preferred within the Campus Backbone submodule because it is appropriate for aggregate flows. Complex traffic classification and conditioning are performed at network edges (Building Access submodule or the WAN module) resulting in per-packet QoS handling within the Campus Backbone submodule, where performance and scalability are the primary requirements.

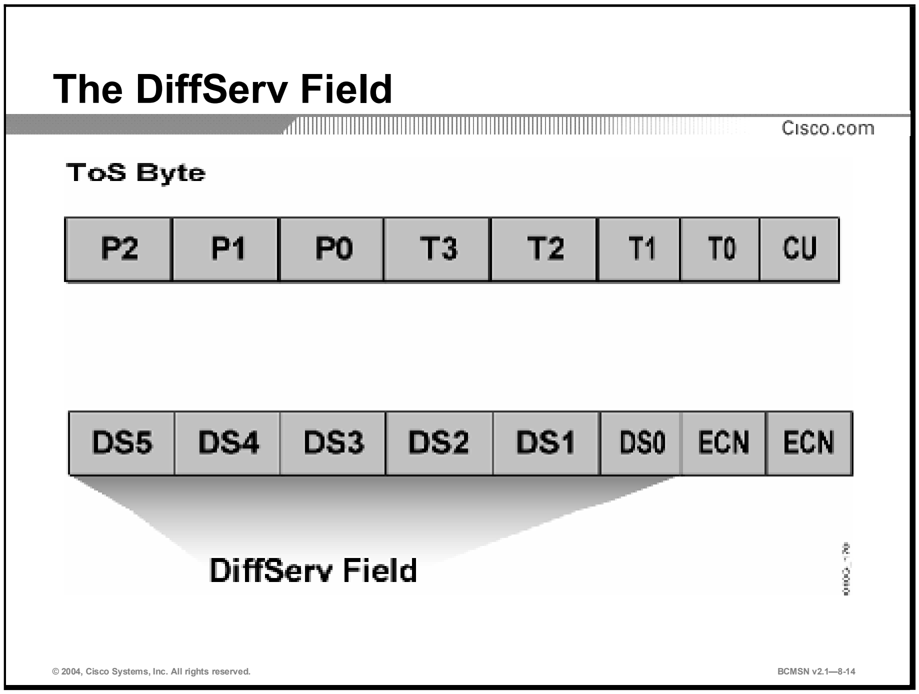

The DiffServ Field

Cisco.com

BCMSN v2.1—8-14

Switches at the edge of the network identify packets based on the IP precedence or on the DSCP fields in the header. Network devices that support DiffServ use DSCP in the IP header to select a per-hop behavior (PHB) for a packet.

The six most significant bits of the ToS byte form the DiffServ field. The last two bits are used as Early Congestion Notification (ECN) bits. IP precedence uses three bits, while DSCP, an extension of IP precedence, uses six bits to select the PHB for the packet at each network node.

The figure shows a comparison between the ToS byte with IP precedence and with the DiffServ field.

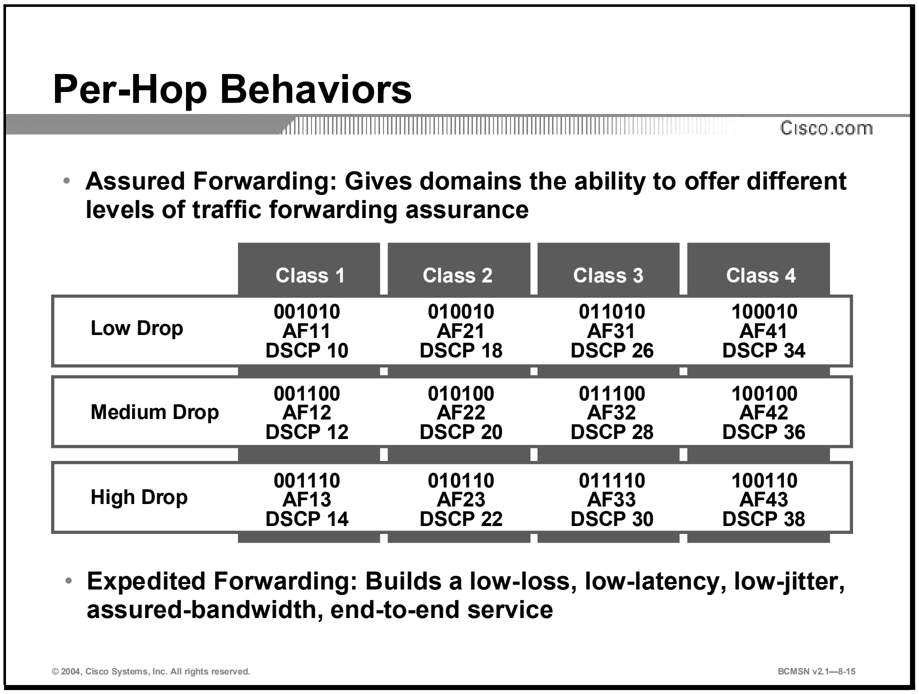

Per-Hop Behaviors

Cisco.com

- **Assured Forwarding: Gives domains the ability to offer different levels of traffic forwarding assurance**

	Class 1	Class 2	Class 3	Class 4
Low Drop	001010 AF11 DSCP 10	010010 AF21 DSCP 18	011010 AF31 DSCP 26	100010 AF41 DSCP 34
Medium Drop	001100 AF12 DSCP 12	010100 AF22 DSCP 20	011100 AF32 DSCP 28	100100 AF42 DSCP 36
High Drop	001110 AF13 DSCP 14	010110 AF23 DSCP 22	011110 AF33 DSCP 30	100110 AF43 DSCP 38

- **Expedited Forwarding: Builds a low-loss, low-latency, low-jitter, assured-bandwidth, end-to-end service**

BCMSN v2.1—8-15

The Assured Forwarding (AF) PHB group is a way for a service provider DiffServ domain to offer different levels of forwarding assurances for IP packets received from a customer DiffServ domain. There are four AF classes, AF1x through AF4x. Within each class, there are three drop probabilities. Depending on the policy of a given network, packets can be selected for a PHB based on required throughput, delay, jitter, or loss, or according to priority of access to network services.

AF PHB is nearly equivalent to the controlled load service available in the IntServ model. AF PHB defines a method by which traffic classes can be given different forwarding assurances. For example, network traffic can be divided into the following classes:

- **Gold:** 50 percent of the available bandwidth
- **Silver:** 30 percent of the available bandwidth
- **Bronze:** 20 percent of the available bandwidth

The table illustrates the DSCP coding for specifying the AF class with the probability. Bits 0, 1, and 2 define the class; bits 3 and 4 specify the drop probability; bit 5 is always 0.

	Class 1	Class 2	Class 3	Class 4
Low Drop	001010 AF11 DSCP 10	010010 AF21 DSCP 18	011010 AF31 DSCP 26	100010 AF41 DSCP 34
Medium Drop	001100 AF12 DSCP 12	010100 AF 22 DSCP 20	011100 AF32 DSCP 28	100100 AF42 DSCP 36
High Drop	001110 AF13 DSCP 14	010110 AF23 DSCP 22	011110 AF33 DSCP 30	100110 AF43 DSCP 38

Expedited Forwarding (RFC 2598) defines the Expedited Forwarding (EF) PHB. The EF PHB can be used to build a low-loss, low-latency, low-jitter, assured-bandwidth, end-to-end service through DiffServ domains. Such a service appears to the endpoints like a point-to-point connection. This service has also been described as premium service.

When implemented in a DiffServ network, EF PHB provides a premium service.

Classification and Marking

Classification mechanisms distinguish a frame or packet with a specific priority or predetermined criteria. Marking identifies which frames or packets are processed to meet a specific level of service or QoS policy for end-to-end service. This topic helps you select the classification and marking components of a Cisco QoS solution, given specific quality and application requirements.

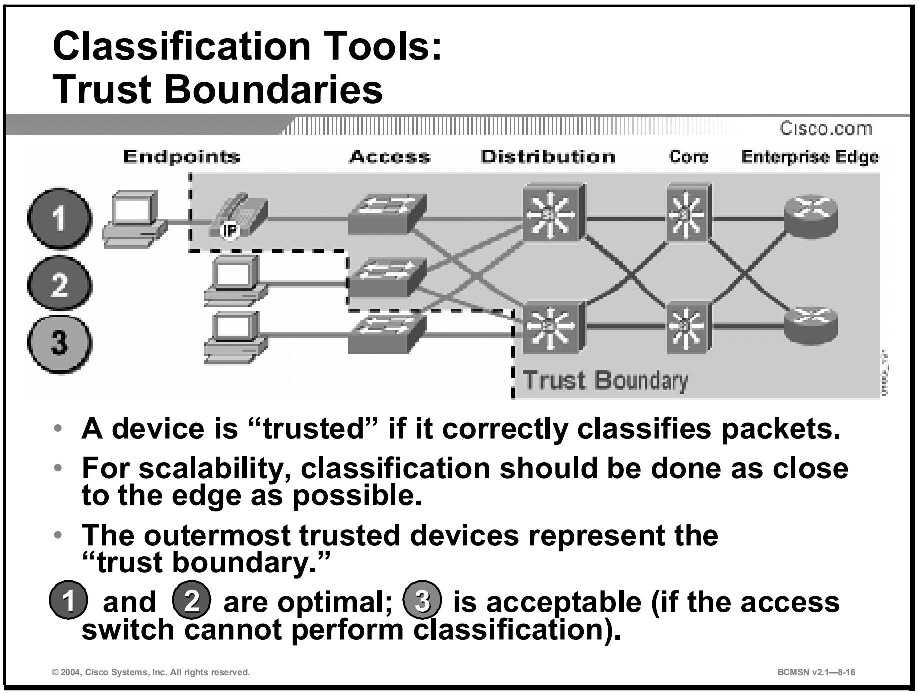

Classification Tools: Trust Boundaries

- A device is "trusted" if it correctly classifies packets.
- For scalability, classification should be done as close to the edge as possible.
- The outermost trusted devices represent the "trust boundary."
- ① and ② are optimal; ③ is acceptable (if the access switch cannot perform classification).

BCMSN v2.1—8-16

The first task of a QoS policy is to identify the traffic that is to be treated with preference. Classification entails using a traffic descriptor to categorize a frame within a specific group to define that frame and make it accessible for QoS handling on the network on a per-device basis. Using classification, you can partition network traffic into multiple priority levels or classes of service. Common methods of identifying types of traffic include access control lists (ACLs) and NBAR.

The location where classifications are accepted (or rejected) is referred to as the "trust boundary." Trust port configurations establish a trust boundary that subsequent network devices or elements in the network will enforce. Current methods of marking traffic with its classification allow you to set information in Layer 2, 3, or 4 headers. The network administrator sets trust boundaries based on the following:

- **Layer 2 parameters:** 802.1Q class of service (CoS) bits, MAC address, Multiprotocol Label Switching (MPLS), ATM cell loss priority (CLP) bit, Frame Relay discard eligible (DE) bit, ingress interface
- **Switching:** Ingress interface
- **Layer 3 parameters:** IP precedence, DSCP, QoS group, IP address, ingress interface
- **Layer 4 parameters:** TCP or UDP ports, ingress interface
- **Layer 7 parameters:** application signatures, ingress interface

Best-practice design recommendations are to identify and mark traffic as close to the source of the traffic as possible.

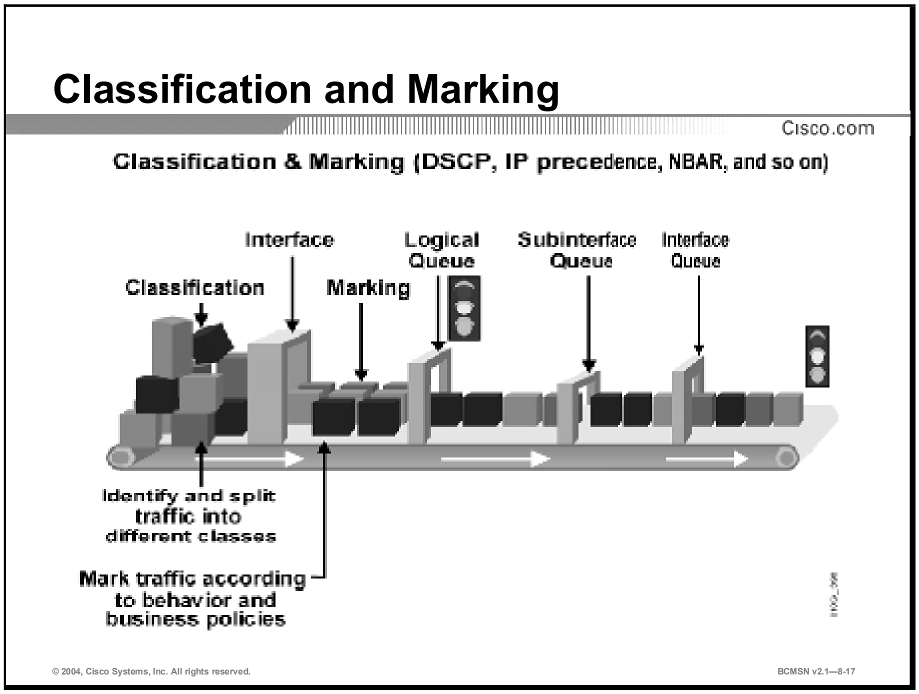

Classification and Marking

Cisco.com

Classification & Marking (DSCP, IP precedence, NBAR, and so on)

Interface
Logical Queue
Subinterface Queue
Interface Queue
Classification
Marking
Identify and split traffic into different classes
Mark traffic according to behavior and business policies

BCMSN v2.1—8-17

Packet classification features provide the capability to partition network traffic into multiple priority levels or classes of service. For example, by using the three precedence bits in the ToS field of the IP packet header (two of the values are reserved for other purposes), you can categorize packets into a limited set of up to six (0 to 5) traffic classes. After you classify packets, you can use other QoS features to assign the appropriate traffic handling policies, including congestion management, bandwidth allocation, and delay boundaries for each traffic class.

For non-IP traffic, you have these classification options:

- Use the port default. If the packet is a non-IP packet, assign the default port DSCP value to the incoming packet.
- Trust the CoS value in the incoming frame (configure the port to trust CoS). Then use the configurable CoS-to-DSCP map to generate the internal DSCP value. Layer 2 Inter-Switch Link (ISL) frame headers carry the CoS value in the three least significant bits of the 1-byte User field. Layer 2 802.1Q frame headers carry the CoS value in the three most significant bits of the Tag Control Information (TCI) field. CoS values range from 0 for low priority to 7 for high priority. If the frame does not contain a CoS value, assign the default port CoS to the incoming frame.
- The trust DSCP configuration applies only to IP traffic. If you configure a port with trust DSCP and non-IP traffic is received, the switch assigns the default port DSCP.

Most QoS mechanisms within the IOS software include some type of classification. Some mechanisms classify packets automatically; some require manual configuration. The table lists configuration options for various QoS mechanisms.

QoS Mechanism	Configuration Options
Committed access rate (CAR) Class-based policing	Access lists Rate limit access list QoS group DSCP
Policy-based routing (PBR)	Route map
Priority queuing (PQ) and custom queuing (CQ)	Access list Packet size Input interface Protocol
All mechanisms available	Class map that can use another class map, access list, protocol (including NBAR), input interface, source or destination MAC address, IP precedence, DSCP, QoS group, MPLS experimental bits, and so on

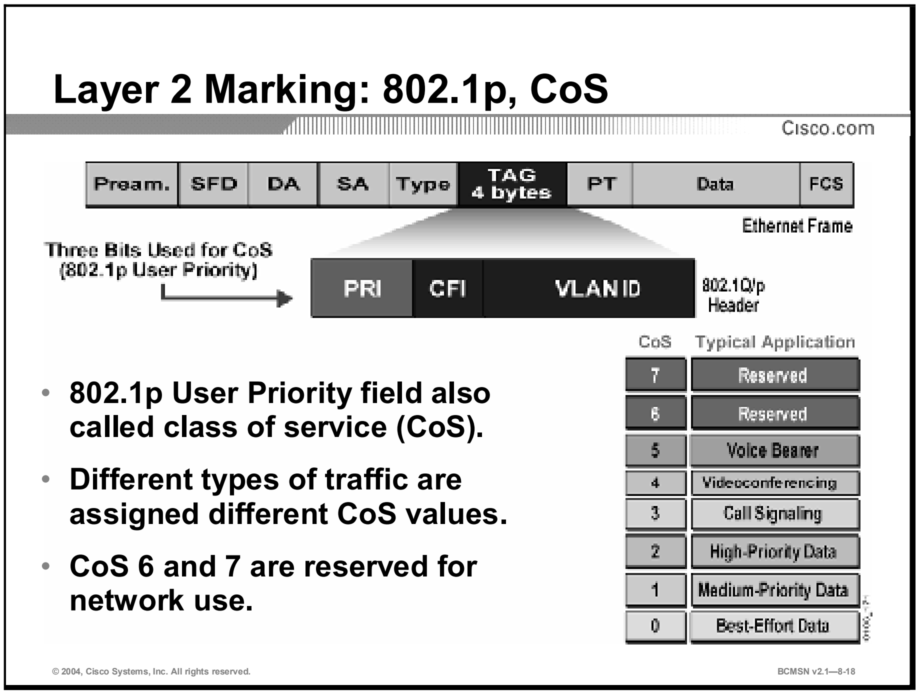

Layer 2 Marking: 802.1p, CoS

- 802.1p User Priority field also called class of service (CoS).
- Different types of traffic are assigned different CoS values.
- CoS 6 and 7 are reserved for network use.

CoS	Typical Application
7	Reserved
6	Reserved
5	Voice Bearer
4	Videoconferencing
3	Call Signaling
2	High-Priority Data
1	Medium-Priority Data
0	Best-Effort Data

BCMSN v2.1—8-18

With a Layer 2 switching engine, Layer 2 QoS supports classification using Layer 2 destination MAC addresses, and VLANs and marking using Layer 2 CoS values on the ingress port or interface. Classification and marking at Layer 2 does not use or set Layer 3 IP precedence or DSCP values.

Layer 2 QoS includes the following:

- **Input queue scheduling:** When the frame enters the port, it can be assigned to one of a number of port-based queues prior to being scheduled for switching to an egress port. Typically, multiple queues are used where different traffic requires different service levels, or where switch latency must be kept to a minimum. For instance, IP-based video and voice data requires low latency, so there may be a need to switch this data prior to switching other data.
- **Classification:** The process of classification involves inspecting different fields in the Ethernet Layer 2 header to assist in determining the level of service that will be applied to the frame as it transits the switch.
- **Policing:** Policing is the process of inspecting an Ethernet frame to see if it has exceeded a predefined rate of traffic within a certain time frame (typically, this time frame is a fixed number internal to the switch). If that frame is determined to be out-of-profile (that is, it is part of a data stream in excess of the predefined rate limit), it can either be dropped or the CoS value can be marked down.
- **Output queue scheduling:** The switch will place the Ethernet frame into an appropriate outbound (egress) queue for switching. The switch will perform buffer management on this queue by ensuring that the buffer does not overflow.

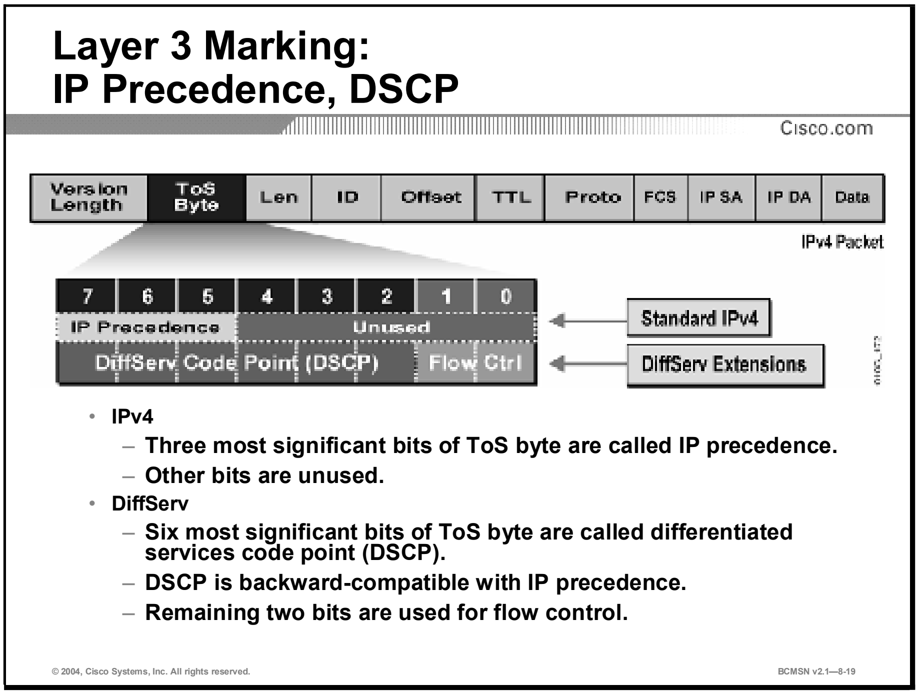

Layer 3 Marking: IP Precedence, DSCP

Cisco.com

- IPv4
  - Three most significant bits of ToS byte are called IP precedence.
  - Other bits are unused.
- DiffServ
  - Six most significant bits of ToS byte are called differentiated services code point (DSCP).
  - DSCP is backward-compatible with IP precedence.
  - Remaining two bits are used for flow control.

BCMSN v2.1—8-19

With a Layer 3 switching engine, QoS supports classification, marking, and policing using IP, Internetwork Packet Exchange (IPX), and MAC ACLs. ACLs contain access control entries (ACEs) that specify Layer 2, 3, and 4 classification criteria, a marking rule, and policing rules. Marking sets the Layer 3 IP precedence or DSCP values and the Layer 2 CoS value to either received or configured Layer 2 or Layer 3 values. Policing uses bandwidth limits to either drop or mark nonconforming traffic.

Queuing and Congestion Management

Congestion management features control congestion when it occurs. One way that network elements handle an overflow of arriving traffic is to use a queuing algorithm to sort the traffic, and then determine a method of prioritizing it onto an output interface. Each queuing algorithm solves a specific type of network traffic condition and has a particular effect on network performance. This topic explains and helps you select congestion management components of a Cisco QoS solution for a multilayer switched environment.

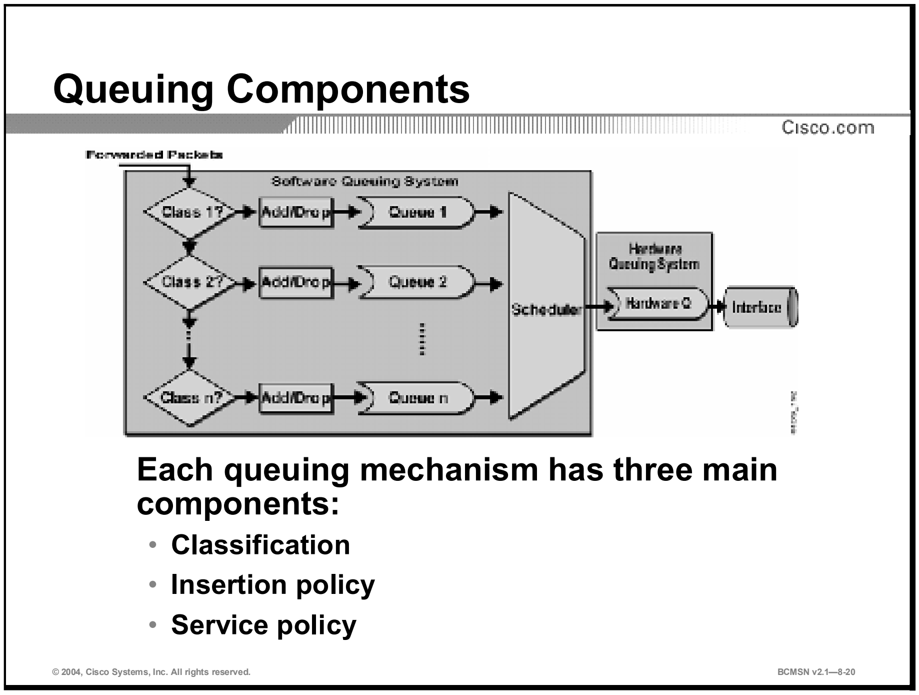

Queuing Components

Cisco.com

Each queuing mechanism has three main components:

- **Classification**
- **Insertion policy**
- **Service policy**

BCMSN v2.1—8-20

The figure illustrates the actions that a QoS-enabled switch must take before transmitting a frame. It takes these actions:

- Most queuing mechanisms include classification of packets.
- After a packet is classified, a network device must determine whether it can put the packet into the queue or it has to drop the packet. By default, queuing mechanisms will drop a packet only if the corresponding queue is full (tail drop).
- If the packet is queued, it is placed into the queue for that particular class.
- Packets are then taken from the individual per-class software queues and put into a FIFO hardware or transmit (Tx) queue.

Queuing methods differ in these ways:

- **Classification options:** Some mechanisms classify packets automatically while other mechanisms require manual configuration of classification.
- **Insertion policy:** Most queuing mechanisms, if configured, use the tail-dropping scheme.
- **Scheduling policy:** This component is the most important part of every queuing mechanism because it determines the order in which the packets will leave the switch.

These queuing mechanisms directly relate to the DiffServ model and specifically to multilayer switching:

- EF PHB implementations:
  - PQ
  - IP Real-Time Transport Protocol (RTP) prioritization
  - Class-based low latency queuing (CBLLQ)
- AF PHB implementations:
  - Class-based weighted fair queuing (CBWFQ), four classes with weighted random early detection (WRED) within each class
  - Optionally CQ (does not support differentiated dropping)

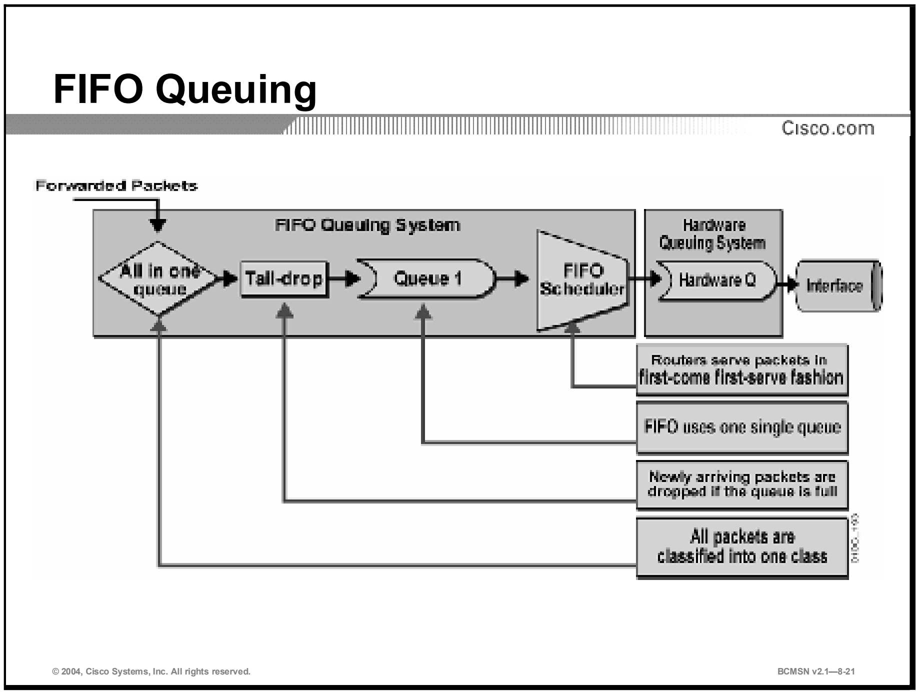

FIFO Queuing

Cisco.com

Forwarded Packets
FIFO Queuing System
All in one queue
Tail-drop
Queue 1
FIFO Scheduler
Hardware Queuing System
Hardware Q
Interface
Routers serve packets in first-come first-serve fashion
FIFO uses one single queue
Newly arriving packets are dropped if the queue is full
All packets are classified into one class

BCMSN v2.1—8-21

FIFO provides basic store-and-forward capability. FIFO is the default queuing algorithm in most Catalyst switches, requiring no configuration.

FIFO queuing has no classification because all packets belong to the same class. The scheduler services packets in the order they arrive.

Software FIFO queuing is basically an extension of the hardware FIFO queue and is used when QoS is not enabled. All of the queue RAM is allocated to queue 1 (no-expedite queue). After QoS is enabled, the four software queues are all defaulted to best-effort delivery service with essentially four FIFO queues (for Layer 3 QoS). Each queue is serviced by weighted round robin (WRR). Each queue has the same weight on Gigabit Ethernet interfaces or the same initial weight on Fast Ethernet interfaces, and there is no expedite queue.

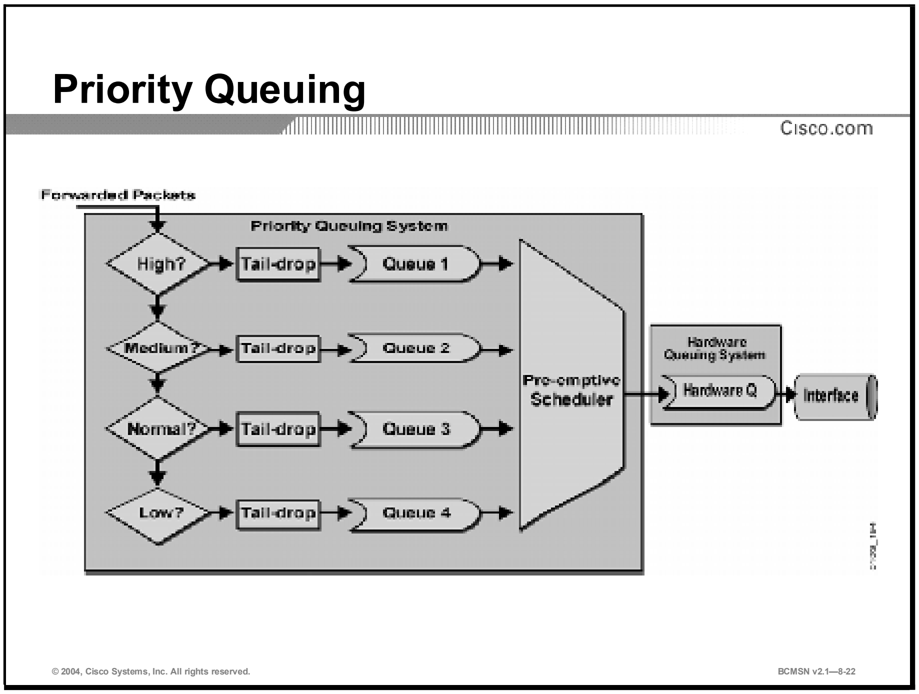

Priority Queuing

Cisco.com

BCMSN v2.1—8-22

PQ is designed to give strict priority to important traffic. PQ ensures that priority traffic is serviced most often and strictly prioritizes according to network protocol (such as IP, IPX, or AppleTalk), incoming interface, packet size, source or destination address, and so on.

PQ is statically configured and implemented through the use of the expedite queue. Priority queuing has one major drawback in that the low queues may never be sampled as long as higher-priority traffic is being processed, resulting in queue starvation.

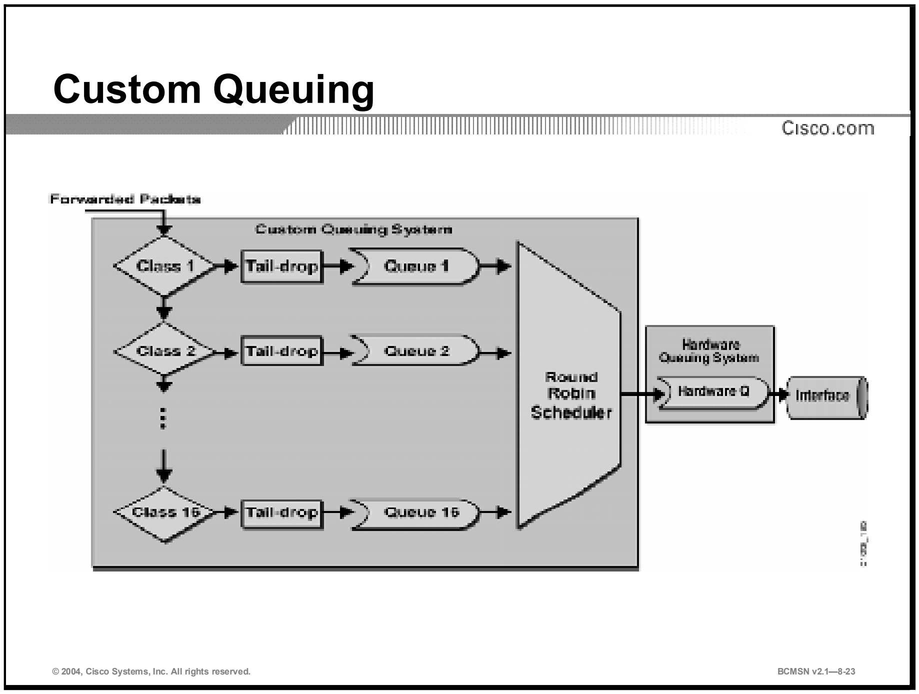

Custom Queuing

Cisco.com

BCMSN v2.1—8-23

CQ reserves a percentage of the available bandwidth of an interface for each selected traffic type. If a particular type of traffic is not using the bandwidth reserved for it, other traffic types may use the remaining reserved bandwidth.

CQ is statically configured and does not provide for automatic adaptation for changing network conditions.

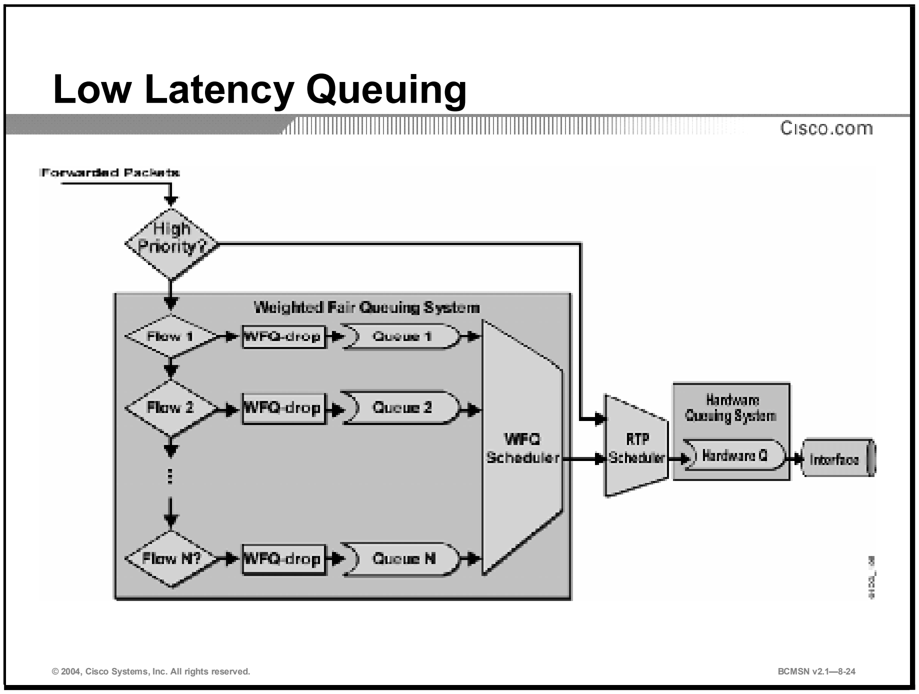

Low Latency Queuing

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—8-24

Low latency queuing (LLQ) provides strict PQ. This feature allows you to configure the priority status for a class within CBWFQ, and is not limited to UDP port numbers.

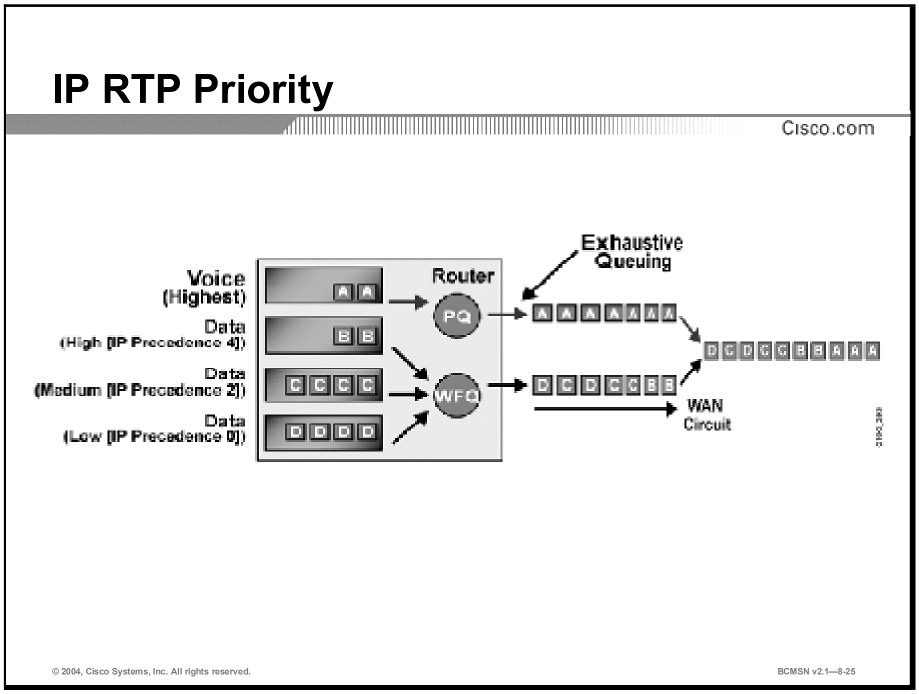

IP RTP Priority

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—8-25

The IP RTP Priority feature provides a strict PQ scheme that allows delay-sensitive data such as voice to be dequeued and sent before packets in other queues. Use this feature on serial interfaces in conjunction with either weighted fair queuing (WFQ) or CBWFQ on the same outgoing interface. In either case, traffic matching the range of UDP ports specified for the priority queue is guaranteed strict priority over other CBWFQ classes or WFQ flows; packets in the priority queue are always serviced first.

Voice traffic can be identified by its RTP port numbers and classified into a priority queue. The result of using this feature is that voice traffic is serviced as strict priority in preference to nonvoice traffic.

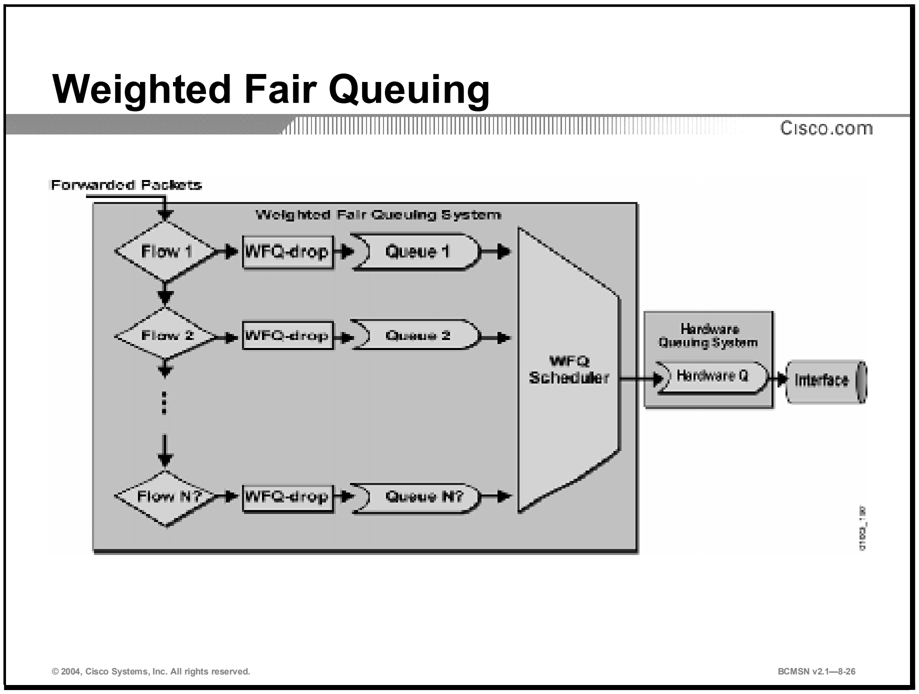

Weighted Fair Queuing

Cisco.com

BCMSN v2.1—8-26

WFQ classifies traffic into different flows based on such characteristics as source and destination address, protocol, and port and socket of the session. WFQ is the default queuing mechanism for E1 and slower links.

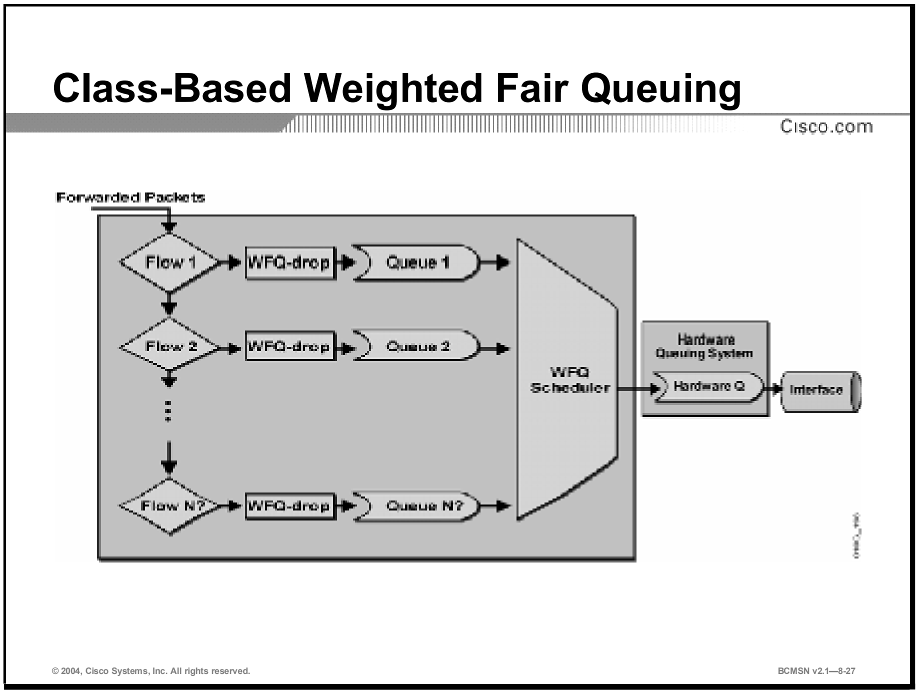

Class-Based Weighted Fair Queuing

Cisco.com

BCMSN v2.1—8-27

CBWFQ extends the standard WFQ functionality to provide support for user-defined traffic classes. They allow you to specify the exact amount of bandwidth to be allocated for a specific class of traffic. Taking into account available bandwidth on the interface, you can configure up to 64 classes and control distribution among them.

Weighted Round Robin

Cisco.com

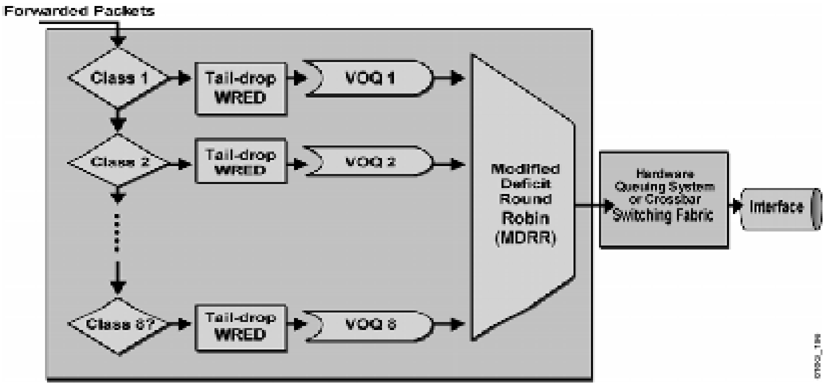

BCMSN v2.1—8-28

WRR scheduling is used on Layer 3 switches on egress ports to manage the queuing and sending of packets. WRR places a packet in one of four queues based on IP precedence, from which it derives a delay priority. The Layer 3 switches automatically use WRR on egress ports. Unlike other queuing properties, you do not configure WRR through the device interface properties. Instead, you configure WRR through policies defined at the device level. With WRR, each queue is given a weight. This weight is used when congestion occurs on the port to give weighted priority to high-priority traffic without starving low-priority traffic. The weights provide the queues with an implied bandwidth for the traffic on the queue. The higher the weight, the greater the implied bandwidth. The queues are not assigned specific bandwidth, however, and when the port is not congested, all queues are treated equally.

Consideration of the behavior of congested systems is not simple because traffic rates do not simply rise to a level, stay there for a while, and then subside. Periods of traffic congestion can be quite long, with heavily concentrated losses. Linear increases in buffer size do not result in large decreases in packet drop rates. A slight increase in the number of active connections can result in a large increase in the packet loss rate. This behavior of congested networks suggests that, because the level of busy period traffic is not predictable, it would be difficult to efficiently size networks to reduce congestion adequately. Observers of network congestion report that it is traffic “spikes” that cause the most losses.

WRR queuing uses the scheduler on the switches going from software queues to hardware queues. If the expedite queue is enabled, WRR services traffic until the queue is empty before servicing the other three queues. There are four software queues per port on most Cisco multilayer switches.

Congestion Avoidance

Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottleneck points. Congestion avoidance is achieved through packet dropping using more complex techniques than simple tail drop. This topic helps you select the congestion avoidance components of a Cisco QoS solution, given specific quality and application requirements.

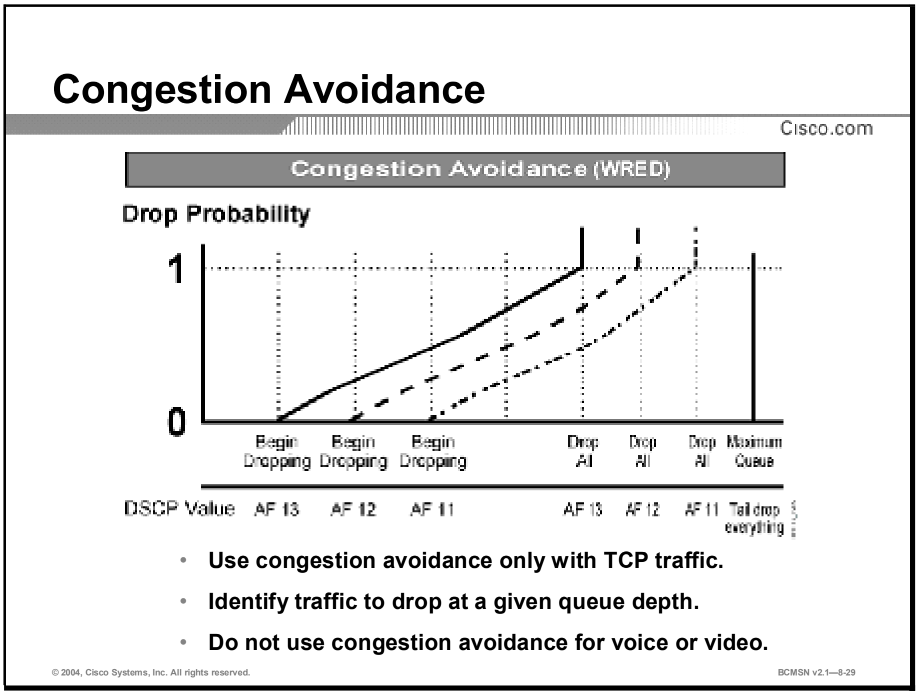

Congestion Avoidance

Cisco.com

- **Use congestion avoidance only with TCP traffic.**
- **Identify traffic to drop at a given queue depth.**
- **Do not use congestion avoidance for voice or video.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—8-29

When an interface on a router cannot transmit a packet immediately, the packet is queued, either in an interface Tx ringer or the interface output hold queue, depending on the switching path used. Packets are then taken out of the queue and eventually transmitted on the interface.

If the arrival rate of packets to the output interface exceeds the router capability to buffer and forward traffic, the queues increase to their maximum length and the interface becomes congested. Tail drop is the router default queuing response to congestion. When the output queue is full and tail drop is in effect, all packets trying to enter (at the tail of) the queue are dropped until the congestion is eliminated and the queue is no longer full.

Tail drop treats all traffic equally and does not differentiate between classes of service. The simple tail-dropping scheme unfortunately does not work very well in environments with a large number of TCP flows or in environments in which selective dropping is desired. Understanding of the interaction between TCP stack intelligence and dropping in the network is required to implement a more efficient and fair dropping scheme, especially in service provider environments.

Tail drop has these shortcomings:

- When congestion occurs, dropping affects most of the TCP sessions, which simultaneously back off and then restart again. This behavior causes inefficient link utilization at the congestion point (TCP global synchronization).
- TCP starvation, where all buffers are temporarily seized by aggressive flows, and normal TCP flows experience buffer starvation, can occur.
- Buffering at the point of congestion can introduce delay and jitter, because packets are stuck waiting in queues.
- There is no differentiated drop mechanism, and therefore, premium traffic is dropped in the same way as best-effort traffic.
- Even in the event of a single TCP stream across an interface, the presence of other non-TCP traffic can congest the interface and TCP traffic will also be dropped. In this scenario, the feedback to the TCP protocol is very poor, and therefore it cannot adapt properly to a congested network.

A router can handle multiple concurrent TCP sessions. There is a high probability that when traffic exceeds the queue limit, it vastly exceeds the limit due to the bursty nature of packet networks. However, there is also a high probability that excessive traffic depth caused by packet bursts is temporary, and that traffic does not stay congested except at points where traffic flows merge or at edge routers.

If the receiving router drops all traffic that exceeds the queue limit, as is done by default (with tail drop), many TCP sessions then simultaneously go into slow start. Consequently, traffic temporarily slows down to the extreme and then all flows slow-start again. This activity creates a condition called global synchronization.

Global synchronization occurs as waves of congestion crest only to be followed by troughs during which the transmission link is not fully utilized. Global synchronization of TCP hosts, for example, can occur because packets are dropped all at once. Global synchronization manifests itself when multiple TCP hosts reduce their transmission rates in response to packet dropping, and then increase their transmission rates once again when the congestion is reduced. The most important point is that the waves of transmission known as global synchronization result in significant link underutilization.

WFQ, if configured on an interface, has a more elaborate scheme for dropping traffic, because it is able to punish the most aggressive flows via its Congestion Discard Threshold (CDT)-based dropping algorithm. Unfortunately, WFQ does not scale to backbone speeds.

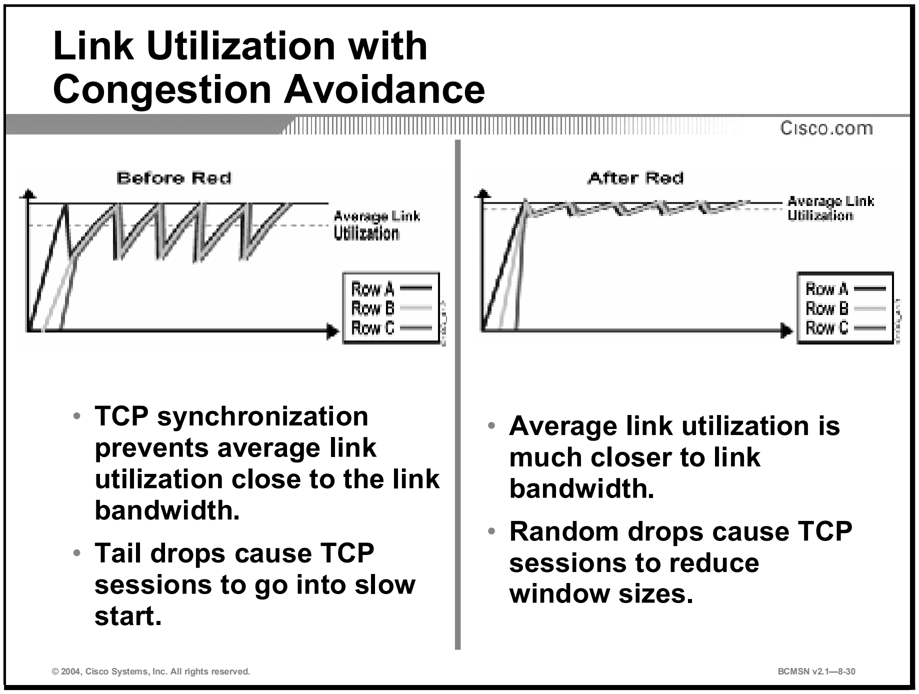

Link Utilization with Congestion Avoidance

Cisco.com

- TCP synchronization prevents average link utilization close to the link bandwidth.
- Tail drops cause TCP sessions to go into slow start.

- Average link utilization is much closer to link bandwidth.
- Random drops cause TCP sessions to reduce window sizes.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—8-30

WRED is a congestion avoidance mechanism that randomly drops packets with a certain IP precedence when the buffers are reaching a defined threshold. WRED is a combination of two features: tail drop and random early detection (RED).

RED is not precedence- or CoS-aware. RED uses one of the single thresholds when that threshold value for the buffer fills. RED will start to randomly drop packets (not all packets, as in tail drop) until the maximum threshold is reached, where all packets are then dropped. The probability of dropping a packet rises linearly with the increase of buffer filling above the threshold.

RED and WRED are very useful congestion avoidance mechanisms when the traffic type is TCP-based. For other types of traffic, RED is not very efficient. This is because RED takes advantage of the windowing mechanism used by TCP to manage congestion. RED avoids the typical congestion that occurs on a router when multiple TCP sessions are going through the same router port.

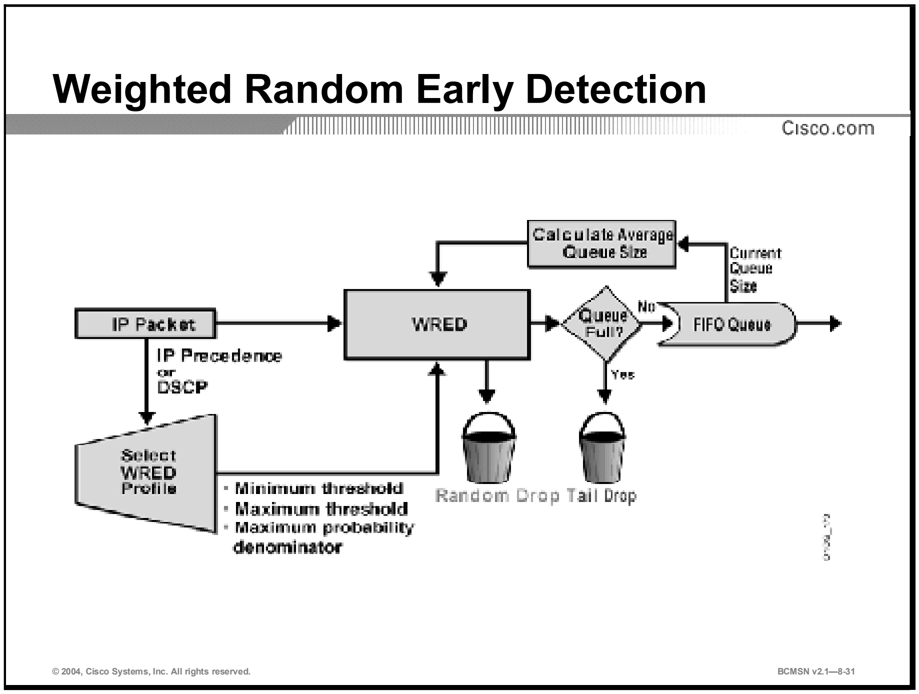

Weighted Random Early Detection

Cisco.com

BCMSN v2.1—8-31

WRED is similar to RED in the fact that both define some threshold, and that threshold starts to randomly drop packets. WRED is also CoS-aware, which means that a CoS value is added to each of the thresholds. When the threshold is exceeded, WRED randomly drops packets with the CoS assigned. Consider this example, where there are two thresholds in the queue:

- CoS 0 and 1 are assigned to threshold 1, which is set to 50 percent of buffer filling
- CoS 2 and 3 are assigned to threshold 2, which is set to 80 percent of buffer filling

As soon as the buffer exceeds 50 percent utilization, packets with CoS 0 and 1 will start to be randomly dropped. More packets will be dropped when the buffer utilization grows. If 80 percent is reached, packets with CoS 2 and 3 will start to drop, but WRED will also continue to randomly drop packets with CoS 0 and 1. More packets will always be dropped with CoS 0 and CoS 1 than with CoS 2 and CoS 3.

Traffic Conditioning

The IOS QoS software solutions include two traffic-shaping tools, generic traffic shaping (GTS) and Frame Relay traffic shaping (FRTS), to manage traffic and congestion on the network. The IOS policing tool is CAR. This topic helps you select the traffic conditioning components of a Cisco QoS solution, given specific quality and application requirements.

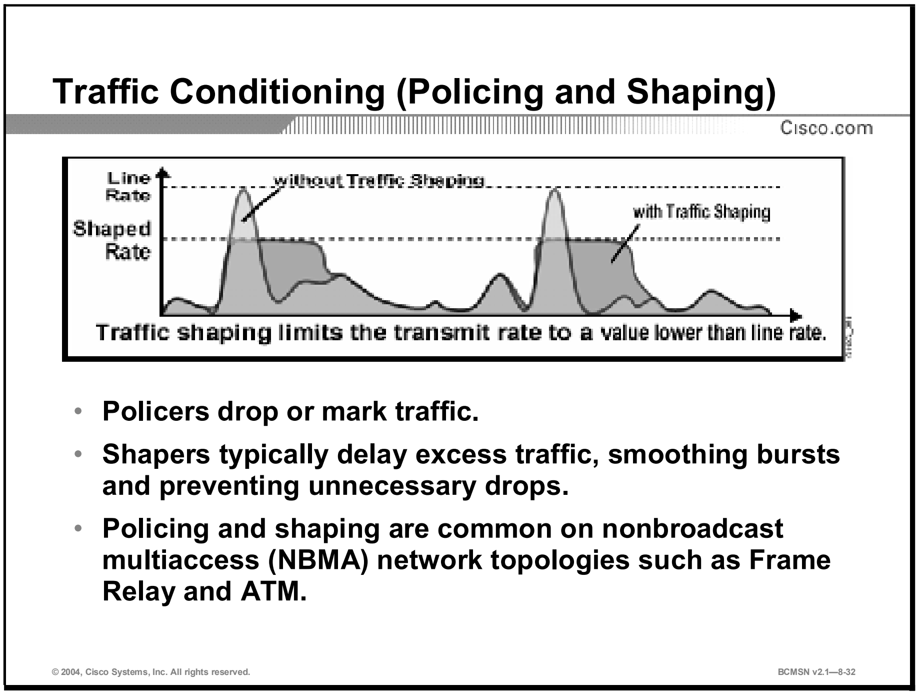

Traffic Conditioning (Policing and Shaping)

Cisco.com

Traffic shaping limits the transmit rate to a value lower than line rate.

- Policers drop or mark traffic.
- Shapers typically delay excess traffic, smoothing bursts and preventing unnecessary drops.
- Policing and shaping are common on nonbroadcast multiaccess (NBMA) network topologies such as Frame Relay and ATM.

BCMSN v2.1—8-32

Both shaping and policing mechanisms are used in a network to control the rate at which traffic is admitted into the network. Both mechanisms use classification so they can differentiate traffic. They also use metering to measure the rate of traffic and compare it to the configured shaping or policing policy.

The difference between shaping and policing is described as follows:

- Shaping meters the traffic rate and delays excessive traffic so that it stays within the desired rate limit. With shaping, traffic bursts are smoothed out, producing a steadier flow of data. Reducing traffic bursts helps reduce congestion in the core of the network.
- Policing drops excess traffic in order to control traffic flow within specified limits. Policing does not introduce any delay to traffic that conforms to traffic policies. It can, however, cause more TCP retransmissions, because traffic in excess of specified limits is dropped.

Rate limiting is typically used to satisfy one of these requirements:

- Prevent and manage congestion in ATM and Frame Relay networks, where asymmetric bandwidths are used along the traffic path. This requirement prevents the Layer 2 network from dropping large amounts of traffic, by dropping excess traffic at ingress to the ATM or Frame Relay networks based on Layer 3 information (for example: IP precedence, DSCP, access list, protocol type, and so on).
- Limit the access rate on an interface when the high-speed physical infrastructure is used in transport, but subrate access is desired.
- Engineer bandwidth so that traffic rates to certain applications or classes of traffic follow a specified traffic-rate policy.
- Implement a virtual time-division multiplexing (TDM) system where an IP network is used but has the bandwidth characteristics of a TDM system, that is, fixed maximum available bandwidth. Inbound and outbound policing can, for example, be used on one router to split a single point-to-point link into two or more virtual point-to-point links by assigning a portion of the bandwidth to each class, thus preventing any class from monopolizing the link in either direction.

A shaper typically delays excess traffic using a buffer, or mechanism, to hold packets and shape the flow when the data rate of the source is higher than expected. Traffic shaping smoothes traffic by storing traffic above the configured rate in a queue. Therefore, shaping increases buffer utilization on a router but causes nondeterministic packet delays. Shaping can also interact with a Frame Relay network, adapting to indications of Layer 2 congestion in the WAN.

A policer typically does the following:

- Drops nonconforming traffic
- Supports marking of traffic
- Is more efficient in terms of memory utilization (no additional buffering of packets is needed)
- Does not increase buffer usage

Both policing and shaping ensure that traffic does not exceed a bandwidth limit, but they have different impacts on the traffic:

- Policing drops packets more often, generally causing more retransmissions of connection-oriented protocols.
- Shaping adds variable delay to traffic, possibly causing jitter.

To perform rate limiting, routers must meter (or measure) traffic rates through their interfaces. To enforce a rate limit, metered traffic is said to do one of the following:

- Conform to the rate limit, if the rate of traffic is below or equal to the configured rate limit
- Exceed the rate limit, if the rate of traffic is above the configured rate limit

The metering is usually performed with an abstract model called a token bucket, which is used when processing each packet. The token bucket can calculate whether the current packet conforms or exceeds the configured rate limit on an interface.

IOS software supports five token-bucket-based rate-limiting methods. Three of these are shaping mechanisms:

- GTS
- FRTS
- Class-based shaping

Two token-bucket-based rate-limiting methods are policing mechanisms:

- CAR
- Class-based policing

Link Efficiency Mechanisms

QoS software offers three link efficiency mechanisms that work in conjunction with queuing and traffic shaping to improve efficiency and predictability of the application service levels: link fragmentation and interleaving (LFI), payload compression, and header compression. This topic helps you select the link efficiency components of a Cisco QoS solution, given specific quality and application requirements.

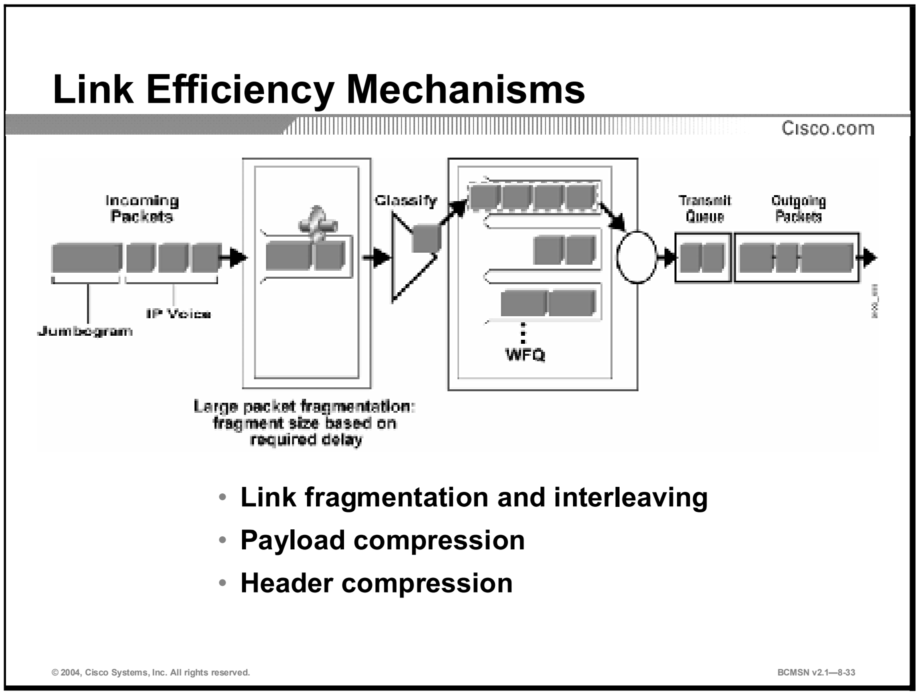

Link Efficiency Mechanisms

Cisco.com

- Link fragmentation and interleaving
- Payload compression
- Header compression

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—8-33

While many mechanisms exist for optimizing throughput and reducing delay in network traffic within the QoS portfolio, QoS does not create bandwidth. QoS optimizes the use of existing resources, and enables the differentiation of traffic according to the operator policy.

Payload compression does create additional bandwidth, because it squeezes packet payloads, and therefore increases the amount of data that can be sent through a transmission resource in a given time period. Payload compression is mostly performed on Layer 2 frames and, therefore, compresses the entire Layer 3 packet.

Note that IP Payload Compression Protocol (PCP) is a fairly new technique for compressing payloads on Layer 3, and can handle out-of-order data.

As compression squeezes payloads, it both increases the perceived throughput and decreases perceived latency in transmission because smaller packets with compressed payloads take less time to transmit than the larger, uncompressed packets.

Compression is a CPU-intensive task and can add per-packet delay due to the application of the compression method to each frame. The transmission (serialization) delay, however, is reduced, because the resulting frame is smaller. Depending on the complexity of the payload compression algorithm, overall latency might be reduced, especially on low-speed links.

IOS software supports three different compression algorithms used in Layer 2 compression:

- STAC (or Stacker)
- Microsoft Point-to-Point Compression (MPPC)
- Predictor

These algorithms differ vastly in their compression efficiency, and in utilization of router resources. Verify the support for compression on your Catalyst switch.

All compression methods are based on eliminating redundancy when sending the same or similar data over a transmission medium. One piece of data, which is often repeated, is the protocol header. In a flow, the header information of packets in the same flow does not change much over the lifetime of that flow. Therefore, most of the header information could be sent at the beginning of a session only, stored in a dictionary, and then referenced in later packets by a short dictionary index.

The IETF has standardized two methods for use with IP protocols:

- **TCP header compression (also known as Van Jacobson, or VJ, header compression):** Used to compress the packet TCP headers over slow links, thus considerably improving the interactive application performance
- **RTP header compression:** Used to compress UDP and RTP headers, thus lowering the delay for transporting real-time data, such as voice and video over slower links

It is important to note that header compression is performed on a link-by-link basis. Header compression cannot be performed across multiple routers, because routers need full Layer 3 header information to be able to route packets to the next hop.

LFI is a Layer 2 technique, where all Layer 2 frames are broken into small, equal-size fragments, and transmitted over the link in an interleaved fashion.

When fragmentation and interleaving are in effect, all frames waiting in the queuing system are fragmented, smaller frames are prioritized, and a mixture of fragments is sent over the link. Small frames may be scheduled behind larger frames in the WFQ system. LFI fragments all frames, thus reducing the queuing delay of small frames, because they are sent almost immediately. Link fragmentation therefore reduces delay and jitter by expediting transfer of smaller frames through the hardware queue.

These LFI mechanisms are implemented in IOS software:

- Multilink PPP (MLP) with interleaving is by far the most common and widely used form of LFI.
- FRF.11 Annex C LFI is used with Voice over Frame Relay (VoFR).
- FRF.12 Frame Relay LFI is used with Frame Relay data connections.

In an ATM network, using separate permanent virtual circuits (PVCs) carrying voice and data can be used to interleave packets when they are output on an interface.

QoS Within the Campus Network

The QoS implementation for a campus network differs at the Campus Backbone, Building Access, and Building Distribution submodules. This topic helps you select QoS solutions for the campus network, given specific network and application needs.

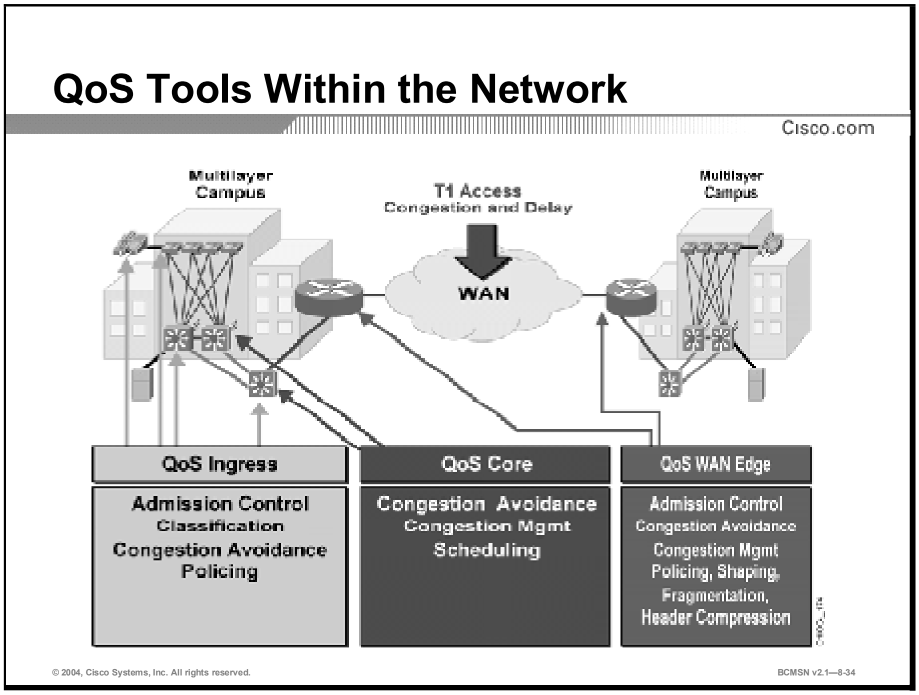

QoS Tools Within the Network

Cisco.com

QoS Ingress	QoS Core	QoS WAN Edge
Admission Control Classification Congestion Avoidance Policing	Congestion Avoidance Congestion Mgmt Scheduling	Admission Control Congestion Avoidance Congestion Mgmt Policing, Shaping, Fragmentation, Header Compression

BCMSN v2.1—8-34

Not all QoS tools are appropriate in all areas of the network. The diagram points out what types of QoS should be implemented in each module and submodule of the Enterprise Composite Network model. As pictured in the diagram, you can pick different QoS solutions based on where in the network QoS is being implemented.

In general, Layer 3 edge devices perform these QoS functions:

- Packet classification
- Admission control
- Configuration management

In general, backbone Layer 3 devices perform these QoS functions:

- Congestion management
- Congestion avoidance

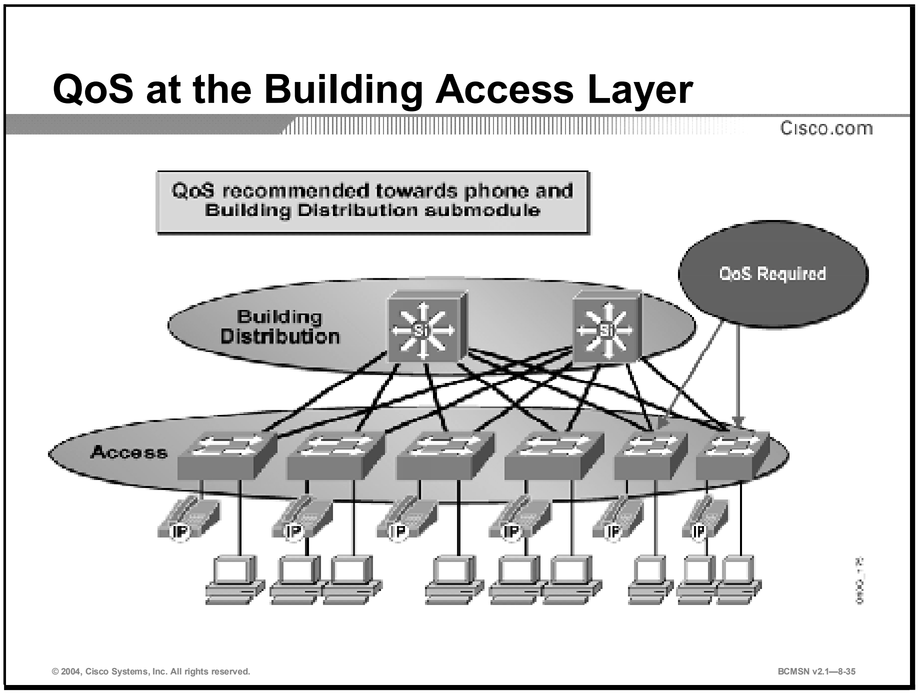

QoS at the Building Access Layer

Cisco.com

QoS recommended towards phone and Building Distribution submodule

BCMSN v2.1—8-35

The Building Access submodule is typically where the trust boundary is formed. That is where the precedence is set for the packets and then trusted throughout the rest of the network. A switch at the access layer must support the following:

- Multiple VLANs on the access port to which an end user is attached
- Manipulation of the QoS or CoS values provided by the end device
- Extension of the trust boundary for the CoS or DSCP marking toward the end devices

There are times when the devices attached to the campus network do not classify their traffic with the appropriate Layer 2 and Layer 3 markings. When considering your choices for access-layer devices, consider the ability of the switch to classify and mark traffic at the edge of the network using ACLs and service policies. Doing so allows you to offer QoS as a service throughout the network, and have it administered at the edge of the network where CPU resources are plentiful, rather than at the Campus Backbone and Building Distribution submodules where QoS classification and marking could adversely affect network responsiveness and utilization.

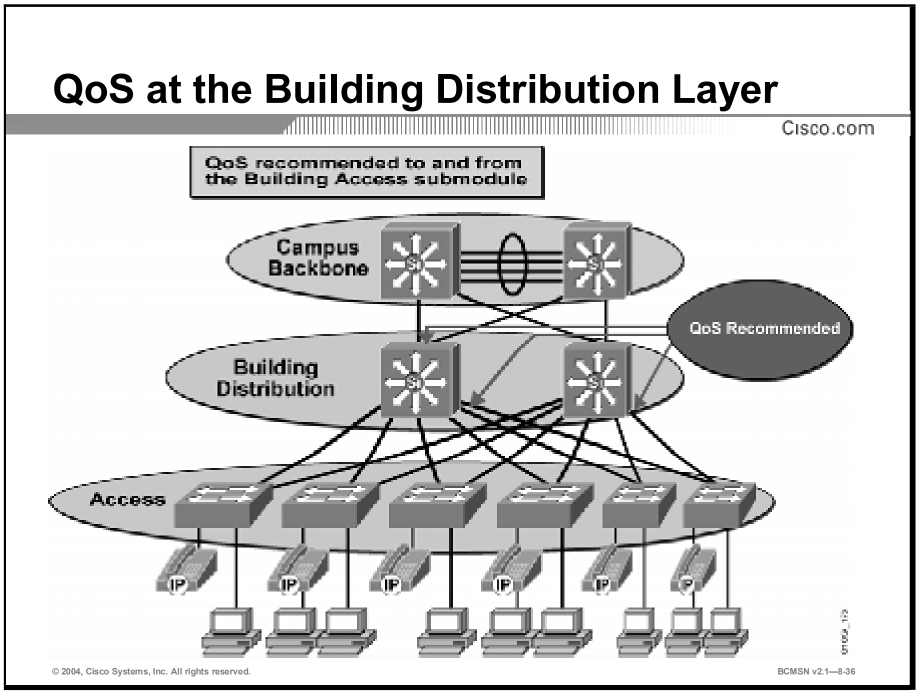

QoS at the Building Distribution Layer

Cisco.com

QoS recommended to and from the Building Access submodule

BCMSN v2.1—8-36

QoS at the Building Distribution submodule requires these changes to the implementation of the switches:

- Enable QoS.
- Change the default CoS-to-DSCP table so that CoS and DSCP behavior can be maintained throughout the network.
- Configure service policies to classify traffic that does not contain a CoS-to-DSCP marking that you can trust.
- Enable CoS or DSCP trust on the ports where trust is appropriate.

DSCP is for Layer 3-aware access, and CoS is for Layer 2 access only.

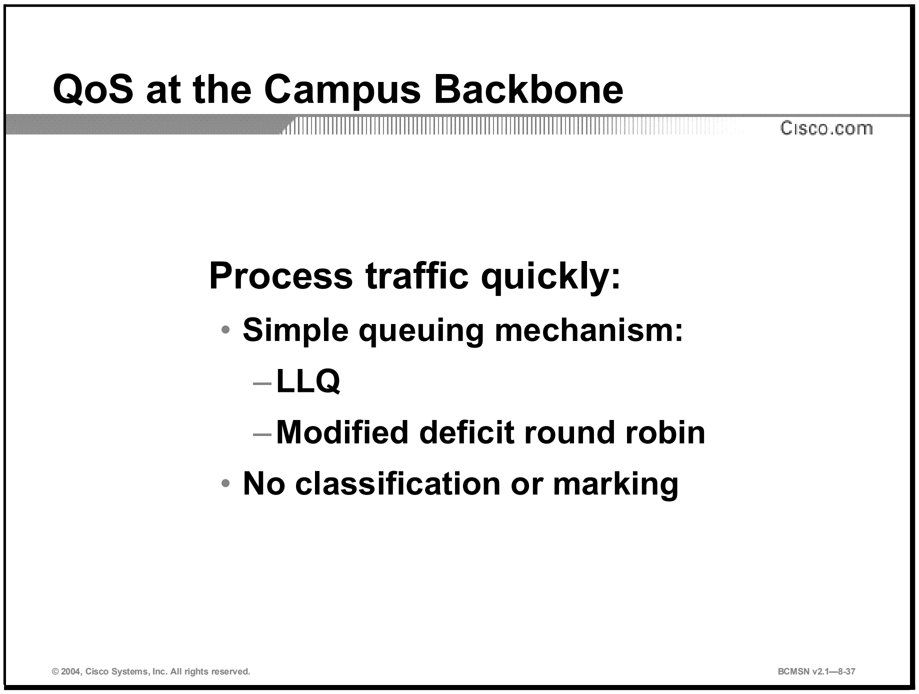

QoS at the Campus Backbone

Cisco.com

Process traffic quickly:

- **Simple queuing mechanism:**
  - **LLQ**
  - **Modified deficit round robin**
- **No classification or marking**

BCMSN v2.1—8-37

To maintain the QoS policy throughout the network, the core device can provide some QoS congestion management features. The traffic going to the Campus Backbone submodule should already be classified and marked at the Building Access or Building Distribution submodules, so the Campus Backbone submodule should be able to process the traffic quickly using a simple queuing mechanism. There should be no need to run QoS classification and marking tools within the Campus Backbone submodule of the network.

If QoS is implemented in the Campus Backbone submodule, keep it to a minimum to facilitate high-speed queuing with some form of intelligent dropping. The typical queue service mechanisms are LLQ or modified deficit round robin (MDRR) scheduling.

Summary of Key Cisco IOS Software QoS Categories and Features

Each QoS feature has an important role to play in the network. This topic summarizes the key IOS software QoS features so you can determine when to use each.

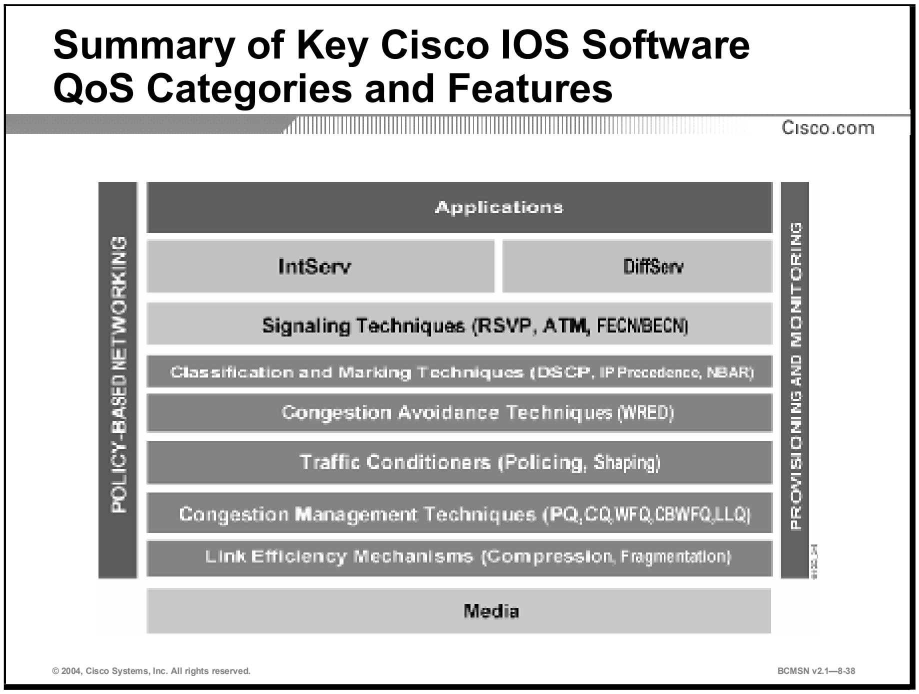

Summary of Key Cisco IOS Software QoS Categories and Features

Cisco.com

POLICY-BASED NETWORKING

Applications

IntServ

DiffServ

Signaling Techniques (RSVP, ATM, FECN/BECN)

Classification and Marking Techniques (DSCP, IP Precedence, NBAR)

Congestion Avoidance Techniques (WRED)

Traffic Conditioners (Policing, Shaping)

Congestion Management Techniques (PQ,CQ,WFQ,CBWFQ,LLQ)

Link Efficiency Mechanisms (Compression, Fragmentation)

Media

PROVISIONING AND MONITORING

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—8-38

The figure summarizes the key IOS software QoS categories and features. The base layer contains the transport protocols.

The middle section lists the tools that are used in deploying QoS: classification and marking (including NBAR), congestion avoidance, traffic conditioners, congestion management, and link efficiency.

The top row contains the different applications that benefit from QoS.

Below the top line are IntServ and DiffServ. In the bars below are RSVP and DSCP, the two marking tools of IntServ and DiffServ, respectively. The IntServ architecture defines fine-grained (flow-based) methods of performing IP traffic admission control that uses RSVP. The DiffServ architecture defines methods of classifying IP traffic into coarse-grained service classes and defines forwarding treatment based on these classifications.

Summary

This topic summarizes the key points discussed in this lesson.

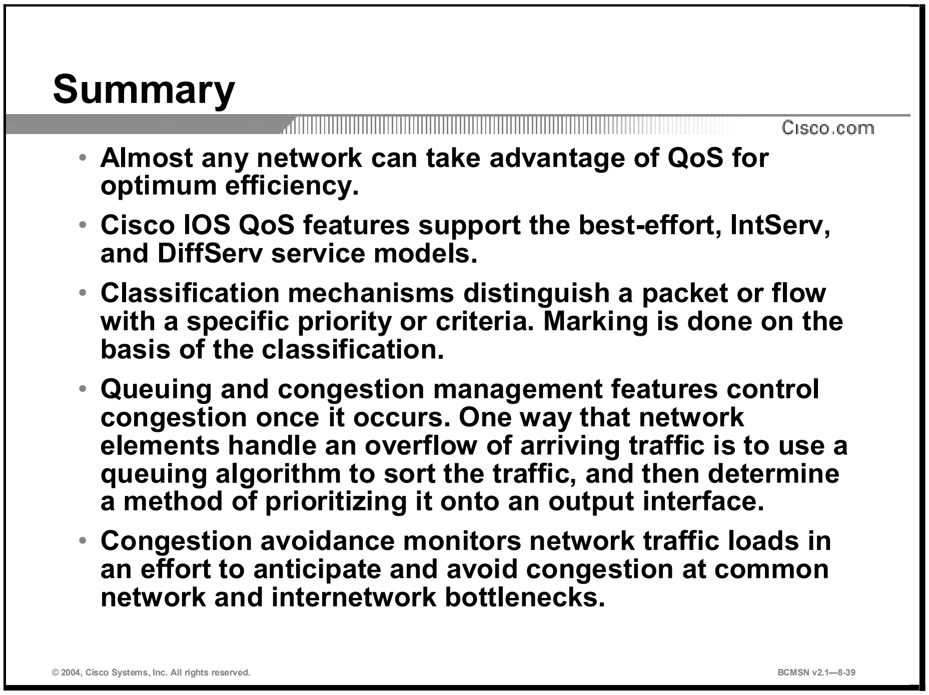

Summary

Cisco.com

- Almost any network can take advantage of QoS for optimum efficiency.
- Cisco IOS QoS features support the best-effort, IntServ, and DiffServ service models.
- Classification mechanisms distinguish a packet or flow with a specific priority or criteria. Marking is done on the basis of the classification.
- Queuing and congestion management features control congestion once it occurs. One way that network elements handle an overflow of arriving traffic is to use a queuing algorithm to sort the traffic, and then determine a method of prioritizing it onto an output interface.
- Congestion avoidance monitors network traffic loads in an effort to anticipate and avoid congestion at common network and internetwork bottlenecks.

BCMSN v2.1—8-39

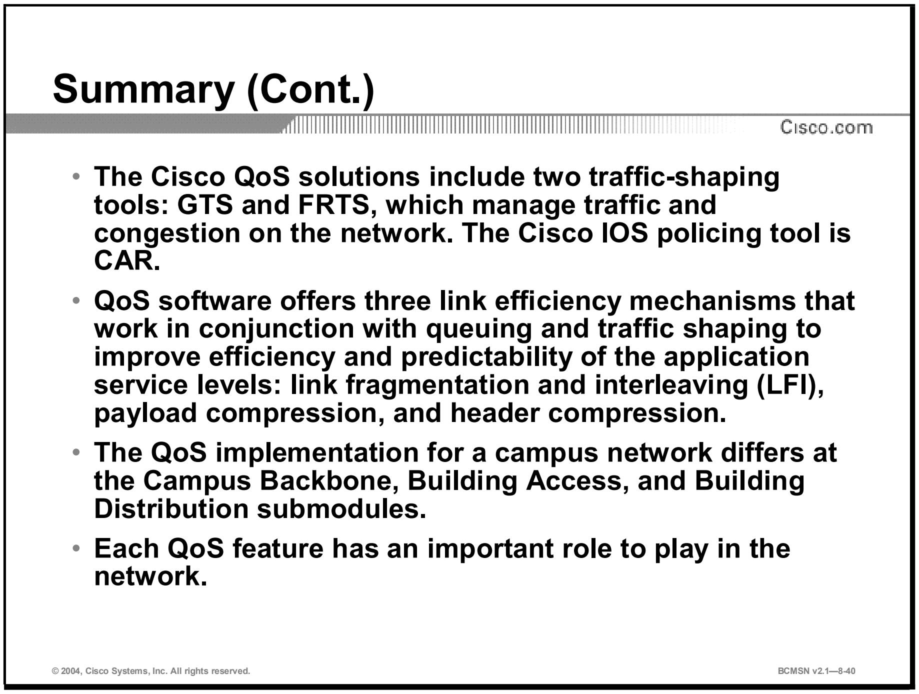

Summary (Cont.)

Cisco.com

- The Cisco QoS solutions include two traffic-shaping tools: GTS and FRTS, which manage traffic and congestion on the network. The Cisco IOS policing tool is CAR.
- QoS software offers three link efficiency mechanisms that work in conjunction with queuing and traffic shaping to improve efficiency and predictability of the application service levels: link fragmentation and interleaving (LFI), payload compression, and header compression.
- The QoS implementation for a campus network differs at the Campus Backbone, Building Access, and Building Distribution submodules.
- Each QoS feature has an important role to play in the network.

BCMSN v2.1—8-40

References

For additional information, refer to these resources:

- "Cisco IOS Quality of Service" at http://www.cisco.com/warp/public/732/Tech/qos/
- "Implementing QoS Solutions for H.323 Video Conferencing Over IP" at http://www.cisco.com/warp/public/105/video-qos.html
- "VoIP Interoperability with Cisco Express Forwarding and Policy Based Routing" at http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t11/ft_cef26.htm

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three quality issues does QoS address? (Choose three.)

A) jitter

B) delay

C) reliability

D) bandwidth

E) packet loss

Q2) Which two QoS models provide effective quality of service on an enterprise network? (Choose two.)

A) RSVP

B) IntServ

C) queuing

D) DiffServ

E) best-effort

Q3) What is the purpose of trust boundaries?

A) to determine where in the network to trust the packet

B) to determine where in the network packets are dropped

C) to determine where in the network to reclassify packets

D) It is an administrative configuration where Layer 2 or Layer 3 priority designations of frames or packets are either accepted or not.

Q4) Which three features are congestion management tools? (Choose three.)

A) LFI

B) LLQ

C) WFQ

D) CLLLQ

E) CBWFQ

F) CCBWFQ

Q5) Which type of traffic does WRED discard?

A) TCP

B) UDP

C) LDP

D) CDP

Q6) Which function does traffic shaping perform?

A) discards traffic based on outbound traffic restrictions

B) buffers traffic based on inbound bandwidth restrictions

C) buffers traffic based on destination bandwidth restrictions

D) discards traffic based on destination bandwidth restrictions

Q7) Which QoS mechanism reduces delay and jitter on slower-speed links by breaking up large datagrams and interleaving low-delay traffic packets with the resulting smaller packets?

A) LFI

B) CRTP

C) dCRTP

D) queuing

Q8) Which QoS feature should access lists support?

A) inline power

B) support for voice signaling

C) support for multiple VLANs per port

D) support for BECN and FECN acknowledgments

Q9) Which QoS model requires end-to-end configuration?

A) LLQ

B) IntServ

C) DiffServ

D) traffic conditioning

Quiz Answer Key

Q1) A, B, E

Relates to: Quality Issues on a Multilayer Switched Network

Q2) B, D

Relates to: QoS Service Models

Q3) D

Relates to: Classification and Marking

Q4) B, C, E

Relates to: Queuing and Congestion Management

Q5) A

Relates to: Congestion Avoidance

Q6) C

Relates to: Traffic Conditioning

Q7) A

Relates to: Link Efficiency Mechanisms

Q8) C

Relates to: QoS Within the Campus Network

Q9) B

Relates to: Summary of Key Cisco IOS Software QoS Categories and Features

Configuring QoS in Multilayer Switched Networks

Overview

QoS selects network traffic (both unicast and multicast), prioritizes it according to its relative importance, and uses congestion avoidance to provide priority-indexed treatment. QoS can also limit the bandwidth used by specific network traffic. QoS can make network performance more predictable and bandwidth utilization more effective.

The QoS implementation for Catalyst switches is based on the DiffServ architecture, a standard from the IETF. With DiffServ, each packet is classified upon entry into the network. The classification is carried in the IP packet header, using six bits from the IP ToS field to carry the classification information.

Relevance

QoS is crucial for many networks that transport critical data, particularly voice and video.

Objectives

Upon completing this lesson, you will be able to:

- Configure the QoS policy using the modular QoS user interface
- Configure classification and marking
- Configure queuing and congestion management
- Configure congestion avoidance
- Troubleshoot the QoS configuration

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Configuring the QoS Policy with Modular QoS
- Configuring Classification and Marking
- Configuring Queuing and Congestion Management
- Configuring Congestion Avoidance
- Troubleshooting the QoS Configuration
- Summary
- Quiz

Configuring the QoS Policy with Modular QoS

The Modular QoS CLI (MQC) is a provisioning mechanism in Cisco IOS software that allows for separation of packet classification (class maps) and policies (policy maps) applied on the defined classes, from the application of those policies on interfaces and subinterfaces (service policy). This topic explains how to configure the QoS policy using the modular QoS user interface.

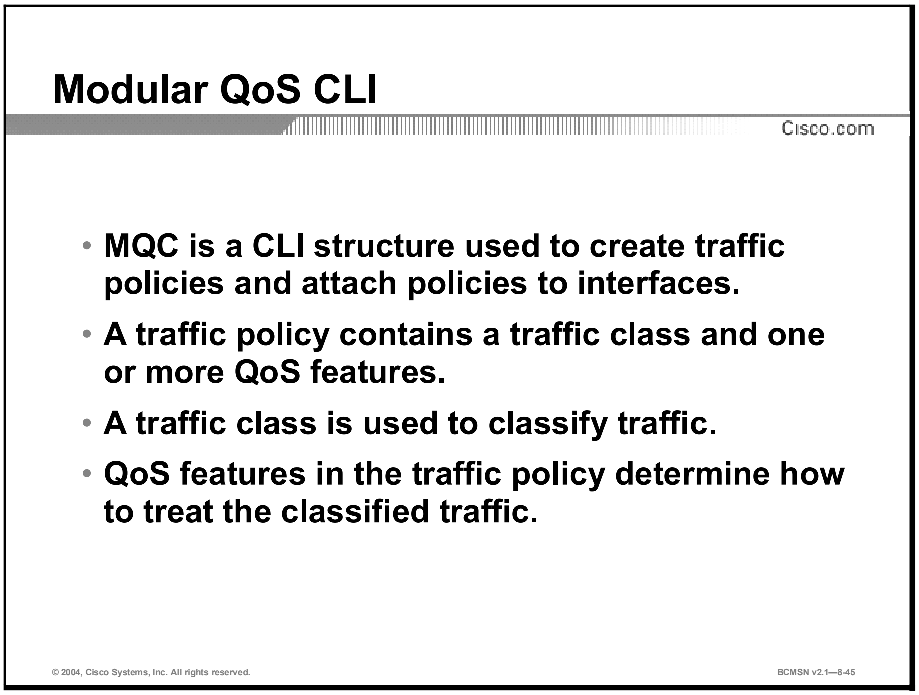

Modular QoS CLI

Cisco.com

- **MQC is a CLI structure used to create traffic policies and attach policies to interfaces.**
- **A traffic policy contains a traffic class and one or more QoS features.**
- **A traffic class is used to classify traffic.**
- **QoS features in the traffic policy determine how to treat the classified traffic.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—8-45

The MQC forms the basis for provisioning DiffServ, and all the QoS mechanisms are part of the class maps (classification) or policy maps (policing, shaping, queuing, congestion avoidance, packet marking, Layer 2 CoS marking, and so on).

Packets entering a DiffServ domain can be metered, marked, shaped, or policed to implement traffic policies (as defined by the administrative authority). In IOS software, classifying and marking is done using the MQC class maps. Metering is done using a token bucket algorithm, shaping is done using GTS or FRTS, and policing is done using class-based policing, or CAR. Cisco provides for the per-class accounting Management Information Base (MIB), where statistics for each class (regardless of congestion) are available for management purposes.

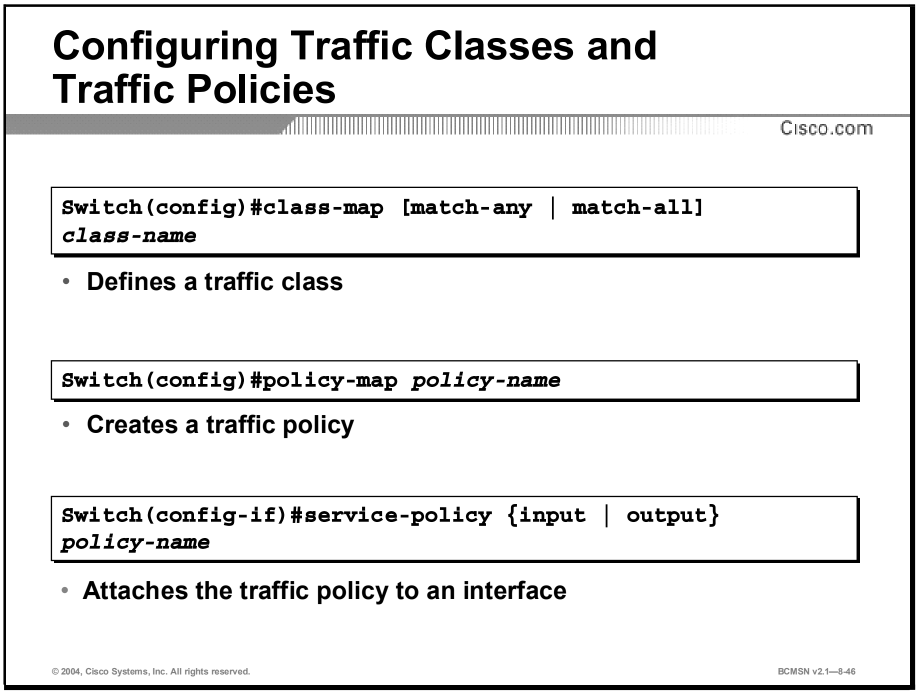

Configuring Traffic Classes and Traffic Policies

```
Switch(config)#class-map [match-any | match-all]
class-name
```

- Defines a traffic class

```
Switch(config)#policy-map policy-name
```

- Creates a traffic policy

```
Switch(config-if)#service-policy {input | output}
policy-name
```

- Attaches the traffic policy to an interface

BCMSN v2.1—8-46

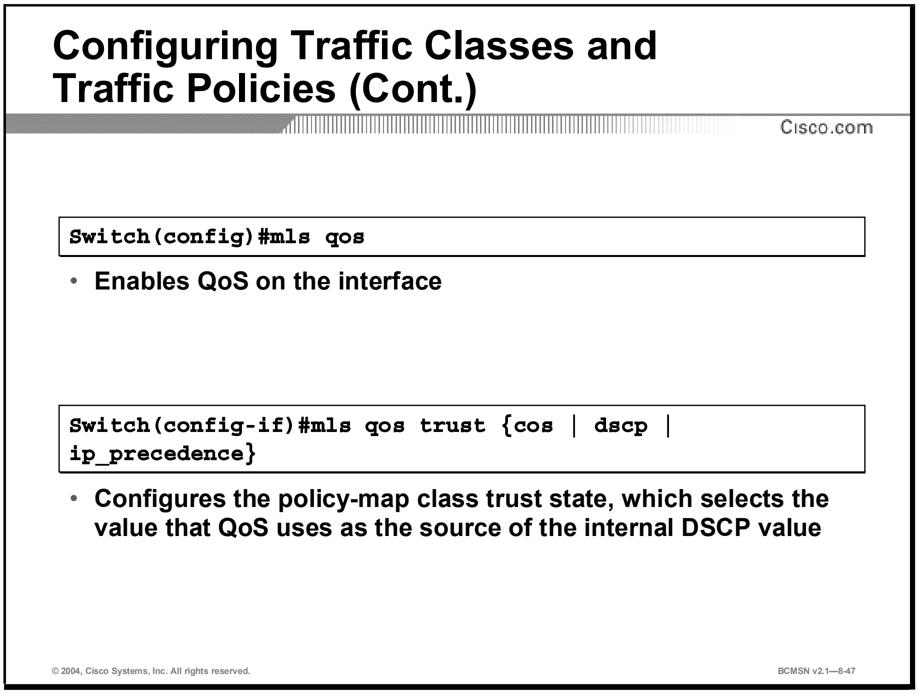

Configuring Traffic Classes and Traffic Policies (Cont.)

```
Switch(config)#mls qos
```

- Enables QoS on the interface

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

- Configures the policy-map class trust state, which selects the value that QoS uses as the source of the internal DSCP value

BCMSN v2.1—8-47

Configuring Traffic Classes and Traffic Policies

To configure and enable class-based QoS features, you create a traffic class and a traffic policy and then attach the traffic policy to an interface. To configure and enable class-based QoS features, perform these tasks:

Step	Command	Notes and Comments
1.	Define a traffic class. `Switch(config)#class-map [match-any	match-all] class-name`	Specify the **match-all** and **match-any** options only if more than one match criterion is configured in the traffic class. The **class-map match-all** command is used when all of the match criteria in the traffic class must be met in order for a packet to match the specified traffic class. The **class-map match-any** command is used when only one of the match criteria in the traffic class must be met in order for a packet to match the specified traffic class. If neither the **match-all** nor **match-any** keyword is specified, the traffic class will behave in a manner consistent with the **class-map match-all** command.
2.	Create a traffic policy by associating the traffic class with one or more QoS features. `Switch(config)#policy-map policy-name`	Specify the traffic policy name using the **policy-map** command.
3.	Select an interface to configure. `Switch(config)#interface {vlan vlan_ID	{fastethernet	gigabitethernet} slot/interface	port-channel number}`	Select the interface to configure on the Catalyst switch.
4.	Attach the traffic policy to the interface. `Switch(config-if)#service-policy {input	output} policy-name`	Attach a traffic policy to an interface and specify the direction in which the policy should be applied (either on packets coming into the interface or packets leaving the interface).
5.	Enable QoS in the switch. `Switch(config)#mls qos`	Globally enable QoS.
6.	Configure the trust state of an interface. `Switch(config-if)#mls qos trust [dscp	cos]`	When configuring the trust state of an interface, note the following: ■ By default, the trust state of an interface, when QoS is enabled, is set to "untrusted." When QoS is disabled on an interface, the trust state resets to "trust dscp." If the trust state is modified using the **qos trust** command on an interface, the trust state remains the same whether or not QoS is enabled. Use the default **mls qos trust** command to revert to the default behavior. ■ You can use the **no mls qos trust** command to set the interface state to "untrusted." ■ For traffic received on an ingress interface configured to trust CoS using the **mls qos trust cos** command, the transmit CoS is always the incoming packet CoS (or the ingress interface default CoS if the packet is

Step	Command	Notes and Comments
		received untagged). ■ When the interface trust state is not configured to "trust dscp" using the **mls qos trust dscp** command, the security and QoS ACL classification will always use the interface DSCP and not the incoming packet DSCP.
7.	Exit configuration mode. `Switch(config-if)#`**`end`**	Return to configuration mode.

Example: Defining Traffic Classes

In the example, two traffic classes are created and their match criteria are defined. For the first traffic class, called "class1," ACL 101 is used as the match criterion. For the second traffic class, called "class2," ACL 102 is used as the match criterion. Packets are checked against the contents of these ACLs to determine if they belong to the class.

```
Switch(config)#class-map class1
Switch(config-cmap)#match access-group 101
Switch(config-cmap)#exit

Switch(config)#class-map class2
Switch(config-cmap)#match access-group 102
Switch(config-cmap)#exit
```

Example: Creating a Traffic Policy

In the example, a traffic policy called "policy1" is defined to contain policy specifications for the two classes: class1 and class2. The match criteria for these classes were defined in the traffic classes.

For class1, the policy includes a bandwidth allocation request and a maximum packet count limit for the queue reserved for the class. For class2, the policy specifies only a bandwidth allocation request.

```
Switch(config)#policy-map policy1
Switch(config-pmap)#class class1
Switch(config-pmap-c)#bandwidth 3000
Switch(config-pmap-c)#queue-limit 30
Switch(config-pmap)#exit

Switch(config-pmap)#class class2
Switch(config-pmap-c)#bandwidth 2000
Switch(config-pmap)#exit
```

Example: Attaching a Traffic Policy to an Interface

The example shows how to attach an existing traffic policy to an interface. After you define a traffic policy with the **policy-map** command, you can attach it to one or more interfaces to specify the traffic policy for those interfaces by using the **service-policy** command in interface configuration mode. Although you can assign the same traffic policy to multiple interfaces, each interface can have only one traffic policy attached at the input and only one traffic policy attached at the output.

```
Switch(config)#interface fastethernet 1/0/0
Switch(config-if)#service-policy output policy1
Switch(config-if)#exit

Switch(config)#interface gigabitethernet 1/1
Switch(config-if)#service-policy output policy1
Switch(config-if)#exit
```

Example: Configuring a Gigabit Ethernet Interface

The example shows how to configure Gigabit Ethernet interface 1/1 with the **mls qos trust cos** keywords:

```
Switch#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)#interface gigabitethernet 1/1
Switch(config-if)#mls qos qos trust cos
Switch(config-if)#end
```

Default QoS Configuration

The table shows the QoS default configuration.

Feature	Default Value
Global QoS state	Disabled
Interface CoS value	0
Interface DSCP value	0
CoS-to-DSCP map (DSCP set from CoS values)	CoS 0 = DSCP 0 CoS 1 = DSCP 8 CoS 2 = DSCP 16 CoS 3 = DSCP 24 CoS 4 = DSCP 32 CoS 5 = DSCP 40 CoS 6 = DSCP 48 CoS 7 = DSCP 56
DSCP-to-CoS map (CoS set from DSCP values)	DSCP 0-7 = CoS 0 DSCP 8-15 = CoS 1 DSCP 16-23 = CoS 2 DSCP 24-31 = CoS 3 DSCP 32-39 = CoS 4 DSCP 40-47 = CoS 5 DSCP 48-55 = CoS 6 DSCP 56-63 = CoS 7
Marked-down DSCP from DSCP map (policed DSCP)	The marked-down DSCP value equals the original DSCP value (pass-through with no markdown).
Policers	None
Policy maps	None
Tx queue sharing	1/4 of the link bandwidth
Tx queue size	1/4 of the Tx queue entries for the port. The Tx queue size of a port depends on the type of port, ranging from 240 packets per Tx queue to 1920 packets per Tx queue.
Tx queue shaping	None
DCSP-to-Tx-queue map	DSCP 0-15 Queue 1 DSCP 16-31 Queue 2 DSCP 32-47 Queue 3 DSCP 48-63 Queue 4
High priority Tx queue	Disabled
With QoS disabled	
Interface trust state	Trust DSCP
With QoS enabled	With QoS enabled and all other QoS parameters at default values, QoS sets the IP DSCP to 0 and the Layer 2 CoS to 0 in all traffic transmitted.
Interface trust state	Untrusted

Configuration Guidelines

Before beginning the QoS configuration, you should be aware of this information:

- If you have EtherChannel ports configured on your switch, you must configure QoS classification and policing on the EtherChannel. The Tx queue configuration must be configured on the individual physical ports that form the EtherChannel.
- It is not possible to match IP fragments against configured IP extended ACLs to enforce QoS. IP fragments are transmitted as best-effort. IP fragments are denoted by fields in the IP header.
- It is not possible to match IP options against configured IP extended ACLs to enforce QoS. These packets are sent to the CPU and processed by software.
- IP options are denoted by fields in the IP header.
- Control traffic, such as spanning tree bridge protocol data units (BPDUs) and routing update packets received by the switch, are subject to all ingress QoS processing.
- If you want to use the **set** command in the policy map, you must enable IP routing (disabled by default) and configure an IP default route to send traffic to the next-hop device that is capable of forwarding.

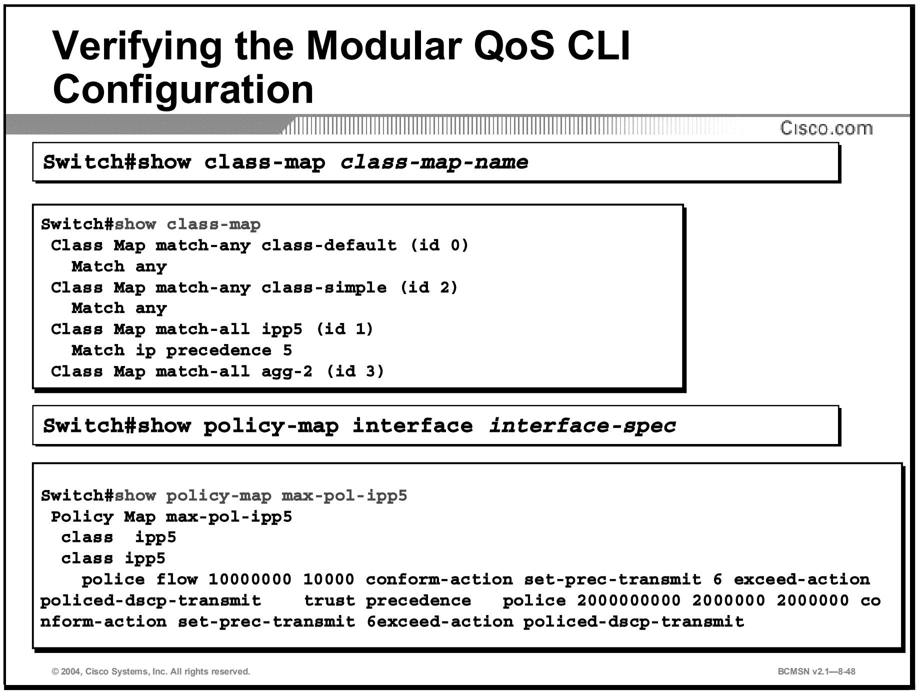

Verifying the Modular QoS CLI Configuration

```
Switch#show class-map class-map-name
```

```
Switch#show class-map
 Class Map match-any class-default (id 0)
   Match any
 Class Map match-any class-simple (id 2)
   Match any
 Class Map match-all ipp5 (id 1)
   Match ip precedence 5
 Class Map match-all agg-2 (id 3)
```

```
Switch#show policy-map interface interface-spec
```

```
Switch#show policy-map max-pol-ipp5
 Policy Map max-pol-ipp5
  class  ipp5
  class ipp5
    police flow 10000000 10000 conform-action set-prec-transmit 6 exceed-action
policed-dscp-transmit    trust precedence   police 2000000000 2000000 2000000 co
nform-action set-prec-transmit 6exceed-action policed-dscp-transmit
```

BCMSN v2.1—8-48

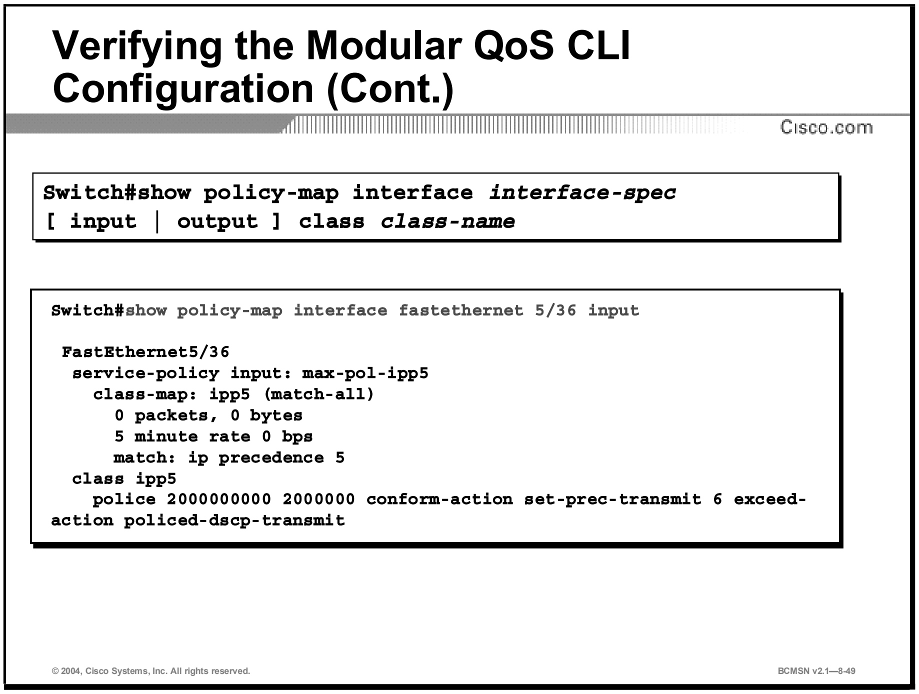

Verifying the Modular QoS CLI Configuration (Cont.)

```
Switch#show policy-map interface interface-spec
[ input | output ] class class-name
```

```
Switch#show policy-map interface fastethernet 5/36 input

 FastEthernet5/36
  service-policy input: max-pol-ipp5
    class-map: ipp5 (match-all)
      0 packets, 0 bytes
      5 minute rate 0 bps
      match: ip precedence 5
  class ipp5
    police 2000000000 2000000 conform-action set-prec-transmit 6 exceed-
action policed-dscp-transmit
```

BCMSN v2.1—8-49

To display the information relating to a traffic class or traffic policy, use one of the following commands in EXEC mode, as needed. To display the configuration of a traffic policy and its associated traffic class, use the **show policy-map** EXEC command.

Command	Purpose
show class-map	Displays all traffic class information.
show class-map *class-map-name*	Displays the traffic class information for the user-specified traffic class.
show policy-map	Displays all configured traffic policies.
show policy-map *policy-name*	Displays the user-specified traffic policy.
show policy-map interface	Displays the configurations and statistics of all input and output policies attached to an interface.
show policy-map interface *interface-spec*	Displays the configuration and statistics of the input and output policies attached to a particular interface.
show policy-map interface *interface-spec* **input**	Displays the configuration and statistics of the input policy attached to an interface.
show policy-map interface *interface-spec* **output**	Displays the configuration and statistics of the output policy attached to an interface.
show policy-map [**interface** [*interface-spec* [*input* | *output*] [**class** *class-name*]]]	Displays the configuration and statistics of the class name configured in the policy.

Example: show policy-map interface Command Output

In the example, 1.25 Mbps of traffic is sent ("offered") to a policer class:

```
Switch#show policy-map interface s3/0
 Serial3/0

  Service-policy output: policy1

   Class-map: police (match all)
    148803 packets, 36605538 bytes
    30 second offered rate 1249000 bps, drop rate 249000 bps
    Match: access-group 101
    police:
     cir 500000 bps, conform-burst 10000, pir 1000000, peak-
burst 100000
     conformed 59538 packets, 14646348 bytes; action: transmit
     exceeded 59538 packets, 14646348 bytes; action: set-prec-
transmit 2
     violated 29731 packets, 7313826 bytes; action: drop
     conformed 499000 bps, exceed 500000 bps violate 249000
bps

   Class-map: class-default (match-any)
    19 packets, 1990 bytes
    30 seconds offered rate 0 bps, drop rate 0 bps
    Match: any
```

Configuring Classification and Marking

Classification is the process of identifying packets that belong to flows for which reservations have been made so that they can receive the appropriate QoS. The classification in a DiffServ network may be done using just a few bits in the IP header, while IntServ classification may examine up to five fields in the packet: the source address, destination address, protocol number, source port, and destination port. This topic explains how to configure classification and marking.

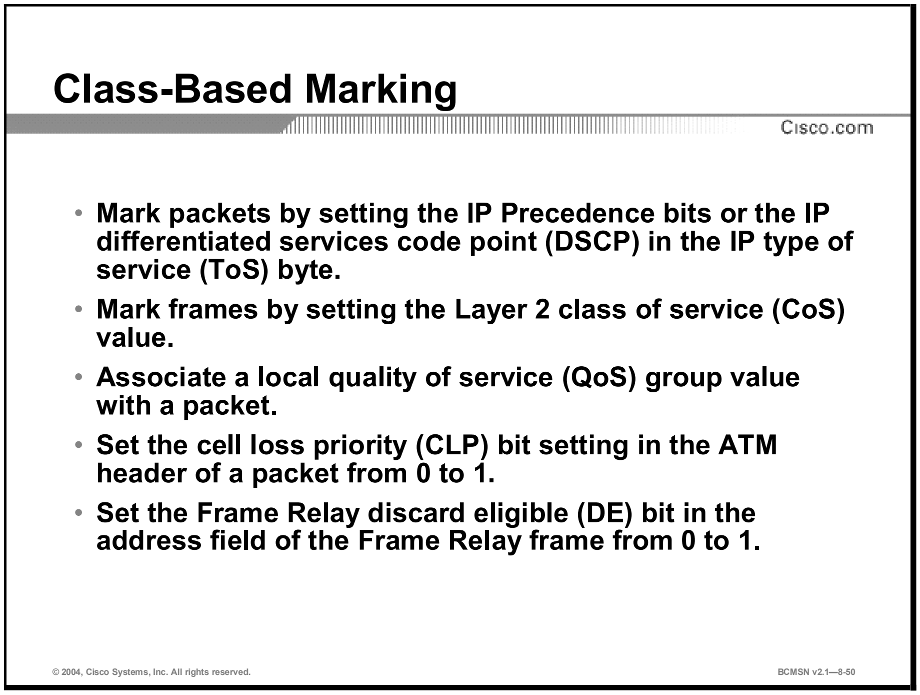

Class-Based Marking

Cisco.com

- **Mark packets by setting the IP Precedence bits or the IP differentiated services code point (DSCP) in the IP type of service (ToS) byte.**
- **Mark frames by setting the Layer 2 class of service (CoS) value.**
- **Associate a local quality of service (QoS) group value with a packet.**
- **Set the cell loss priority (CLP) bit setting in the ATM header of a packet from 0 to 1.**
- **Set the Frame Relay discard eligible (DE) bit in the address field of the Frame Relay frame from 0 to 1.**

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—8-50

The class-based packet-marking feature allows users to perform these tasks:

- Mark packets by setting the IP precedence bits or the IP DSCP in the IP ToS byte.
- Mark packets by setting the Layer 2 CoS value.
- Associate a local QoS group value with a packet.
- Set the CLP bit setting in the ATM header of a packet from 0 to 1.
- Set the Frame Relay DE bit in the address field of the Frame Relay frame from 0 to 1.

IP Precedence

IP precedence is one of the service methods for allocating network resources by classifying how the network should treat an IP packet.

IP precedence provides the capability to partition the traffic into multiple classes of service. The network operator may define up to six classes of service and then utilize extended ACLs to define network policies in terms of congestion handling and bandwidth allocation for each class. The IP precedence feature utilizes the three precedence bits in the ToS field in the IP header to specify CoS assignment for each packet. In addition, the IP precedence feature provides considerable flexibility for precedence and network assignment based on IP or MAC address, physical port, and application.

The IP precedence feature enables the network to act either in passive mode, accepting precedence assigned by the customer, or in active mode, using defined policies to either set or override the precedence assignment. IP precedence can be mapped into adjacent technologies to deliver end-to-end QoS policies in a heterogeneous network environment. Thus, IP precedence enables service classes to be established with no changes to existing applications and with no complicated network signaling requirements.

IP Precedence Implementation

You can manipulate three bits in the IP header designated as the ToS field to inform network devices to give priority to the IP packet as it traverses the network. This designation is also called “coloring” a data stream, because it causes IP packets with the ToS set to stand out from other IP packets.

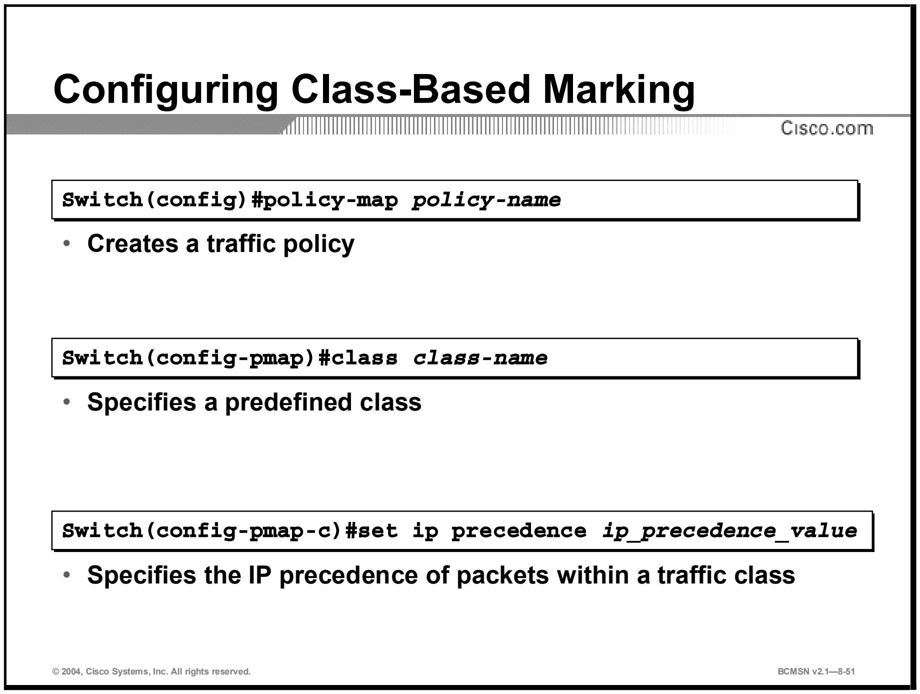

Configuring Class-Based Marking

```
Switch(config)#policy-map policy-name
```

- Creates a traffic policy

```
Switch(config-pmap)#class class-name
```

- Specifies a predefined class

```
Switch(config-pmap-c)#set ip precedence ip_precedence_value
```

- Specifies the IP precedence of packets within a traffic class

BCMSN v2.1—8-51

To mark a packet by setting the IP precedence bits in the ToS byte, use the commands in this table, beginning in global configuration mode.

Step	Command	Notes and Comments
1.	Create a traffic policy by associating the traffic class with one or more QoS features. `Switch(config)#`**`policy-map`** *`policy-name`*	Specify the traffic policy name using the **policy-map** command.
2.	Specify the name of a predefined class. `Switch(config-pmap)#`**`class`** *`class-name`*	The name of a predefined class was defined with the **class-map** command, included in the service policy.
3.	Specify the IP precedence of packets within a traffic class. `Switch(config-pmap-c)#`**`set ip precedence`** *`ip-precedence-value`*	The *ip-precedence-value* is in the range from 0 to 7. You must set IP precedence for delay-sensitive applications, such as voice, to a higher setting, such as 5. The IP precedence setting is critical for voice or if traffic is going over a multivendor network.

IP Precedence Considerations

Use the **set ip precedence** command if the following conditions are present:

- The IP link utilization is high, and the QoS for priority packets needs to have a higher priority than that of other IP packets.
- RSVP is not enabled, and you would like to give priority packets a higher priority over other IP data traffic.

Instead of requesting best-effort, controlled-load, or guaranteed-delay, IP precedence forwards packets by prioritizing classes of service.

IP precedence settings 6 and 7 are reserved for network control information.

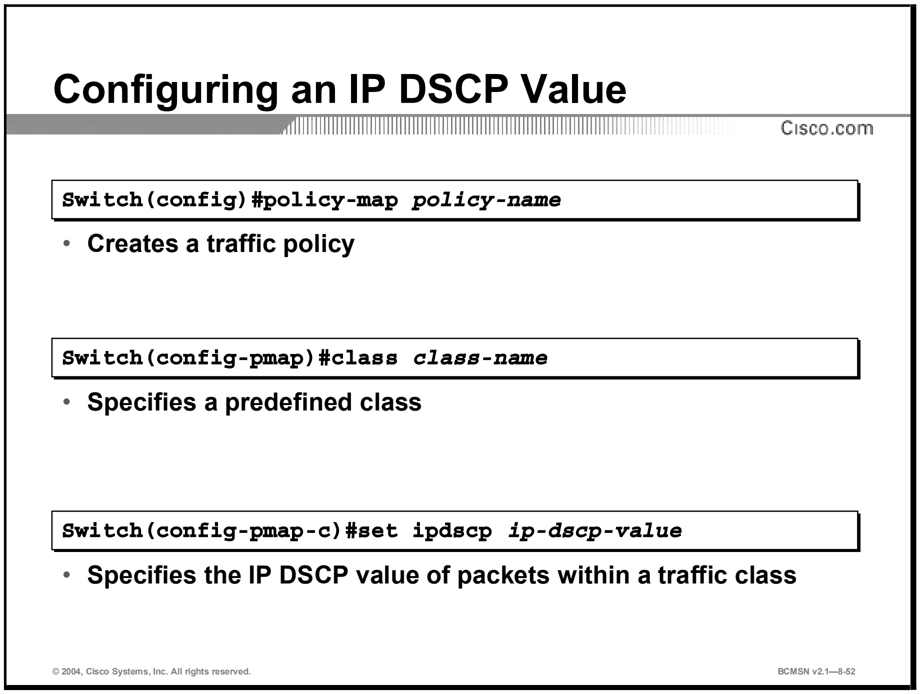

Configuring an IP DSCP Value

Cisco.com

```
Switch(config)#policy-map policy-name
```

- Creates a traffic policy

```
Switch(config-pmap)#class class-name
```

- Specifies a predefined class

```
Switch(config-pmap-c)#set ipdscp ip-dscp-value
```

- Specifies the IP DSCP value of packets within a traffic class

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—8-52

To mark a packet by setting the IP DSCP value, use these commands in this table, beginning in global configuration mode.

Step	Command	Notes and Comments
1.	Create a traffic policy by associating the traffic class with one or more QoS features. `Switch(config)#policy-map policy-name`	Specify the traffic policy name using the **policy-map** command.
2.	Specify the name of a predefined class. `Switch(config-pmap)#class class-name`	The name of a predefined class was defined with the **class-map** command, included in the service policy.
3.	Specify the IP DSCP value of packets within a traffic class. `Switch(config-pmap-c)#set ipdscp ip-dscp-value`	The number is in the range from 0 to 63. Reserved keywords such as “EF” and “AF11” can be specified instead of numeric values.

Example: Configuring an IP Precedence Value

In the example, a service policy called “policy1” is created. This service policy is associated to a previously defined classification policy through the use of the **class** command. This example assumes that a classification policy called “class1” was previously configured.

In the example, the IP precedence bit in the ToS byte is set to 1:

```
Switch(config)#policy-map policy1
Switch(config-pmap)#class class1
Switch(config-pmap-c)#set ip precedence 1
```

Example: Configuring an IP DSCP Value

In the example, a service policy called “policy1” is created. This service policy is associated with a previously defined classification policy through the use of the **class** command. This example assumes that a classification policy called “class1” was previously configured.

In the example, the IP DSCP value in the ToS byte is set to 5:

```
Switch(config)#policy-map policy1
Switch(config-pmap)#class class1
Switch(config-pmap-c)#set ip dscp 5
Switch(config-pmap-c)#class class2
Switch(config-pmap-c)#set ip dscp ef
```

After you configure the settings shown for voice packets at the edge, all intermediate switches are configured to provide low-latency treatment to the voice packets, as follows:

```
Switch(config)#class-map voice
Switch(config-cmap)#match ip dscp ef
Switch(config)#policy qos-policy
Switch(config-pmap)#class voice
Switch(config-pmap-c)#priority 24
```

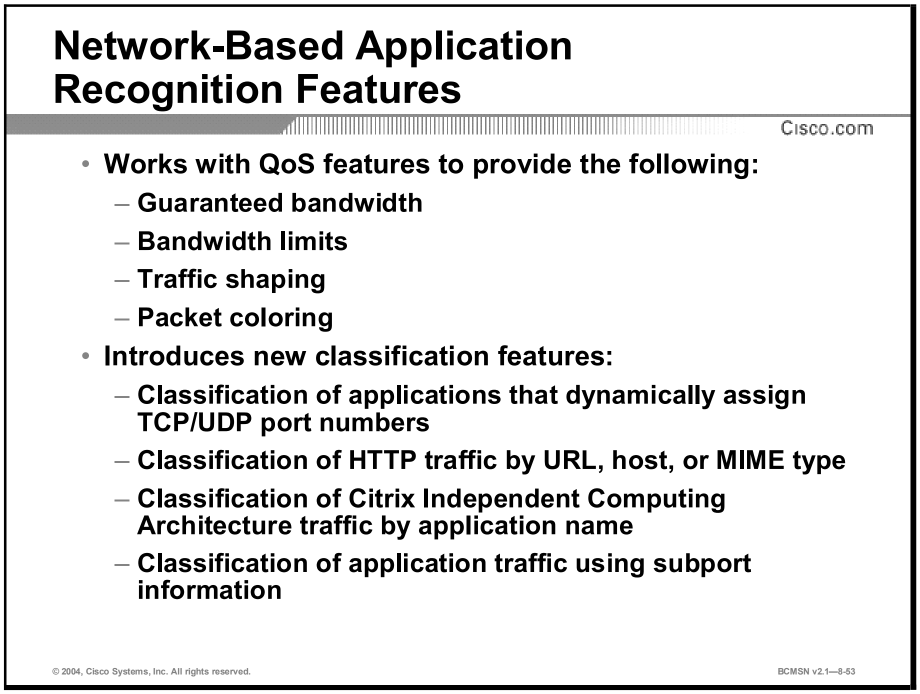

Network-Based Application Recognition Features

- **Works with QoS features to provide the following:**
  - **Guaranteed bandwidth**
  - **Bandwidth limits**
  - **Traffic shaping**
  - **Packet coloring**
- **Introduces new classification features:**
  - **Classification of applications that dynamically assign TCP/UDP port numbers**
  - **Classification of HTTP traffic by URL, host, or MIME type**
  - **Classification of Citrix Independent Computing Architecture traffic by application name**
  - **Classification of application traffic using subport information**

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—8-53

NBAR adds intelligent network classification to network infrastructures. NBAR is a new classification engine that recognizes a wide variety of applications, including web-based and other difficult-to-classify protocols that use dynamic TCP/UDP port assignments.

When an application is recognized and classified by NBAR, a network can invoke services for that specific application. NBAR ensures that network bandwidth is used efficiently by working with QoS features to provide these features:

- Guaranteed bandwidth
- Bandwidth limits
- Traffic shaping
- Packet coloring

NBAR introduces several new classification features:

- Classification of applications that dynamically assign TCP/UDP port numbers
- Classification of HTTP traffic by URL, host, or Multipurpose Internet Mail Extensions (MIME) type
- Classification of Citrix Independent Computing Architecture (ICA) traffic by application name
- Classification of application traffic using subport information

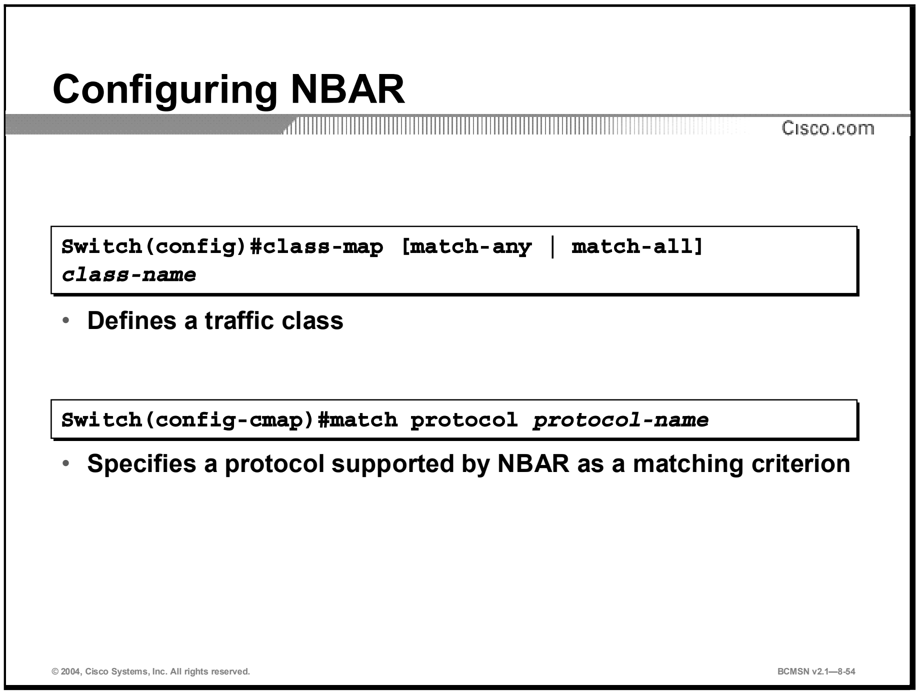

Configuring NBAR

```
Switch(config)#class-map [match-any | match-all]
class-name
```

- Defines a traffic class

```
Switch(config-cmap)#match protocol protocol-name
```

- Specifies a protocol supported by NBAR as a matching criterion

BCMSN v2.1—8-54

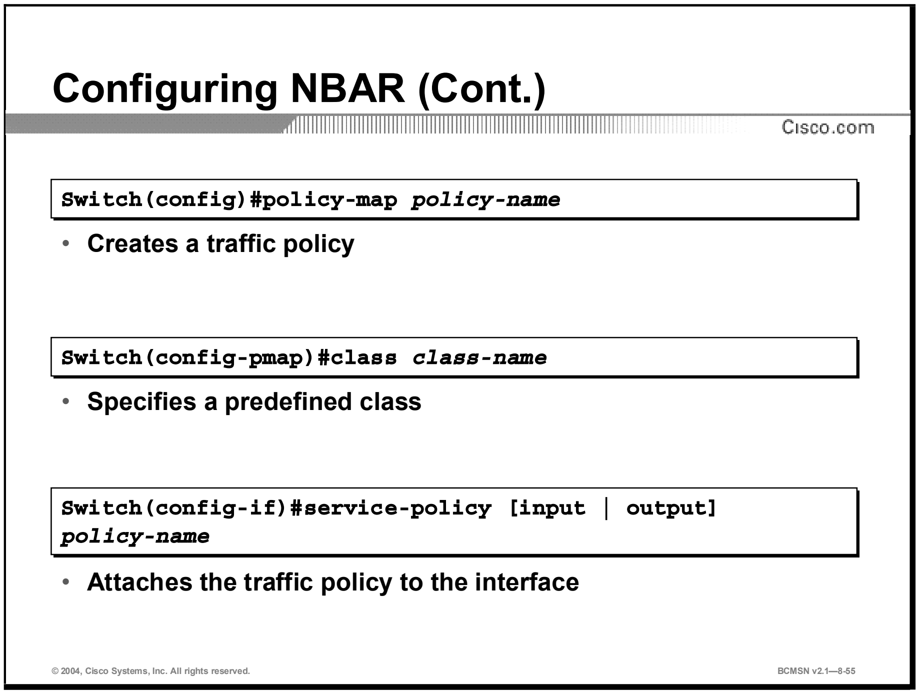

Configuring NBAR (Cont.)

```
Switch(config)#policy-map policy-name
```

- Creates a traffic policy

```
Switch(config-pmap)#class class-name
```

- Specifies a predefined class

```
Switch(config-if)#service-policy [input | output]
policy-name
```

- Attaches the traffic policy to the interface

BCMSN v2.1—8-55

Complete these steps in this table to configure NBAR.

Step	Command	Notes and Comments
1.	Specify the user-defined name of the class map. `Switch(config)#`**`class-map`** `[`**`match-all`** `|` **`match-any`**`]` *`class-name`*	The **match-all** option specifies that all match criteria in the class map must be matched. The **match-any** option specifies that one or more match criteria must match.
2.	Specify a protocol supported by NBAR as a matching criterion. `Switch(config-cmap)#`**`match protocol`** *`protocol-name`*	Use this command to match protocols that are known to NBAR.
3.	Create a traffic policy by associating the traffic class with one or more QoS features. `Switch(config)#`**`policy-map`** *`policy-name`*	Specify the traffic policy name using the **policy-map** command.
4.	Specify the name of a predefined class. `Switch(config-pmap)#`**`class`** *`class-name`*	The name of a predefined class was defined with the **class-map** command, included in the service policy.
5.	Enter QoS policies.	Enter QoS policies in policy map class configuration mode.
6.	Attach the traffic policy to the interface. `Switch(config-if)#`**`service-policy input`** *`policy-name`*	Attach a traffic policy to an interface to apply to packets coming into the interface.
7.	Attach the traffic policy to the interface. `Switch(config-if)#`**`service-policy output`** *`policy-name`*	Attach a traffic policy to an interface to apply to packets leaving the interface.

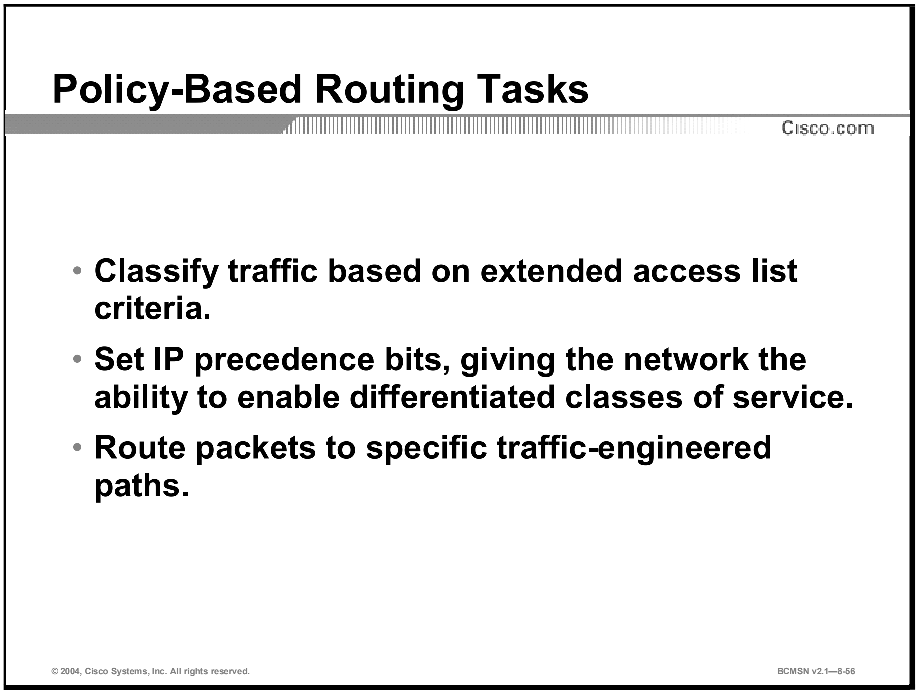

Policy-Based Routing Tasks

Cisco.com

- Classify traffic based on extended access list criteria.
- Set IP precedence bits, giving the network the ability to enable differentiated classes of service.
- Route packets to specific traffic-engineered paths.

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—8-56

PBR gives you a flexible means of routing packets by allowing you to configure a defined policy for traffic flows, lessening reliance on routes derived from routing protocols. It gives you more control over routing by extending and complementing the existing mechanisms provided by routing protocols. PBR allows you to set the IP precedence. It also allows you to specify a path for certain traffic, such as priority traffic over a high-cost link.

PBR allows you to accomplish the following:

- Classify traffic based on extended access list criteria. Access lists, then, establish the match criteria.
- Set IP precedence bits, giving the network the ability to enable differentiated classes of service.
- Route packets to specific traffic-engineered paths; you might need to route them to allow a specific QoS through the network.

Policies can be based on IP address, port numbers, protocols, or size of packets. For a simple policy, you can use any one of these descriptors; for a complicated policy, you can use all of them.

For example, classification of traffic through PBR allows you to identify traffic for different classes of service at the edge of the network and then implement QoS defined for each CoS in the core of the network, using PQ, CQ, or WFQ techniques. This process obviates the need to classify traffic explicitly at each WAN interface in the backbone network.

Configuring Queuing and Congestion Management

Congestion management features allow you to control congestion by determining the order in which packets are sent out on an interface based on priorities assigned to those packets. Congestion management entails the creation of queues, assignment of packets to those queues based on the classification of the packet, and scheduling of the packets in a queue for transmission. This topic explains how to configure congestion management.

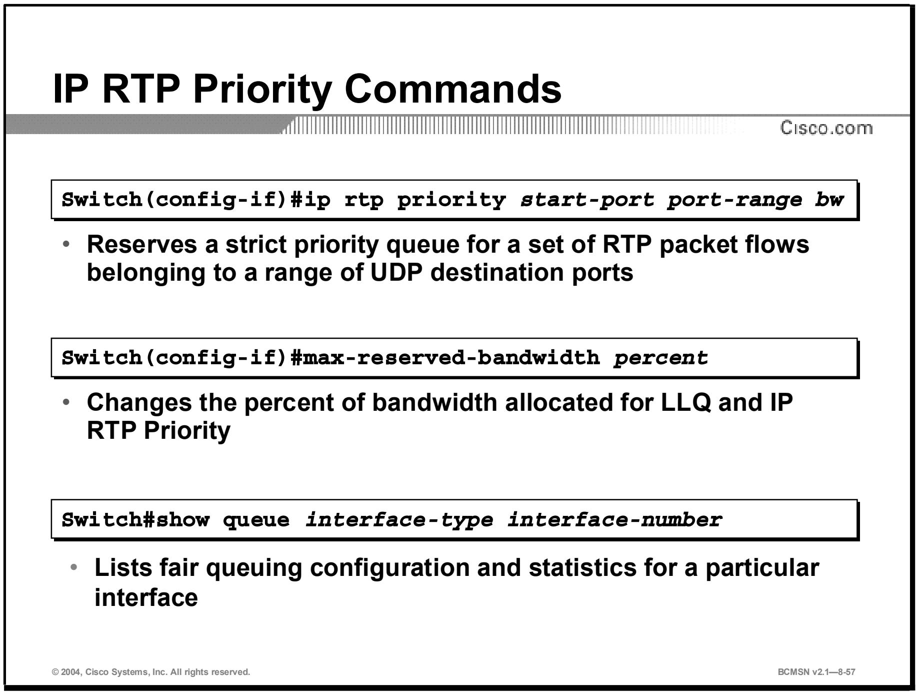

IP RTP Priority Commands

```
Switch(config-if)#ip rtp priority start-port port-range bw
```

- **Reserves a strict priority queue for a set of RTP packet flows belonging to a range of UDP destination ports**

```
Switch(config-if)#max-reserved-bandwidth percent
```

- **Changes the percent of bandwidth allocated for LLQ and IP RTP Priority**

```
Switch#show queue interface-type interface-number
```

- **Lists fair queuing configuration and statistics for a particular interface**

BCMSN v2.1—8-57

IP RTP priority, also known as priority queuing-weighted fair queuing (PQ-WFQ), creates a strict-priority queue for a set of RTP packet flows belonging to a range of UDP destination ports. When the switch recognizes the traffic, it places it into the strict-priority queue. When the priority queue is empty, the other queues are processed according to standard WFQ. IP RTP priority does not become active until there is congestion on the interface.

The IP RTP priority feature provides a strict-priority queuing scheme for delay-sensitive data. You can identify traffic by its RTP port numbers and classify it into a priority queue configured by the **ip rtp priority** command. The result is that priority data is serviced as strict priority in preference to other traffic.

The IP RTP priority feature allows you to specify a range of UDP and RTP ports whose traffic is guaranteed strict-priority service over any other queues or classes using the same output interface. “Strict priority” means that if packets exist in the priority queue, they are dequeued and sent first, that is, before packets in other queues are dequeued.

IP RTP priority is especially useful on slow-speed links whose speed is less than 1.544 Mbps.

You can use this feature in conjunction with either WFQ or CBWFQ on the same outgoing interface. In either case, traffic matching the range of ports specified for the priority queue is guaranteed strict priority over other CBWFQ classes or WFQ flows; packets in the priority queue are always serviced first.

Setting Bandwidth Limitations

When you configure the priority queue with the IP RTP priority feature, you specify a strict bandwidth limitation. This amount of bandwidth is guaranteed to traffic queued in the priority queue. (This is the case whether you use the IP RTP priority feature with CBWFQ or WFQ.)

Note IP RTP priority does not have per-call Call Admission Control (CAC). The admission control is on an aggregate basis. For example, if configured for 96 kbps, IP RTP priority guarantees that 96 kbps is available for reservation. It does not ensure that only four calls of 24 kbps are admitted. A fifth call of 24 kbps could be admitted, but because the five calls will get only 96 kbps, the call quality will deteriorate. (Each call would get 96 / 5 = 19.2 kbps.) In this example, it is the responsibility of the user to ensure that only four calls are placed at one time.

IP RTP priority closely polices use of bandwidth for the priority queue, ensuring that the allocated amount is not exceeded in the event of congestion. In fact, IP RTP priority polices the flow every second. IP RTP priority prohibits transmission of additional packets after the allocated bandwidth is consumed. If it discovers that the configured amount of bandwidth is exceeded, IP RTP priority drops packets. (Enable debugging to watch for this condition.) Close policing allows for fair treatment of other data packets queued in other CBWFQ or WFQ queues. To avoid packet drop, be certain to allocate to the priority queue the optimum amount of bandwidth, taking into consideration the type of coder-decoder (codec) used and interface characteristics. IP RTP priority will not allow traffic beyond the allocated amount.

It is always safest to allocate to the priority queue slightly more than the known required amount of bandwidth.

The sum of all bandwidth allocation for voice and data flows on an interface cannot exceed 75 percent of the total available bandwidth. Bandwidth allocation takes into account the payload plus the IP, RTP, and UDP headers, but again, not the Layer 2 header. Allowing 25 percent bandwidth for other overhead is conservative and safe. On a PPP link, for instance, overhead for Layer 2 headers assumes 4 kbps. You can change the amount of configurable bandwidth for IP RTP priority using the **max-reserved-bandwidth** command on the interface.

If you know how much bandwidth is required for additional overhead on a link, under aggressive circumstances in which you want to give some traffic as much bandwidth as possible, you can override the 75 percent maximum allocation for the bandwidth sum allocated to all classes or flows by using the **max-reserved-bandwidth** command. If you want to override the fixed amount of bandwidth, exercise caution and ensure that you allow enough remaining bandwidth to support best-effort and control traffic, and Layer 2 overhead.

You can configure strict PQ with WFQ by using the IP RTP priority queuing feature. Strict PQ allows delay-sensitive data to be dequeued and sent before packets in other queues are dequeued.

This table presents the two basic commands for configuring IP RTP priority.

Command	Description
ip rtp priority *starting-rtp-port-number port-number-range bandwidth*	Reserves a strict-priority queue for a set of RTP packet flows belonging to a range of UDP destination ports.
max-reserved-bandwidth *percent*	Changes the percent of bandwidth allocated for LLQ and IP RTP priority.

To display information about fair queuing configuration, use the command in this table.

Command	Purpose
show queue *interface-type interface-number*	Lists fair queuing configuration and statistics for a particular interface.

Example: Configuring IP RTP Priority

The example first defines an IP RTP priority configuration and then reserves a strict-priority queue with these values: a starting RTP port number of 16,384, a range of 16,383 UDP ports, and a maximum bandwidth of 40 kbps:

```
! The following command reserves a strict priority queue:
Switch(config-if)#ip rtp priority 16384 16383 40
```

Example: Configuring max-reserved-bandwidth

In this example, the **max-reserved-bandwidth** command changes the maximum bandwidth allocated between LLQ and IP RTP priority from the default (75 percent) to 80 percent.

```
Switch(config)#multilink virtual-template 1
Switch(config)#interface virtual-template 1
Switch(config-if)#ip address 172.16.1.1 255.255.255.0
Switch(config-if)#no ip directed-broadcast
Switch(config-if)#ip rtp priority 16384 16383 25
Switch(config-if)#service-policy output policy1
Switch(config-if)#ppp multilink
Switch(config-if)#ppp multilink fragment-delay 20
Switch(config-if)#ppp multilink interleave
Switch(config-if)#max-reserved-bandwidth 80
Switch(config-if)#end
```

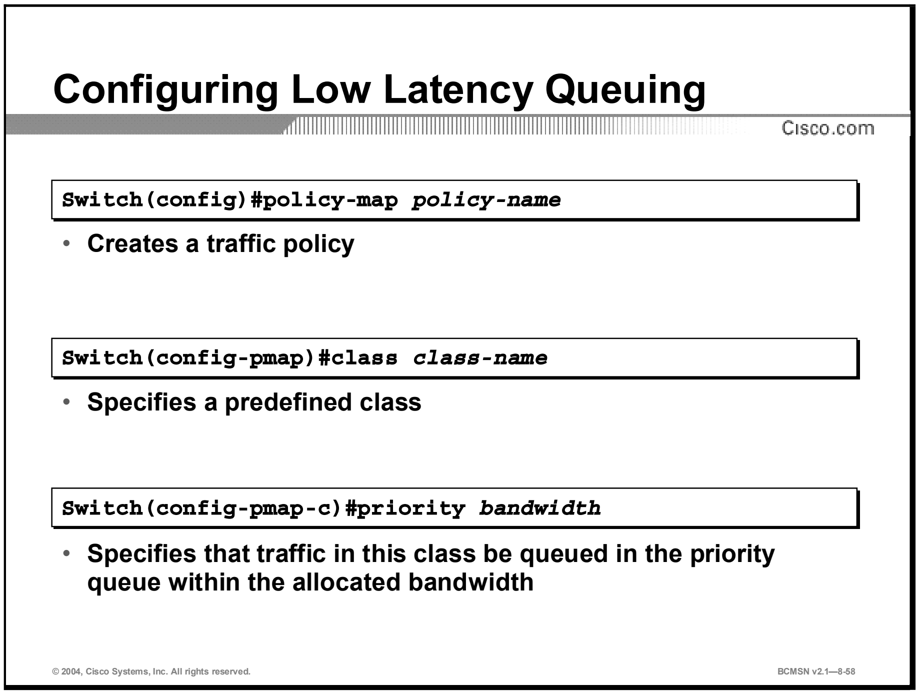

Configuring Low Latency Queuing

Cisco.com

```
Switch(config)#policy-map policy-name
```

- Creates a traffic policy

```
Switch(config-pmap)#class class-name
```

- Specifies a predefined class

```
Switch(config-pmap-c)#priority bandwidth
```

- Specifies that traffic in this class be queued in the priority queue within the allocated bandwidth

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—8-58

The distributed LLQ feature brings the ability to specify low latency behavior for a traffic class. LLQ allows delay-sensitive data to be dequeued and sent first, before packets in other queues are dequeued, giving delay-sensitive data preferential treatment over other traffic.

The **priority** command is used to allow delay-sensitive data to be dequeued and sent first. LLQ enables use of a single priority queue within which individual classes of traffic are placed.

LLQ offers these features:

- LLQ supports multiple traffic types over various Layer 2 technologies, including high-level data link control (HDLC), PPP, ATM, and Frame Relay.
- All classes are policed to bandwidth to ensure that other traffic is serviced.
- The rate limit is per class, even if multiple classes point traffic to a priority queue.
- Oversubscription of bandwidth is not allowed for the priority class.
- WRED support is not provided on priority classes. WRED is allowed only on bandwidth classes.
- Bandwidth and priority are mutually exclusive.

To queue class traffic to the priority queue, you configure the **priority** command for the class after you specify the named class within a policy map. Within a policy map, you can give one or more classes priority status. When multiple classes within a single policy map are configured as priority classes, all traffic from these classes is queued to the same priority queue. CBWFQ handles other traffic, such as data.

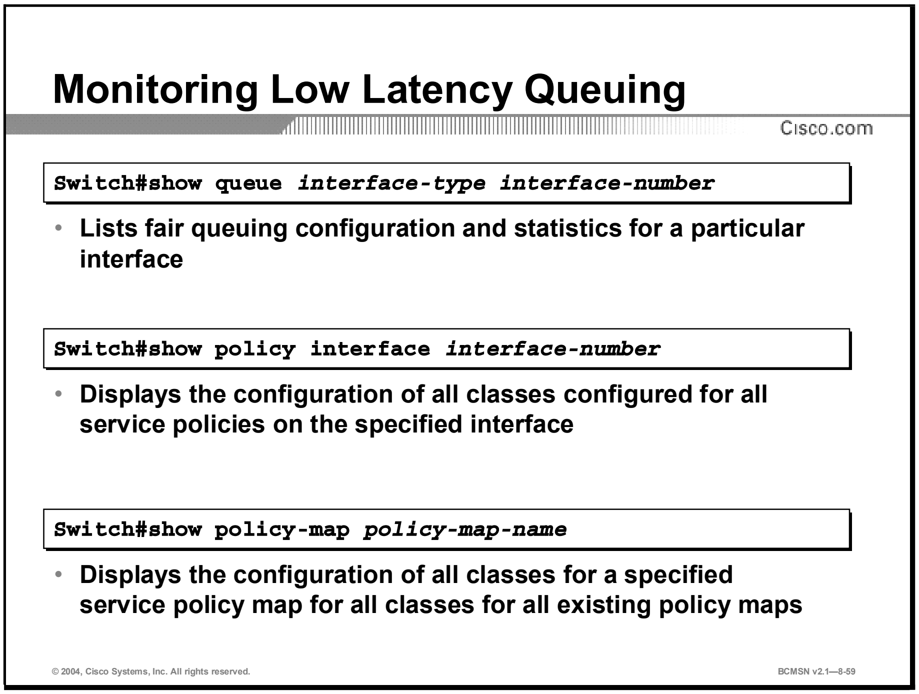

Monitoring Low Latency Queuing

```
Switch#show queue interface-type interface-number
```

- **Lists fair queuing configuration and statistics for a particular interface**

```
Switch#show policy interface interface-number
```

- **Displays the configuration of all classes configured for all service policies on the specified interface**

```
Switch#show policy-map policy-map-name
```

- **Displays the configuration of all classes for a specified service policy map for all classes for all existing policy maps**

BCMSN v2.1—8-59

To tune your RTP bandwidth or decrease RTP traffic if the priority queue is experiencing drops, use one or more of the commands in this table.

Command	Purpose
show queue *interface-type interface-number*	Lists fair queuing configuration and statistics for a particular interface.
show policy interface *interface-name*	Displays the configuration of all classes configured for all service policies on the specified interface. Shows if packets and bytes were discarded or dropped for the priority class in the service policy attached to the interface.
show policy-map *policy-map-name*	Displays the configuration of all classes for a specified service policy map or all classes for all existing policy maps.

Example: LLQ with IP RTP Priority

You can configure LLQ and IP RTP priority at the same time, but IP RTP priority takes precedence. To demonstrate how they work together, consider the following configuration:

```
class-map match-any control
  match ip precedence 3

class-map match-any voice
  match ip precedence 5
!
!
policy-map pol64branch
  class voice
    priority 32
  class control
   bandwidth 8
  class class-default
   fair-queue
!
interface Serial0/0
 bandwidth 64
 no ip address
 encapsulation frame-relay
 no ip mroute-cache
 load-interval 30
 no keepalive
 max-reserved-bandwidth 90
 cdp enable
 frame-relay traffic-shaping
 no frame-relay inverse-arp
!
interface Serial0/0.1 point-to-point
 bandwidth 64
 ip address 10.80.1.1 255.255.255.252
 no ip mroute-cache
 frame-relay class LLQ64k
 frame-relay interface-dlci 100
 frame-relay ip rtp header-compression
!
map-class frame-relay LLQ64k
 no frame-relay adaptive-shaping
 frame-relay cir 60000
 frame-relay bc 640
 frame-relay be 0
 frame-relay mincir 60000
 service-policy output pol64branch
 frame-relay fragment 80
```

In the example, class maps are defined for call control traffic with an IP precedence of 3 and for voice with an IP precedence of 5. The traffic policy specifies that 32 kbps of the available bandwidth be allocated for the voice traffic and that such traffic be sent to the priority queue. Call control traffic is allocated 8 kbps of bandwidth with class-based queuing. All remaining traffic is handled with fair queuing.

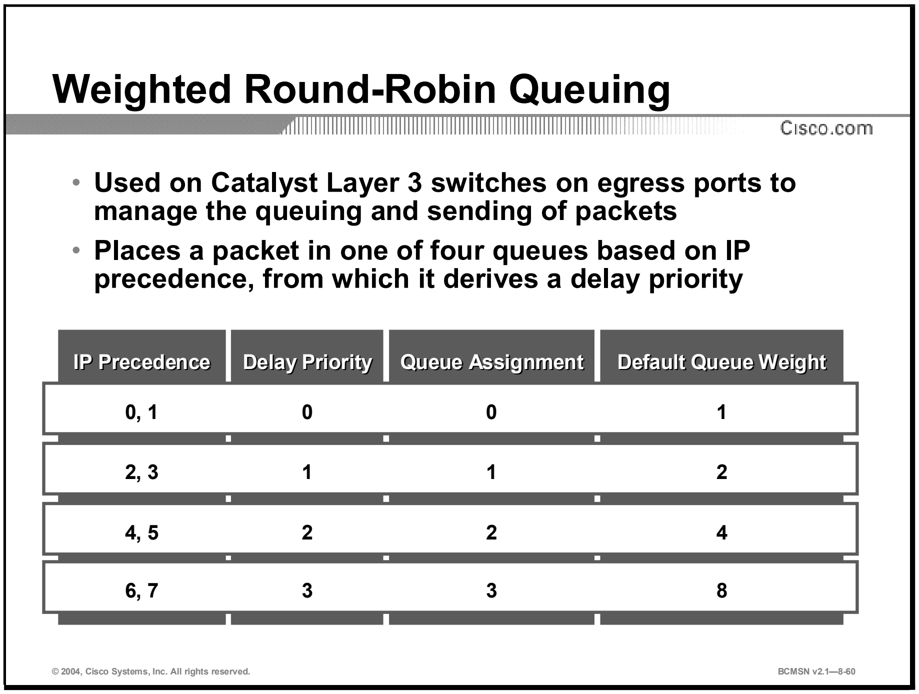

Weighted Round-Robin Queuing

Cisco.com

- Used on Catalyst Layer 3 switches on egress ports to manage the queuing and sending of packets
- Places a packet in one of four queues based on IP precedence, from which it derives a delay priority

IP Precedence	Delay Priority	Queue Assignment	Default Queue Weight
0, 1	0	0	1
2, 3	1	1	2
4, 5	2	2	4
6, 7	3	3	8

BCMSN v2.1—8-60

WRR scheduling is used on Catalyst Layer 3 switches on egress ports to manage the queuing and sending of packets. WRR places a packet in one of four queues based on IP precedence, from which it derives a delay priority. The figure shows the queue assignments based on the IP precedence value, derived delay priority of the packet, and the weight of the queue if you do not change it.

The Layer 3 switches automatically use WRR on egress ports. Unlike other queuing properties, you do not configure WRR through the device interface properties. Instead, you configure WRR through policies defined on the device level.

With WRR, each queue is given a weight. This weight is used when congestion occurs on the port to give weighted priority to high-priority traffic without starving low-priority traffic. The weights provide the queues with an implied bandwidth for the traffic on the queue. The higher the weight, the greater the implied bandwidth. The queues are not assigned specific bandwidth, however, and when the port is not congested, all queues are treated equally.

Devices that use WRR automatically create the four queues with default weights for each interface. You need only define policies on the device if you want to change the queue weights for an interface. QoS Policy Manager (QPM) does not display the interfaces for Layer 3 switches.

WRR is sensitive to the IP precedence settings in the packets. WRR automatically places the packets in queues based on precedence. Although you cannot change the color of a packet on a Layer 3 switch, if you change the packet color on another device before it reaches the Layer 3 switch, that change affects the WRR queuing.

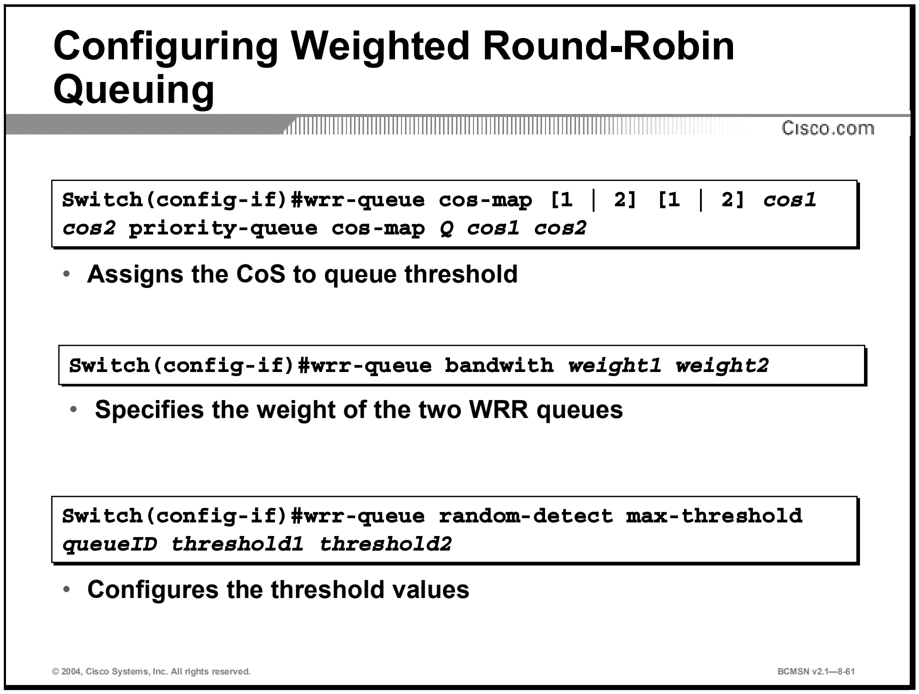

Configuring Weighted Round-Robin Queuing

```
Switch(config-if)#wrr-queue cos-map [1 | 2] [1 | 2] cos1
cos2 priority-queue cos-map Q cos1 cos2
```

- **Assigns the CoS to queue threshold**

```
Switch(config-if)#wrr-queue bandwith weight1 weight2
```

- **Specifies the weight of the two WRR queues**

```
Switch(config-if)#wrr-queue random-detect max-threshold
queueID threshold1 threshold2
```

- **Configures the threshold values**

BCMSN v2.1—8-61

To configure WRR scheduling, perform the tasks in this table.

Step	Command	Notes and Comments
1.	Enable QoS. Switch(config)#**mls qos**	The first thing to do is to enable QoS. Remember that QoS is disabled by default. When QoS is disabled, whatever you have configured as CoS mapping will not affect the outcome. There is one single queue served in a FIFO manner, and all packets get dropped there.
2.	Select an interface to configure. Switch(config)#**interface** {**vlan** *vlan_ID* | {**fastethernet** | **gigabitethernet**} *slot/interface* | **port-channel** *number*}	Select the interface to configure on the Catalyst switch.
3.	Assign the CoS to queue threshold. Switch(config-if)#**wrr-queue cos-map <Q number (***1-2***)> <threshold_number (***1-2***)> <***cos value 1***> <***cos value 2***> priority-queue cos-map <***Q number***> <***cos value 1***> <***cos value 2***>**	You must assign the CoS to queue threshold for all queue types. The queues are always numbered starting with the lowest-priority queue possible, and ending with the strict-priority queue that is available. For example: ■ Queue 1 will be the low-priority WRR queue. ■ Queue 2 will be the high-priority WRR queue. ■ Queue 3 will be the strict-priority queue. Repeat this operation for all types of queues, or you will keep the default CoS assignment.
4.	Configure the WRR weight for the two WRR queues.	Weight 1 relates to queue 1, which should be the low-priority WRR queue. Keep this weight at one level lower than weight 2. The weight

Step	Command	Notes and Comments
	Switch(config-if)#**wrr-queue bandwidth** *<weight for Q1> <weight for Q2>*	can take any value between 1 and 255. Assign the percentage by using the following formulas: ■ To queue 1: [weight 1 / (weight 1 +weight 2)] ■ To queue 2: [weight 2 / (weight 1 +weight 2)] You must define the weight for all types of queues. These weight types do not need to be the same.
5.	Define the Tx queue ratio. Switch(config-if)#**wrr-queue bandwidth** *<weight for Q1> <weight for Q2>*	This ratio determines the way that the buffers are split among the different queues. If you have three queues, you need to set the same level for the high-priority WRR queues and for the strict-priority queue. These levels cannot be different for hardware reasons. Only the bandwidth for the two WRR queues is configured, and you should automatically use the same value as the high WRR queue for the strict-priority queue, if there is any.
6.	Configure the threshold level for the WRED queue or for the tail-drop queue. Switch(config-if)#**wrr-queue random-detect max-threshold <***Q number***> <***threshold 1 value***> <***threshold 2 value***>**	

Configuring Congestion Avoidance

The Cisco WRED combines the capabilities of RED with IP precedence. This combination provides for preferential traffic handling for higher-priority packets. WRED can selectively discard lower-priority traffic when the interface begins to get congested, and provide differentiated performance characteristics for different classes of service. WRED differs from other congestion management techniques, such as queuing, because it attempts to anticipate and avoid congestion rather than controlling congestion after it occurs. This topic explains how to configure congestion avoidance.

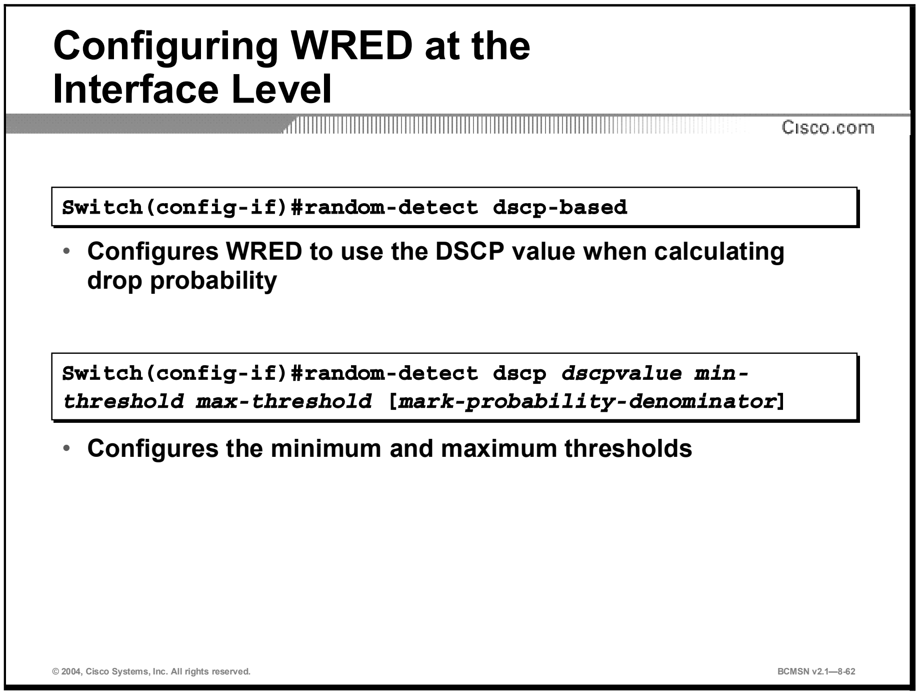

Configuring WRED at the Interface Level

Cisco.com

```
Switch(config-if)#random-detect dscp-based
```

- **Configures WRED to use the DSCP value when calculating drop probability**

```
Switch(config-if)#random-detect dscp dscpvalue min-
threshold max-threshold [mark-probability-denominator]
```

- **Configures the minimum and maximum thresholds**

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—8-62

You can implement WRED at the interface level, at the virtual circuit (VC) level, or at the class level (as part of CBWFQ with policy maps).

To configure WRED to use the DSCP value when it calculates the drop probability, use the commands in this table, beginning in interface configuration mode.

Step	Command	Notes and Comments

1.	Indicate that WRED is to use the DSCP value when it calculates the drop probability for the packet.
	`Switch(config-if)#`
	`random-detect dscp-based`
		Use the **no** form of this command to disable WRED.

2.	Specify the minimum and maximum thresholds.
	`Switch(config-if)#`
	`random-detect dscp` *`dscpvalue min-threshold max-threshold`* `[`*`mark-probability-denominator`*`]`
		Optionally, you can configure the *mark-probability-denominator* for the specified DSCP value.

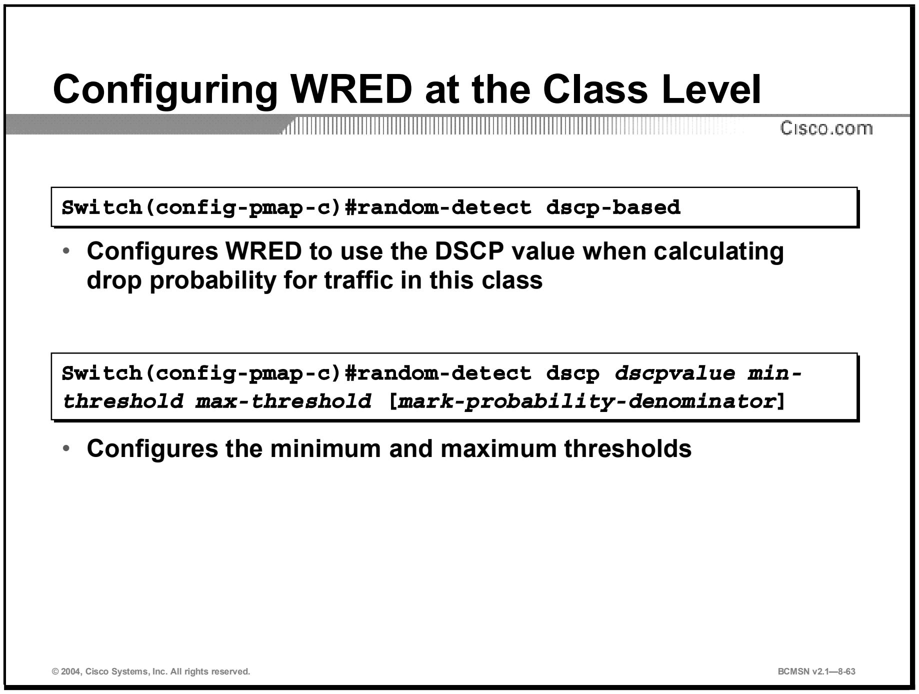

Configuring WRED at the Class Level

```
Switch(config-pmap-c)#random-detect dscp-based
```

- **Configures WRED to use the DSCP value when calculating drop probability for traffic in this class**

```
Switch(config-pmap-c)#random-detect dscp dscpvalue min-
threshold max-threshold [mark-probability-denominator]
```

- **Configures the minimum and maximum thresholds**

BCMSN v2.1—8-63

To configure WRED to use the DSCP value when it calculates the drop probability, use the following commands beginning in interface configuration mode. The table lists the commands to use at the class level, within policy maps.

Step	Command	Purpose
1.	Create a traffic policy by associating the traffic class with one or more QoS features. `Switch(config)#`**`policy-map`** *`policy-name`*	Specify the traffic policy name using the **policy-map** command.
2.	Specify the name of a predefined class. `Switch(config-pmap)#`**`class`** *`class-name`*	The name of a predefined class was defined with the **class-map** command, included in the service policy.

3.	Indicate that WRED is to use the DSCP value when it calculates the drop probability for the packet.
	`Switch(config-pmap-c)#`
	`random-detect dscp-based`
		Use the **no** form of this command to disable WRED.

4.	Specify the minimum and maximum thresholds.
	`Switch(config-pmap-c)#`
	`random-detect dscp` *`dscpvalue min-threshold max-threshold`* `[`*`mark-probability-denominator`*`]`
		Optionally, you can configure the *mark-probability-denominator* for the specified DSCP value.

5.	Attach the traffic policy to the interface.
	`Switch(config-if)#`**`service-policy`**
	`{`**`input`** `|` **`output`**`}` *`policy-name`*
		Attach a traffic policy to an interface and specify the direction in which the policy should be applied (either on packets coming into the interface or packets leaving the interface).

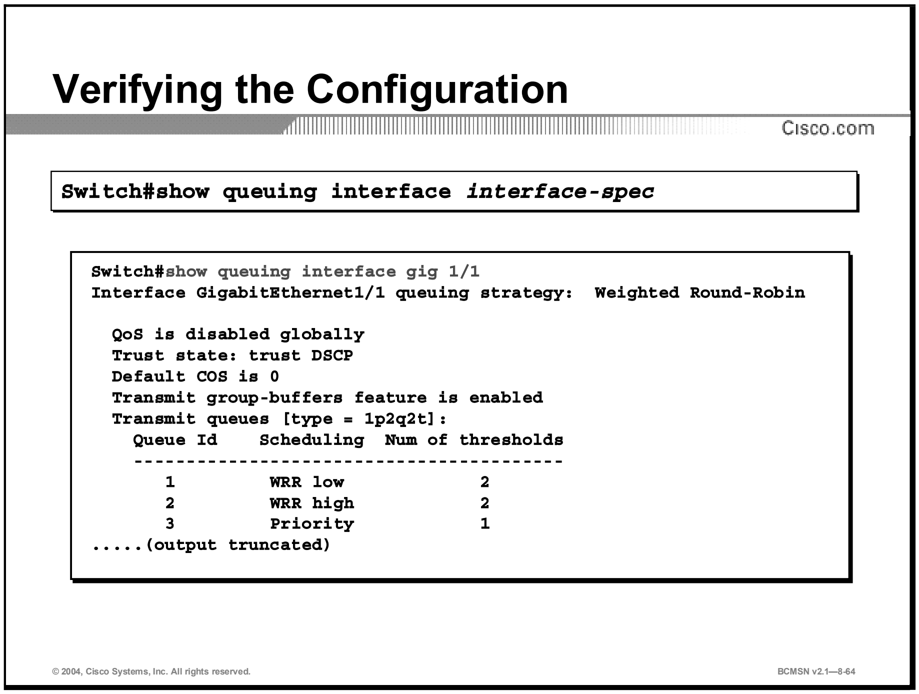

Verifying the Configuration

Cisco.com

```
Switch#show queuing interface interface-spec
```

```
Switch#show queuing interface gig 1/1
Interface GigabitEthernet1/1 queuing strategy:  Weighted Round-Robin

  QoS is disabled globally
  Trust state: trust DSCP
  Default COS is 0
  Transmit group-buffers feature is enabled
  Transmit queues [type = 1p2q2t]:
    Queue Id    Scheduling  Num of thresholds
    -----------------------------------------
       1         WRR low            2
       2         WRR high           2
       3         Priority           1
.....(output truncated)
```

© 2004, Cisco Systems, Inc. All rights reserved.   BCMSN v2.1—8-64

To verify the DSCP value configuration, use either of the commands in this table in global configuration mode.

Command	Description
show queuing interface	Displays the queuing statistics of an interface or virtual circuit (VC).
show policy-map *interface*	Displays the configuration of classes configured for service policies on the specified interface or PVC

Example: Configuring WRED to Use the DSCP Value

The example enables WRED to use the DSCP value 8 for class c1. The minimum threshold for the DSCP value 8 is 24, and the maximum threshold is 40. The last line attaches the service policy to the output interface or VC p1.

```
Switch(config)#class-map c1
Switch(config-cmap)#match access-group 101
Switch(config)#policy-map p1
Switch(config-pmap)#class c1
Switch(config-pmap-c)#bandwidth 48
Switch(config-pmap-c)#random-detect dscp-based
Switch(config-pmap-c)#random-detect dscp 8 24 40
Switch(config-pmap-c)#interface Serial/0
Switch(config-if)#service-policy output p1
```

Troubleshooting the QoS Configuration

Most QoS troubleshooting issues are results of misconfiguration. IOS software provides **debug** commands to help you troubleshoot common problems. This topic explains how to troubleshoot QoS configurations using the **debug** commands.

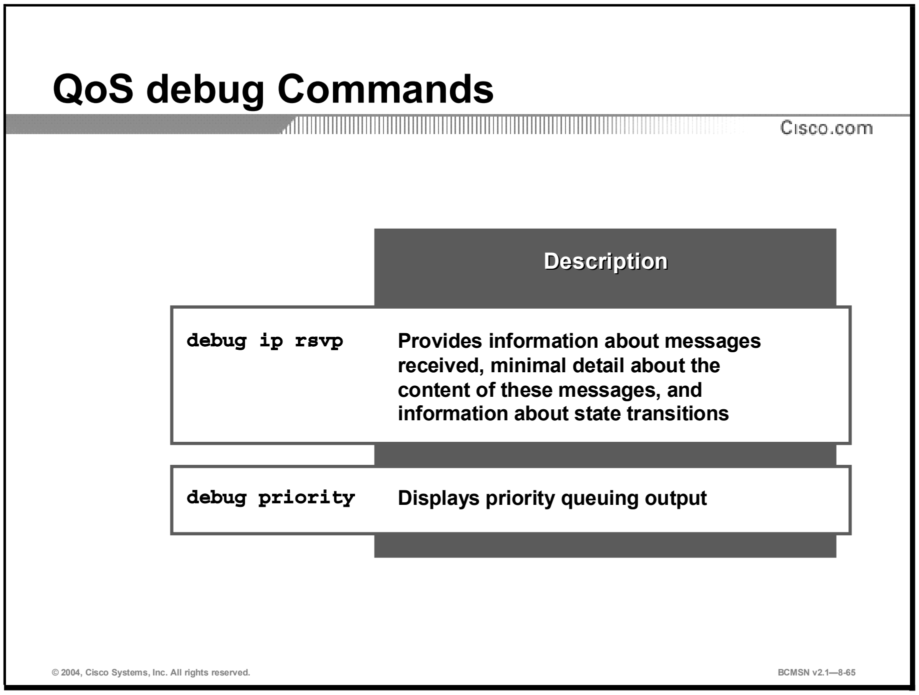

QoS debug Commands

	Description
debug ip rsvp	**Provides information about messages received, minimal detail about the content of these messages, and information about state transitions**
debug priority	**Displays priority queuing output**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—8-65

Use the commands in this table to debug QoS features.

Command	Description
debug ip rsvp	Provides information about messages received, minimal detail about the content of these messages, and information about state transitions.
debug priority	Displays PQ output

Summary

This topic summarizes the key points discussed in this lesson.

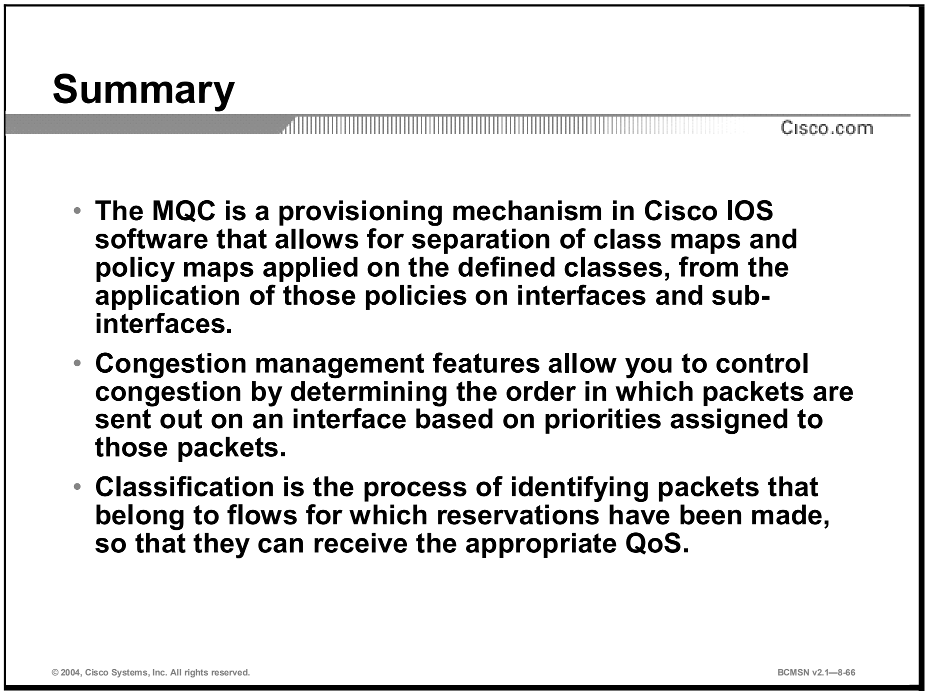

Summary

- The MQC is a provisioning mechanism in Cisco IOS software that allows for separation of class maps and policy maps applied on the defined classes, from the application of those policies on interfaces and sub-interfaces.
- Congestion management features allow you to control congestion by determining the order in which packets are sent out on an interface based on priorities assigned to those packets.
- Classification is the process of identifying packets that belong to flows for which reservations have been made, so that they can receive the appropriate QoS.

BCMSN v2.1—8-66

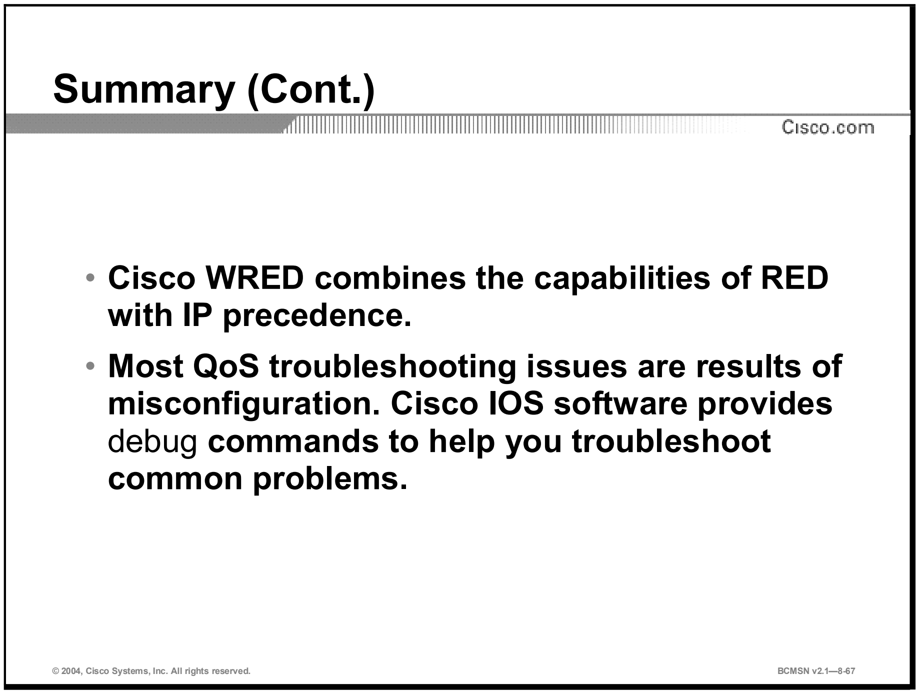

Summary (Cont.)

- Cisco WRED combines the capabilities of RED with IP precedence.
- Most QoS troubleshooting issues are results of misconfiguration. Cisco IOS software provides debug **commands to help you troubleshoot common problems.**

BCMSN v2.1—8-67

References

For additional information, refer to this resource:

- "Cisco IOS Quality of Service" at http://www.cisco.com/warp/public/732/Tech/qos/

Next Steps

For the associated lab exercises, refer to the following section of the course Lab Guide:

- Lab Exercise 8-1: Classifying, Marking, and Implementing QoS Using Policy Maps
- Lab Exercise 8-2: Configuring Egress Queues on Gigabit-Capable Ethernet Ports

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) The _____ is a provisioning mechanism in IOS software that allows for separation of packet classification (class maps) and policies (policy maps) applied on the defined classes, from the application of those policies on interfaces and subinterfaces (service policy).

A) web browser

B) switch console

C) Cisco IOS CLI

D) MQC

Q2) Which IOS command displays the configuration and statistics of the output policy attached to an interface?

A) **show policy-map**

B) **show policy-map** *policy-name*

C) **show policy-map interface** *interface-spec* **output**

D) **show policy map interface** *interface-spec* **output class** *class-name*

Q3) What is the purpose of the ToS field in an IP header?

A) identifies the type of payload in the IP packet

B) identifies network control information for the packet

C) assigns a priority to an IP packet as it traverses the network

D) indicates the proper queue for an IP packet as it traverses the network

Q4) Which three types of queuing can you implement with policy-based routing? (Choose three.)

A) PQ

B) CQ

C) WFQ

D) CBWFQ

E) PQ-WFQ

Q5) Under what two conditions should you use the **set ip precedence** command? (Choose two.)

A) Queuing is enabled and you would like to give video packets a higher priority over other IP data traffic.

B) RSVP is not enabled and you would like to give voice packets a higher priority over other IP data traffic.

C) The Frame Relay link utilization is low, and voice packets do not need to have a higher priority than that of other IP packets.

D) The IP link utilization is high, and the quality of service for voice packets needs to have a higher priority than that of other IP packets.

Q6) Which configuration correctly configures an IP DSCP?

A)
```
policy-map policy1
 class class1
 set ip dscp 5
```

B)
```
class-map class1
 class class1
 set ip dscp 5
```

C)
```
policy-map policy1
 class class1
 set ip dscp 65
```

D)
```
class class1
 policy-map policy1
 set ip dscp 5
```

Q7) On which network links is IP RTP priority especially useful?

A) on fast links whose speed is greater than 1.544 Mbps

B) on slow-speed links whose speed is no more than 64 kbps

C) on slow-speed links whose speed is no more than 784 kbps

D) on slow-speed links whose speed is no more than 1.544 Mbps

Q8) Which value must you specify when you configure the IP RTP priority feature?

A) fair queuing method

B) strict bandwidth limitations

C) packet classification method

D) CAR policies

Q9) Which three features does low latency queuing offer? (Choose three.)

A) makes bandwidth and priority mutually exclusive

B) supports multiple traffic types over any Layer 3 technology

C) polices bandwidth on all classes to ensure that other traffic is serviced

D) provides WRED support on priority classes

E) supports a rate limit per class, even if multiple classes point traffic to a priority queue

Q10) To configure low latency queuing, which command reserves a strict-priority queue for CBWFQ traffic?

A) **fair-queue**

B) **priority** *bandwidth*

C) **class-map** *class-name*

D) **policy-map** *policy-name*

Q11) Which information does the **show queue** *interface-type interface-number* command provide when configuring low latency queuing?

A) lists fair queuing configuration and statistics for a particular interface

B) displays priority queuing output if packets are dropped from the priority queue

C) displays the configuration of all classes configured for all service policies on the specified interface

D) shows if packets and bytes were discarded or dropped for the priority class in the service policy attached to an interface

Q12) WRED generally drops packets selectively based on _____.

A) queue size

B) IP precedence

C) TCP congestion control

D) RED

Q13) When you are configuring WRED at the class level, which command attaches a policy map to an output interface to be used as the service policy for that interface?

A) **policy-map** *policy-map*

B) **random-detect dscp-based**

C) **class-map** *class-map-name*

D) **service-policy output** *policy-map*

Q14) Which IOS command displays priority queuing output?

A) **debug ip rsvp**

B) **debug priority**

C) **debug multilink ppp**

D) **debug ppp multilink fragments**

Quiz Answer Key

Q1) D

Relates to: Configuring the QoS Policy with Modular QoS

Q2) A

Relates to: Configuring the QoS Policy with Modular QoS

Q3) C

Relates to: Configuring Classification and Marking

Q4) A, B, C

Relates to: Configuring Classification and Marking

Q5) B, D

Relates to: Configuring Classification and Marking

Q6) A

Relates to: Configuring Classification and Marking

Q7) D

Relates to: Configuring Queuing and Congestion Management

Q8) B

Relates to: Configuring Queuing and Congestion Management

Q9) A, C, E

Relates to: Configuring Queuing and Congestion Management

Q10) B

Relates to: Configuring Queuing and Congestion Management

Q11) A

Relates to: Configuring Queuing and Congestion Management

Q12) B

Relates to: Configuring Congestion Avoidance

Q13) D

Relates to: Configuring Congestion Avoidance

Q14) B

Relates to: Troubleshooting the QoS Configuration

Lesson Assessments

Overview

Use the lesson assessments here to test what you learned in this module. The correct answers and solutions are found in the Lesson Assessment Answer Key.

Outline

This section includes these assessments:

- Quiz 8-1: Examining the Cisco QoS Solution
- Quiz 8-2: Configuring QoS in Multilayer Switched Networks

Quiz 8-1: Examining the Cisco QoS Solution

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Identify network requirements for QoS
- Describe the IntServ and DiffServ QoS architectures, and explain when to use each
- Identify the classification and marking components of a Cisco QoS solution
- Identify the queuing and congestion management components of a Cisco QoS solution
- Identify the congestion avoidance components of a Cisco QoS solution
- Identify the traffic conditioning components of a Cisco QoS solution
- Identify the link efficiency components of a Cisco QoS solution
- Select QoS solutions for the campus network
- Summarize the key IOS software QoS features and explain when to use each feature

Quiz

Answer these questions:

Q1) Which three tasks do QoS technologies allow you to do? (Choose three.)

A) eliminate packet loss due to congestion

B) eliminate congestion at network convergence points

C) increase the bandwidth of WAN and LAN interfaces

D) manage delay-sensitive traffic, such as real-time voice

E) predict response times for end-to-end network services

F) manage jitter-sensitive applications, such as audio and video playbacks

Q2) What is the difference between IntServ and DiffServ?

A) IntServ is flow-based while DiffServ is type-based.

B) IntServ is type-based while DiffServ is flow-based.

C) IntServ is for voice only, and DiffServ is for data only.

D) IntServ is for data only, and DiffServ is for voice only.

Q3) What is the purpose of the classification QoS feature?

A) to mark packets for deletion

B) to partition traffic into a single priority

C) to partition traffic into multiple priorities

D) to classify packets for insertion into an Rx buffer

Q4) When you are using congestion avoidance mechanisms, when is all traffic discarded?

A) at regular intervals

B) during buffer overflows

C) when a minimum threshold is reached

D) when a maximum threshold is reached

Q5) What is the default congestion management tool for serial links with 256 kbps of bandwidth?

A) LFI

B) LLQ

C) WFQ

D) CBWFQ

Q6) What function does traffic policing perform?

A) discards traffic based on outbound traffic restrictions

B) buffers traffic based on inbound bandwidth restrictions

C) discards traffic based on inbound bandwidth restrictions

D) buffers traffic based on destination bandwidth restrictions

Q7) What can happen to the network if QoS is not applied to the Building Distribution submodule?

A) Nothing, traffic will flow through the switch as it was prioritized at the access layers.

B) While aggregating traffic, the Tx buffer will use a last-in, first-out scheduling mechanism.

C) While aggregating traffic, lower-priority traffic might be sent before high-priority traffic.

D) The Building Distribution submodule will use the QoS configuration at the access switches to ensure QoS compliance.

Q8) What happens when traffic from untrusted sources is allowed in the Campus Backbone submodule?

A) Nothing happens.

B) The untrusted traffic is deleted.

C) It could degrade the anticipated QoS for other traffic.

D) The traffic remains in the queue until the device finishes processing all other traffic.

Q9) When you are designing a QoS solution, what is the first question to ask?

A) Which QoS tools will I use?

B) Which QoS model will I use?

C) What are the critical applications?

D) Where are the QoS features that I will implement?

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Quiz 8-2: Configuring QoS in Multilayer Switched Networks

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Configure the QoS policy using the modular QoS user interface
- Configure classification and marking
- Configure queuing and congestion management
- Configure congestion avoidance
- Troubleshoot the QoS configuration

Quiz

Answer these questions:

Q1) Which IOS command creates a traffic class and specifies that only one match criterion in the traffic class must be met in order to match the specified traffic class?

A) **policy-map policy1**

B) **class-map match-all class1**

C) **class-map match-any class1**

D) **service-policy input policy1**

Q2) What does the **service-policy output policy1** command do?

A) creates a traffic policy called "policy1" and attaches the policy to an interface

B) attaches the traffic policy policy1 to an interface and specifies that the policy should be applied to packets leaving the interface

C) attaches the traffic policy policy1 to an interface and specifies that the policy should be applied to packets entering the interface

D) creates a traffic class where all of the match criteria in the traffic class must be met in order for a packet to match the specified traffic class

Q3) Which IOS command allows you to identify voice traffic by its RTP port numbers and classify it into a priority queue?

A) **ip rtp reserve**

B) **ip rtp priority**

C) **max-reserved-bandwidth**

D) **ip rtp header-compression**

Q4) Which IOS command changes the percent of bandwidth allocated for LLQ and IP RTP priority traffic to 50 percent?

A) **ip rtp reserve**

B) **ip rtp reserve 50**

C) **max-reserved-bandwidth 50**

D) **ip rtp header-compression 50**

Q5) Which feature works well with and uses WFQ to allocate buffer space and schedule packets, and guarantees bandwidth for reserved flows?

A) RTP

B) RSVP

C) MGCP

D) WRED

Q6) Which bits in the ToS field in the IP header are used to specify a class of service assignment for each packet?

A) the first bit

B) the priority bits

C) the last three bits

D) the first three bits

Q7) Which three classification features are supported by network-based application recognition? (Choose three.)

A) classification of applications based on IP precedence

B) classification of application traffic using subport information

C) classification of HTML traffic by URL, MTA, or MIME type

D) classification of applications that dynamically assign TCP and UDP port numbers

E) classification of ICA traffic by application name

Q8) Which IOS command specifies a protocol supported by network-based application routing as a matching criterion?

A) **class**

B) **match protocol**

C) **class-map match-any**

D) **service-policy input**

Q9) When you are using WRED and the average queue size is between the minimum queue threshold and the maximum threshold, what does WRED do?

A) The arriving packet is dequeued.

B) The packet is automatically dropped.

C) The packet is held in a buffer until the threshold returns to normal.

D) The packet is either dropped or queued, depending on the packet drop probability.

Q10) Which configuration correctly configures WRED at the VC level?

A)
```
interface seo/0
random-detect dscp-based
random-detect dscp 8 24 40
```

B)
```
random-detect-group sanjose dscp-based
dscp 9 20 50
random-detect attach sanjose
```

C)
```
class c1
bandwidth 48
random-detect dscp-based
random-detect dscp 8 24 40
```

D)
```
random-detect dscp-based
random-detect dscp 8 24 40
service-policy output p1
```

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Lesson Assessment Answer Key

Quiz 8-1: Examining the Cisco QoS Solution

Q1)	D, E, F
Q2)	A
Q3)	C
Q4)	B
Q5)	C
Q6)	C
Q7)	C
Q8)	C
Q9)	C

Quiz 8-2: Configuring QoS in Multilayer Switched Networks

Q1)	C
Q2)	B
Q3)	B
Q4)	C
Q5)	B
Q6)	D
Q7)	B, D, E
Q8)	B
Q9)	D
Q10)	B

Module 9

Optimizing and Securing Multilayer Switched Networks

Overview

Your multilayer switched network is one of your most valuable assets. In order to gain maximum utilization from your network, you should identify traffic bottlenecks and underutilized resources and then consider implementing optimization features to increase performance. At the same time, you want to make sure that access to your network is secured, not only to protect the network itself but also to protect company assets.

Upon completing this module, you will be able to:

- Enhance multilayer switched networks with hardware modules to provide optimal performance, and to monitor performance
- Secure the multilayer switched network with authentication, and Layer 2 and Layer 3 security

Outline

The module contains these components:

- Optimizing Multilayer Switched Networks
- Securing Multilayer Switched Networks
- Lesson Assessments

Optimizing Multilayer Switched Networks

Overview

Cisco provides software and hardware tools to help you monitor and analyze the traffic on your network, which is the first step to optimizing performance. Traffic monitoring allows you to identify bottlenecks and underutilized resources, and places where you can implement improvements. Port aggregation can be used to increase link bandwidth by combining several physical ports into a single virtual link. Switch Fabric Modules improve switching performance by providing dedicated communication between fabric-enabled switch modules.

Relevance

Optimizing the performance of your multilayer switched network helps maximize utilization of existing resources and reduce unnecessary duplication of network components. This optimization ensures that the network will run efficiently and support multiple solutions.

Objectives

Upon completing this lesson, you will be able to:

- Describe techniques to enhance the performance of a multilayer switched network
- Monitor switch ports using SPAN and VLAN-based SPAN
- Monitor switch ports using RSPAN
- Describe the features and operation of network analysis modules on Catalyst switches to improve network traffic management
- Verify and troubleshoot the operation of network analysis modules

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Techniques to Enhance Performance
- Monitoring Performance with SPAN and VSPAN
- Monitoring Performance with RSPAN
- Network Analysis Modules
- Verifying Network Analysis Modules
- Summary
- Quiz

Techniques to Enhance Performance

Performance management maintains internetwork performance at acceptable levels by measuring and managing various network performance variables. This topic introduces techniques to enhance network performance.

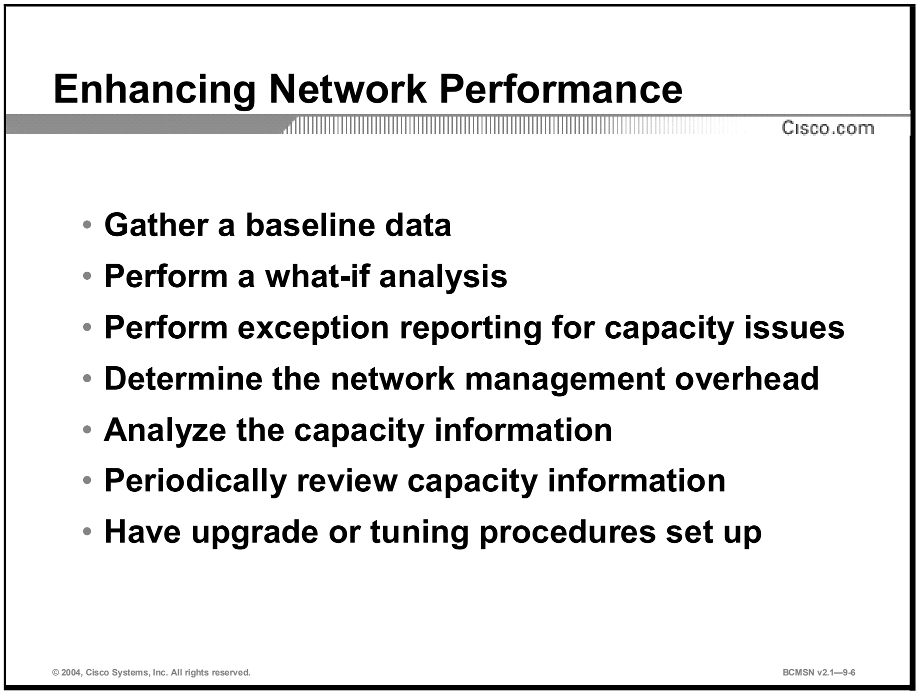

Enhancing Network Performance

Cisco.com

- Gather a baseline data
- Perform a what-if analysis
- Perform exception reporting for capacity issues
- Determine the network management overhead
- Analyze the capacity information
- Periodically review capacity information
- Have upgrade or tuning procedures set up

BCMSN v2.1—9-6

Critical performance management issues are the following:

- **User performance:** For most users, response time is the critical performance success factor. This variable will shape the perception of network success by both your users and application administrators.
- **Application performance:** How quickly a network application responds to user action is a prime indicator of network performance for users.
- **Capacity planning:** Capacity planning is the process of determining the likely future network resource requirements to prevent a performance or availability impact on business-critical applications.
- **Proactive fault management:** Fault management involves both responding to faults as they occur and implementing solutions that prevent faults from impacting performance.

Critical success factors identify the requirements for implementation best practices. To qualify as a critical success factor, a process or procedure must improve availability or the absence of the procedure must decrease availability. The critical success factor should also be measurable so the organization can determine the extent of its success.

The critical success tasks for performance management are as follows:

- **Gather baseline data for both your network and applications:** A typical router and switch baseline report would include capacity issues related to CPU, memory, buffer management, link and media utilization, and throughput. There are other types of baseline data that you may also include, depending on your defined objectives. For instance, an availability baseline would demonstrate the increased stability and availability of the network environment. Perform a baseline comparison between old and new environments to verify solution requirements.
- **Perform a what-if analysis on your network and applications:** A what-if analysis involves modeling and verification of solutions. Before adding a new solution to the network (either a new application or a change in the Cisco IOS release), you should document some of the alternatives. The documentation for this analysis includes the major questions, the methodology, data sets, and configuration files. The main point is that the what-if analysis is an experiment that someone else should be able to re-create with the information provided in the document.
- **Perform exception reporting for capacity issues:** If capacity requirements outstrip available resources, how will you be informed of that fact?
- **Determine the network management overhead for all proposed or potential network management services:** Network management services impact network and application performance. Make sure that your planning and measurement take into account the resources required to perform measurement and management.
- **Analyze the capacity information:** Examine all the data that you have gathered to determine capacity and management requirements.
- **Periodically review capacity information for both the network and applications, as well as baseline and exception:** Make sure that you repeat your analysis periodically to identify changes in network use and growth patterns.
- **Have upgrade or tuning procedures set up to handle capacity issues on both a reactive and longer-term basis:** Set up ahead of time procedures that can be followed to deal with future capacity requirements.

Monitoring Performance with SPAN and VSPAN

Switched Port Analyzer (SPAN) selects and copies network traffic to send to a network analyzer. This topic discusses SPAN.

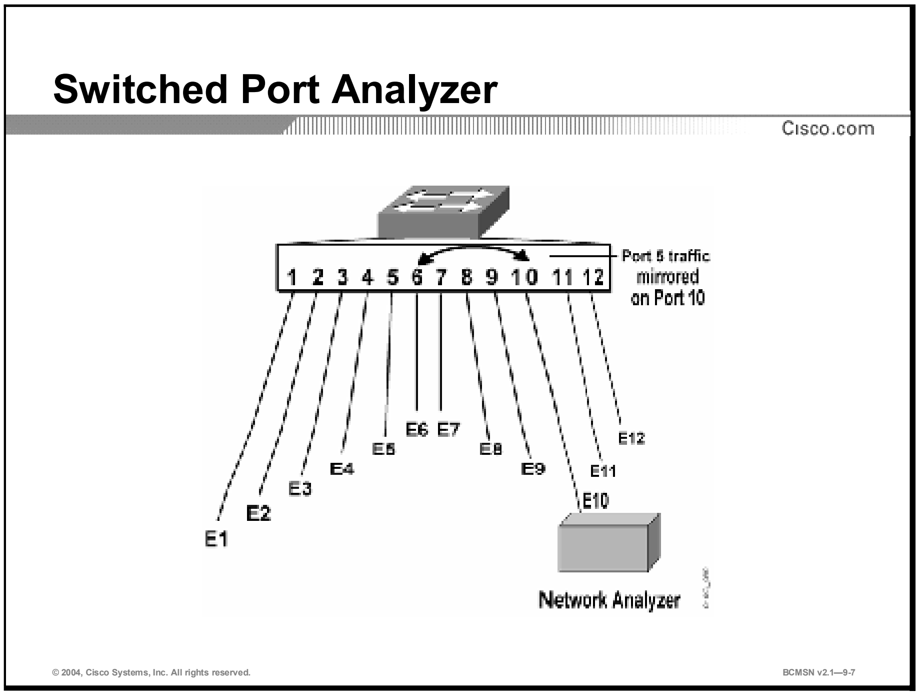

Switched Port Analyzer

BCMSN v2.1—9-7

Local SPAN selects network traffic to send to a network analyzer such as a SwitchProbe device or other Remote Monitoring (RMON) probe. SPAN does not affect the switching of network traffic on source ports or VLANs. SPAN sends a copy of the packets received or transmitted by the source ports and VLANs to the destination port. You must dedicate the destination port for SPAN use.

Local SPAN supports source ports, source VLANs, and destination ports on the same Catalyst switch. Local SPAN copies traffic from one or more source ports in any VLAN or from one or more VLANs to a destination port for analysis. For example, as shown in the figure, all traffic transmitted to Ethernet port 6 (the source port) is copied to Ethernet port 10. A network analyzer on Ethernet port 10 receives all network traffic from Ethernet port 6 without being physically attached to Ethernet port 6.

SPAN sessions allow you to monitor traffic on one or more ports, or on one or more VLANs, and send the monitored traffic to one or more destination ports.

A local SPAN session is an association of a set of source ports and source VLANs with one or more destination ports. You configure a local SPAN session on a single network device. Local SPAN does not have separate source and destination sessions.

You can configure SPAN sessions to monitor ingress network traffic (called ingress SPAN), or to monitor egress network traffic (called egress SPAN), or to monitor traffic flowing in both directions.

Ingress SPAN copies network traffic received by the source ports and VLANs for analysis at the destination port. Egress SPAN copies network traffic transmitted from the source ports and VLANs. When you enter the **both** keyword, SPAN copies the network traffic received and transmitted by the source ports and VLANs to the destination port.

By default, local SPAN monitors all network traffic, including multicast and bridge protocol data unit (BPDU) frames.

A source port is a port monitored for network traffic analysis. You can configure both switched and routed ports as SPAN source ports. SPAN can monitor one or more source ports in a single SPAN session. You can configure source ports in any VLAN. Trunk ports can be configured as source ports and mixed with nontrunk source ports, but SPAN does not copy the encapsulation from a source trunk port.

A destination port is a Layer 2 or Layer 3 LAN port to which SPAN sends traffic for analysis.

When you configure a port as a SPAN destination port using Cisco IOS software (native mode), it can no longer receive any traffic. When you configure a port as a SPAN destination port, the port is dedicated for use by the SPAN feature only. A SPAN destination port does not forward any traffic except that required for the SPAN session.

With Cisco IOS Release 12.1(13)E and later releases, you can configure trunk ports as destination ports, which allows destination trunk ports to transmit encapsulated traffic. With earlier releases, trunk ports stop trunking when you configure them as destination ports.

VSPAN

A source VLAN is a VLAN monitored for network traffic analysis. VLAN-based SPAN (VSPAN) uses a VLAN as the SPAN source. All the ports in the source VLANs become source ports.

Both the SPAN and VSPAN features allow for the destination SPAN port to belong in either the monitored VLAN or another VLAN.

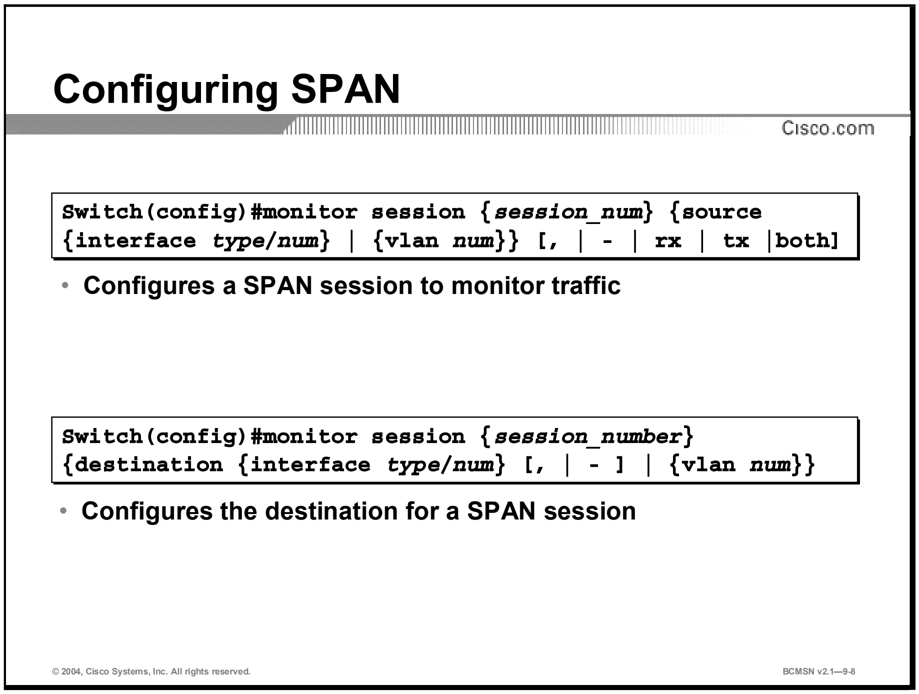

Configuring SPAN

```
Switch(config)#monitor session {session_num} {source
{interface type/num} | {vlan num}} [, | - | rx | tx |both]
```

- Configures a SPAN session to monitor traffic

```
Switch(config)#monitor session {session_number}
{destination {interface type/num} [, | - ] | {vlan num}}
```

- Configures the destination for a SPAN session

BCMSN v2.1—9-8

These guidelines and restrictions apply to local SPAN:

- You need a network analyzer to monitor destination ports.
- You can configure both Layer 2 switched ports (LAN ports configured with the **switchport** command) and Layer 3 ports (LAN ports not configured with the **switchport** command) as sources or destinations.
- A port specified as a destination port in one SPAN session cannot be a destination port for another SPAN session.
- A port configured as a destination port cannot be configured as a source port.
- A port channel interface (an EtherChannel) can be a source.
- With IOS Release 12.1(13)E and later releases, you cannot configure active member ports of an EtherChannel as source ports. Inactive member ports of an EtherChannel can be configured as sources, but they are put into the suspended state and carry no traffic.
- With releases earlier than 12.1(13)E, if you configure a member port of an EtherChannel as a SPAN source port, it is put into the suspended state and carries no traffic.
- A port channel interface (an EtherChannel) cannot be a destination.
- With releases earlier than 12.1(13)E, if you configure a member port of an EtherChannel as a SPAN destination port, it is put into the suspended state and carries no traffic.
- You cannot mix individual source ports and source VLANs within a single session.
- If you specify multiple ingress source ports, the ports can belong to different VLANs.
- You cannot mix source VLANs and filter VLANs within a session. You can have source VLANs or filter VLANs, but not both at the same time.
- When enabled, local SPAN uses any previously entered configuration.
- When you specify sources and do not specify a traffic direction (ingress, egress, or both), “both” is used by default.

- You cannot configure destination ports to receive ingress traffic.
- Destination ports never participate in any spanning tree instance. Local SPAN includes BPDUs in the monitored traffic, so any BPDUs seen on the destination port are from the source port.
- All packets sent through the switch for transmission from a port configured as an egress source are copied to the destination port, including packets that do not exit the switch through the port because STP has put the port into the blocking state, or on a trunk port because STP has put the VLAN into the blocking state on the trunk port.

These are the VSPAN guidelines and restrictions:

- For VSPAN sessions with both ingress and egress configured, two packets are forwarded from the source port if the packets get switched on the same VLAN (one as ingress traffic from the ingress port and one as egress traffic from the egress port).
- VSPAN monitors only traffic that leaves or enters Layer 2 ports in the VLAN.
  - If you configure a VLAN as an ingress source and traffic is routed to the monitored VLAN, the routed traffic will not be monitored because it never appears as ingress traffic entering a Layer 2 port in the VLAN.
  - If you configure a VLAN as an egress source and traffic is routed out of the monitored VLAN, the routed traffic is not monitored because it never appears as egress traffic leaving a Layer 2 port in the VLAN.

To configure a local SPAN session, use the same session number for the sources and the destination ports.

Example: Configuring SPAN

This example shows how to configure session 1 to monitor bidirectional traffic from Fast Ethernet port 5/1:

```
Switch(config)#monitor session 1 source interface fastethernet
5/1 both
```

When configuring a SPAN source port, the specified interfaces can be a single interface, a comma-separated list of interfaces, a range of interfaces, or a combination.

This example shows how to configure Fast Ethernet port 5/48 as the destination for SPAN session 1:

```
Switch(config)#monitor session 1 destination interface
fastethernet 5/48
```

Monitoring Performance with RSPAN

Remote SPAN (RSPAN) is a variation of SPAN that sends monitored traffic through an intermediate switch rather than directly to the Traffic Analyzer. This topic describes RSPAN.

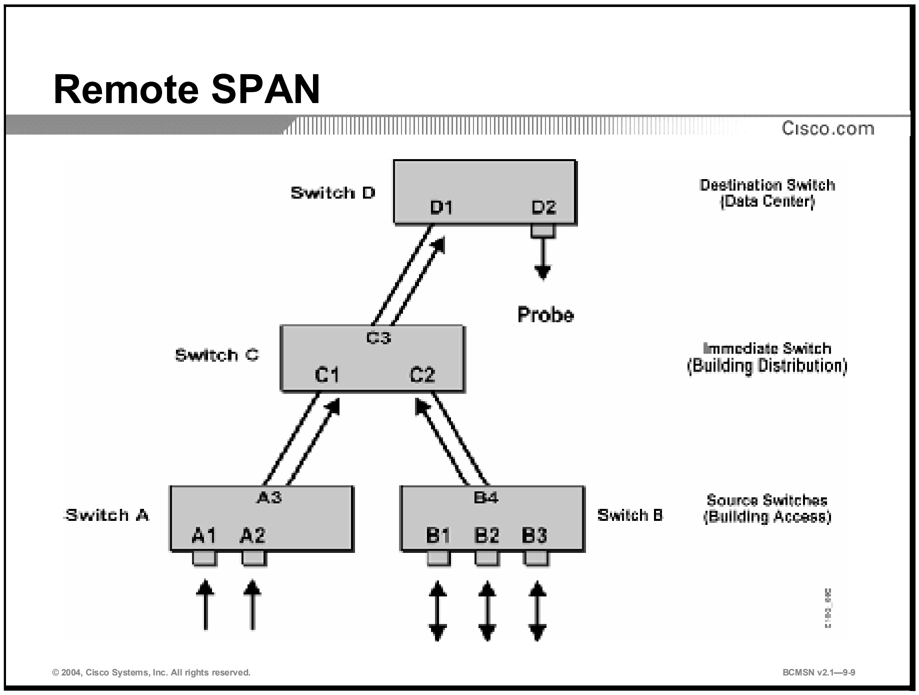

Remote SPAN

BCMSN v2.1—9-9

RSPAN is much like SPAN, but RSPAN supports source ports, source VLANs, and destination ports on different switches. RSPAN provides remote monitoring of multiple switches across your network, as shown in the figure. The traffic for each RSPAN session is carried over a user-specified RSPAN VLAN that is dedicated for that RSPAN session in all participating switches.

The RSPAN source ports can be trunks carrying the RSPAN VLAN. Local SPAN and RSPAN do not monitor the RSPAN traffic in the RSPAN VLAN seen on a source trunk.

The RSPAN traffic from the source ports or source VLANs is switched to the RSPAN VLAN and then forwarded to destination ports, which are in the RSPAN VLAN. The sources (ports or VLANs) in an RSPAN session can be different on different source switches but must be the same for all sources on each RSPAN source switch. Each RSPAN source switch must have either ports or VLANs as RSPAN sources.

RSPAN consists of an RSPAN source session, an RSPAN VLAN, and an RSPAN destination session. You separately configure RSPAN source sessions and destination sessions on different network devices. To configure an RSPAN source session on one network device, you associate a set of source ports and VLANs with an RSPAN VLAN. To configure an RSPAN destination session on another device, you associate the destination port with the RSPAN VLAN.

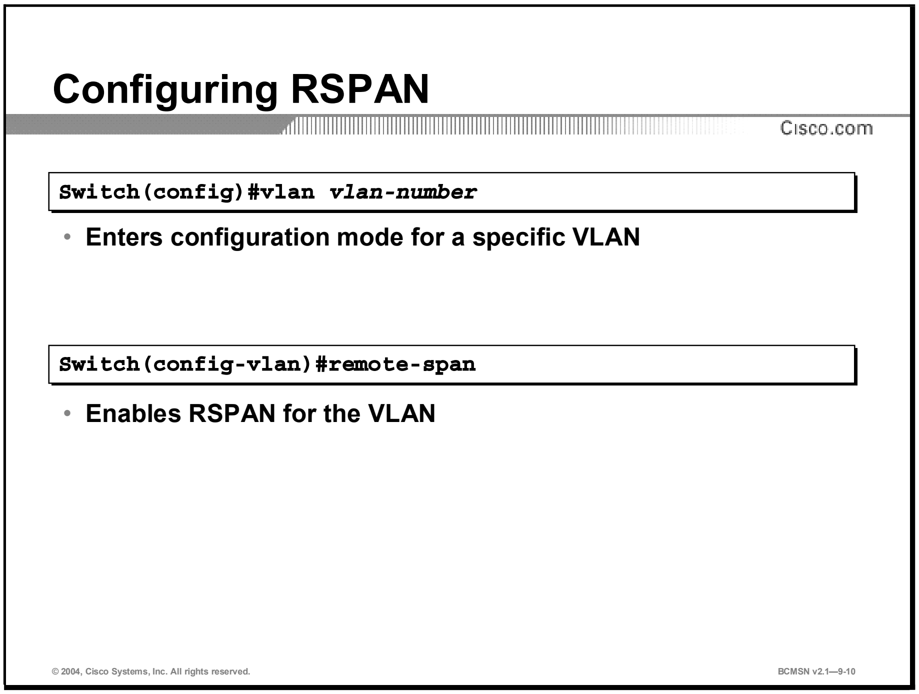

Configuring RSPAN

Cisco.com

```
Switch(config)#vlan vlan-number
```

- Enters configuration mode for a specific VLAN

```
Switch(config-vlan)#remote-span
```

- Enables RSPAN for the VLAN

BCMSN v2.1—9-10

In addition to the guidelines and restrictions that apply to SPAN, these guidelines apply to RSPAN:

- Any network device that supports RSPAN VLANs can be an RSPAN intermediate device.
- Networks impose no limit on the number of RSPAN VLANs that the networks carry.
- Intermediate switches might impose limits on the number of RSPAN VLANs that they can support, based on their capacity.
- You must configure the RSPAN VLANs in all source, intermediate, and destination network devices. If enabled, the VLAN Trunk Protocol (VTP) can propagate configuration of VLANs numbered 1 through 1024 as RSPAN VLANs. You must manually configure VLANs numbered higher than 1024 as RSPAN VLANs on all source, intermediate, and destination network devices.
- If you enable VTP and VTP pruning, RSPAN traffic is pruned in the trunks to prevent the unwanted flooding of RSPAN traffic across the network.
- RSPAN VLANs can be used only for RSPAN traffic.
- Do not configure a VLAN used to carry management traffic as an RSPAN VLAN.
- Do not assign access ports to RSPAN VLANs. RSPAN puts access ports in an RSPAN VLAN into the suspended state.
- Do not configure any ports in an RSPAN VLAN except those selected to carry RSPAN traffic.
- MAC address learning is disabled on the RSPAN VLAN.
- You can use an output access control list (ACL) on the RSPAN VLAN in the RSPAN source switch to filter the traffic sent to an RSPAN destination.
- RSPAN source ports and destination ports must be on different network devices.
- RSPAN does not support BPDU monitoring.
- Do not configure RSPAN VLANs as sources in VSPAN sessions.

- You can configure any VLAN as an RSPAN VLAN as long as all participating network devices support configuration of RSPAN VLANs, and you use the same RSPAN VLAN for each RSPAN session in all participating network devices.
- Entering SPAN configuration commands does not clear previously configured SPAN parameters. You must enter the **no monitor session** command to clear configured SPAN parameters.

Example: Configuring RSPAN

This example shows how to configure RSPAN source session 2:

```
Switch(config)#monitor session 2 source interface
fastethernet1/1 - 3 rx
Switch(config)#monitor session 2 destination remote vlan 901
```

This example shows how to configure an RSPAN source session with multiple sources:

```
Switch(config)# monitor session 2 source interface
fastethernet 5/15 , 7/3 rx
Switch(config)# monitor session 2 source interface
gigabitethernet 1/2 tx
Switch(config)# monitor session 2 source interface port-
channel 102
Switch(config)# monitor session 2 source filter vlan 2 – 3
Switch(config)# monitor session 2 destination remote vlan 901
```

This example shows how to configure an RSPAN destination session:

```
Switch(config)# monitor session 8 source remote vlan 901
Switch(config)# monitor session 8 destination interface
fastethernet 1/2 , 2/3
```

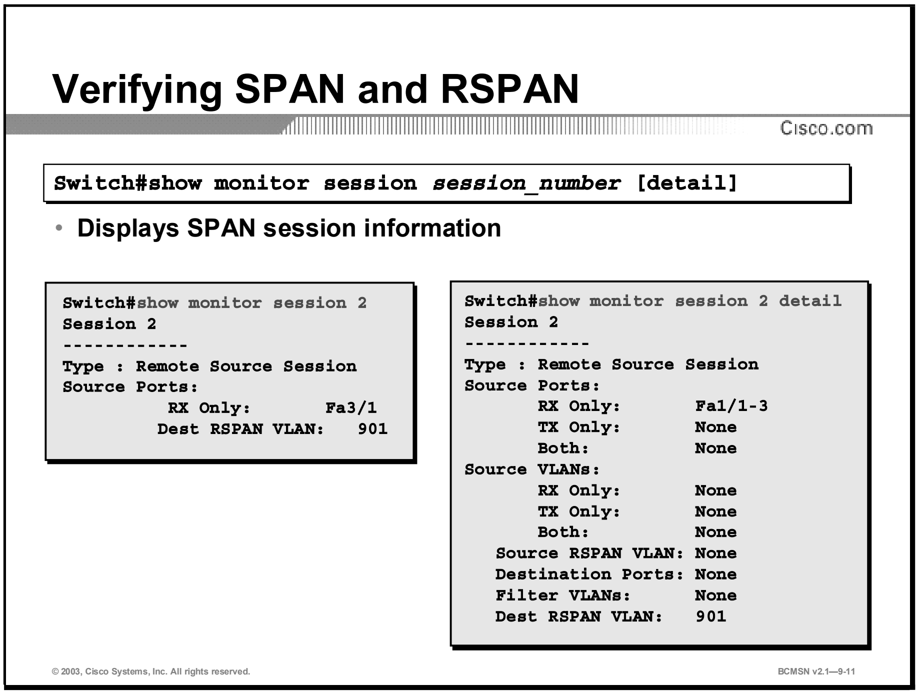

Verifying SPAN and RSPAN

```
Switch#show monitor session session_number [detail]
```

- **Displays SPAN session information**

```
Switch#show monitor session 2
Session 2
------------
Type : Remote Source Session
Source Ports:
             RX Only:       Fa3/1
            Dest RSPAN VLAN:   901
```

```
Switch#show monitor session 2 detail
Session 2
------------
Type : Remote Source Session
Source Ports:
       RX Only:       Fa1/1-3
       TX Only:       None
       Both:          None
Source VLANs:
       RX Only:       None
       TX Only:       None
       Both:          None
   Source RSPAN VLAN: None
   Destination Ports: None
   Filter VLANs:      None
   Dest RSPAN VLAN:   901
```

BCMSN v2.1—9-11

Use the **show monitor** command to verify SPAN and RSPAN configuration.

Example: Verifying SPAN and RSPAN

This example shows how to verify the configuration of session 2:

```
Switch#show monitor session 2
Session 2
------------
Type : Remote Source Session
Source Ports:
          RX Only:       Fa3/1
         Dest RSPAN VLAN:   901
```

This example shows how to display the full details of session 2:

```
Switch#show monitor session 2 detail
Session 2
------------
Type : Remote Source Session
Source Ports:
       RX Only:       Fa1/1-3
       TX Only:       None
       Both:          None
Source VLANs:
       RX Only:       None
       TX Only:       None
       Both:          None
   Source RSPAN VLAN: None
   Destination Ports: None
   Filter VLANs:      None
   Dest RSPAN VLAN:   901
```

Network Analysis Modules

A Network Analysis Module (NAM) uses Simple Network Management Protocol (SNMP) RMON information to monitor and analyze network traffic. This topic describes NAMs.

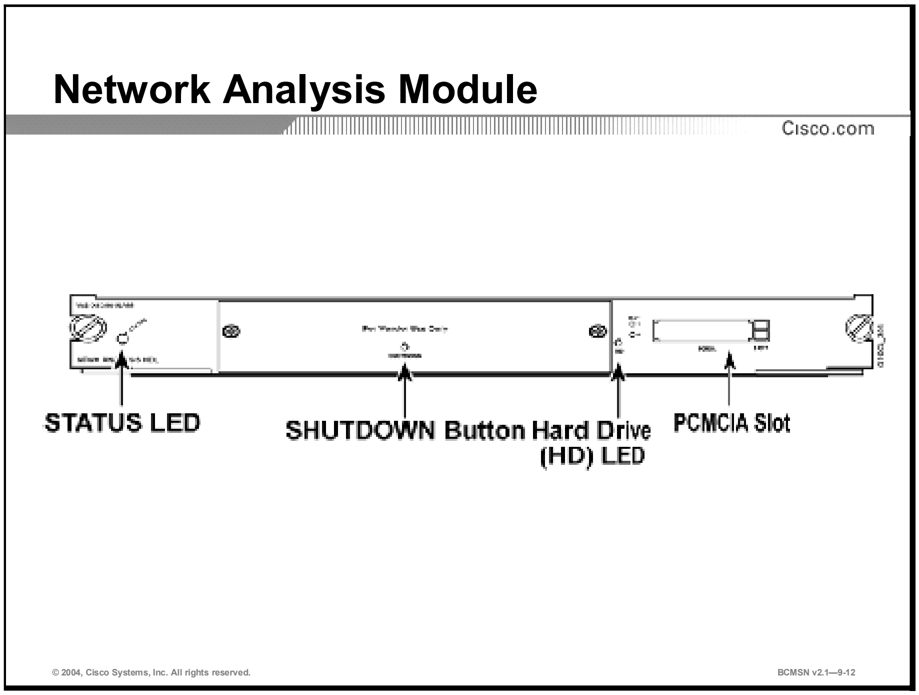

Network Analysis Module

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—9-12

The NAM for the Cisco Catalyst 6500 and 6000 series is part of the end-to-end network management and monitoring solution from Cisco. Managers need to collect statistics about voice or video applications. The NAM gathers multilayer information about data and voice flows that goes all the way to the application layer, helping to simplify the task of managing multiservice switched LANs that support a variety of data, voice, and video applications.

The NAM monitors and analyzes network traffic using RMON, RMON extensions for switched networks, and other Management Information Bases (MIBs).

The NAM supports these RMON groups:

- RMON groups defined in RFC 1757
- RMON2 groups defined in RFC 2021

In addition to extensive MIB support, the NAM can also monitor individual Ethernet VLANs, which allows it to serve as an extension to the basic RMON support provided by the Catalyst Supervisor Engine.

You can use TrafficDirector, or any other Internet Engineering Task Force (IETF)-compliant RMON application, to access link, host, protocol, and response-time statistics for capacity planning, departmental accounting, and real-time application protocol monitoring. You can also use filters and capture buffers to troubleshoot the network.

The NAM can analyze Ethernet VLAN traffic from one or both of these sources:

- Ethernet, Fast Ethernet, Gigabit Ethernet, trunk port, or Fast EtherChannel SPAN or RSPAN source port
- NetFlow Data Export (NDE) (option that makes traffic statistics available for analysis by an external data collector)

The NAM is managed and controlled from either the embedded web-based NAM Traffic Analyzer application or a SNMP management application, such as those bundled with CiscoWorks2000, or both.

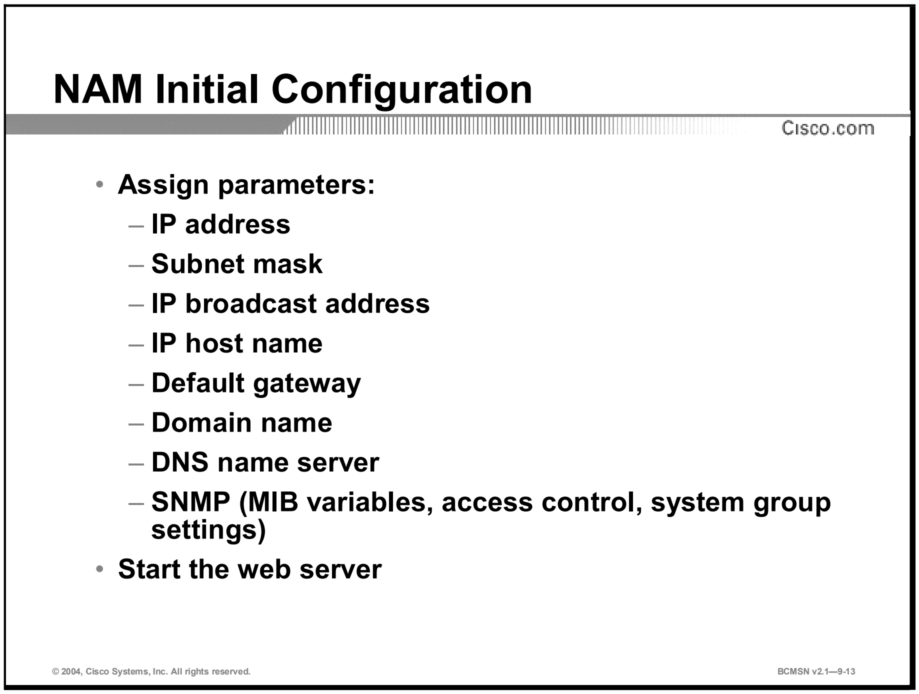

NAM Initial Configuration

- Assign parameters:
  - IP address
  - Subnet mask
  - IP broadcast address
  - IP host name
  - Default gateway
  - Domain name
  - DNS name server
  - SNMP (MIB variables, access control, system group settings)
- Start the web server

BCMSN v2.1—9-13

Before you can use the NAM for network analysis, you must log into the NAM root account and configure the following:

- IP address
- Subnet mask
- IP broadcast address
- IP host name
- Default gateway
- Domain name
- If you are using a Domain Name System (DNS), configure the DNS name server.
- If you are using an external SNMP manager to communicate with the NAM, configure the following:
  - SNMP MIB variables
  - Access control for the SNMP agent
  - System group settings on the NAM
- Start the web server using the **ip http server enable** command.

Example: Configuring the NAM

This example shows the steps that are used to perform the initial configuration of the NAM:

```
Switch#session slot processor 1
Log in as root
root@localhost#ip address ip-address subnet-mask
root@localhost#ip broadcast broadcast-address
root@localhost#ip host name
root@localhost#ip gateway default-gateway
root@localhost#ip domain domain-name
root@localhost#ip nameserver ip-address
root@localhost#snmp location location-string
root@localhost#snmp contact contact-string
root@localhost#snmp name name-string
root@localhost#snmp community community-string rw
root@localhost#snmp community community-string ro
```

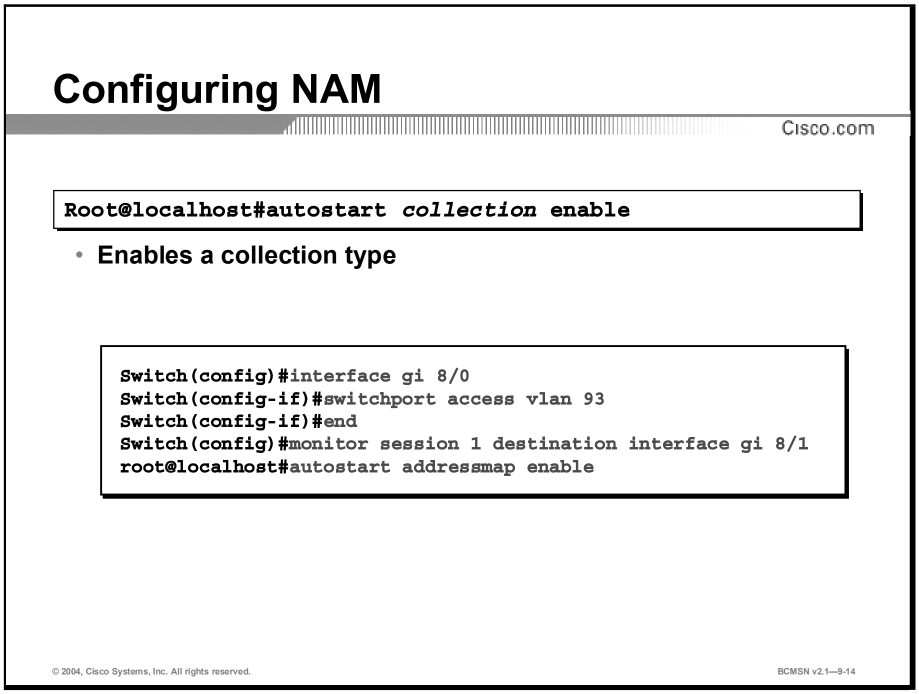

Configuring NAM

Cisco.com

```
Root@localhost#autostart collection enable
```

- **Enables a collection type**

```
Switch(config)#interface gi 8/0
Switch(config-if)#switchport access vlan 93
Switch(config-if)#end
Switch(config)#monitor session 1 destination interface gi 8/1
root@localhost#autostart addressmap enable
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—9-14

You must configure a VLAN for the NAM management port using the **switchport access vlan** *vlan-number* command. You can then configure either the SPAN source port or NDE as a traffic source for the NAM.

To direct SPAN traffic to the NAM for monitoring, configure port 1 on the NAM as the SPAN destination port. You cannot use ports on the NAM module as SPAN source ports.

```
Switch(config)#monitor session 1 destination interface gi 8/1
```

Note The SPAN destination for the NAM must always be port 1.

The NAM can analyze Ethernet traffic from Ethernet, Fast Ethernet, Gigabit Ethernet, trunk port, or Fast EtherChannel SPAN source ports. You can also specify an Ethernet VLAN as the SPAN source.

To use NDE as a traffic source for the NAM, enable the NetFlow Monitor option to allow the NAM to receive the NDE stream. The NAM receives NDE statistics automatically.

NDE makes traffic statistics available for analysis by an external data collector. You can use NDE to monitor all Layer 3 switched and all routed IP unicast traffic.

When you configure a NAM module as an NDE collector, you should use the IP address of the NAM (set up by creating a session on the NAM module).

```
Switch#hw-module module 5 sync nde-info
```

Use the **autostart** command to specify that some RMON collections should be automatically configured on every available data source (including all known VLANs) whenever the NAM is initialized. These collections may also be configured explicitly through SNMP by a management station on some data sources. Collections that are explicitly configured through SNMP take precedence over autostart collections, so if both are configured, only the explicitly configured collections are started on each data source when the NAM initializes.

If you enter the command that instructs the NAM to automatically start a collection, you must reboot the NAM for that command to take effect.

The NAM allows the following collection types to be started automatically:

- **addressMap:** addressMapTable from RMON2-MIB (RFC 2021)
- **art:** artControlTable from draft-warth-rmon2-artmib-01.txt
- **etherStats:** etherStatsTable from RMON-MIB (RFC 1757)
- **prioStats:** smonPrioStatsControlTable from SMON-MIB (RFC 2613)
- **vlanStats:** smonVlanStatsControlTable from SMON-MIB (RFC 2613)

The automatic start process is performed after it sets up any collections that were explicitly created through SNMP by a management station, and stored in the NVRAM in the NAM. Automatic start collections are not configured on data sources that already have a collection of that type configured through SNMP.

Enable a collection type by entering this command from the root account of the NAM:

```
root@localhost#autostart collection enable
```

In this command, *collection* is one of addressMap, prioStats, vlanStats, or etherStats. Disable a collection with the same command, replacing **enable** with **disable**.

After enabling or disabling one or more collection types, you must reboot the NAM before the configuration takes effect.

The Application Response Time (ART) MIB is enabled and disabled globally. When it is enabled, it measures the response time on the network at the transport layer.

Note You must purchase an ART MIB license from Cisco Systems before enabling it and using the ART MIB feature.

To enable the ART MIB, perform this task in privileged EXEC mode:

```
Switch#rmon artmib enable
```

Verifying Network Analysis Modules

Use the **show** commands to verify NAM configuration. This topic describes the **show** commands used to verify the NAM configuration.

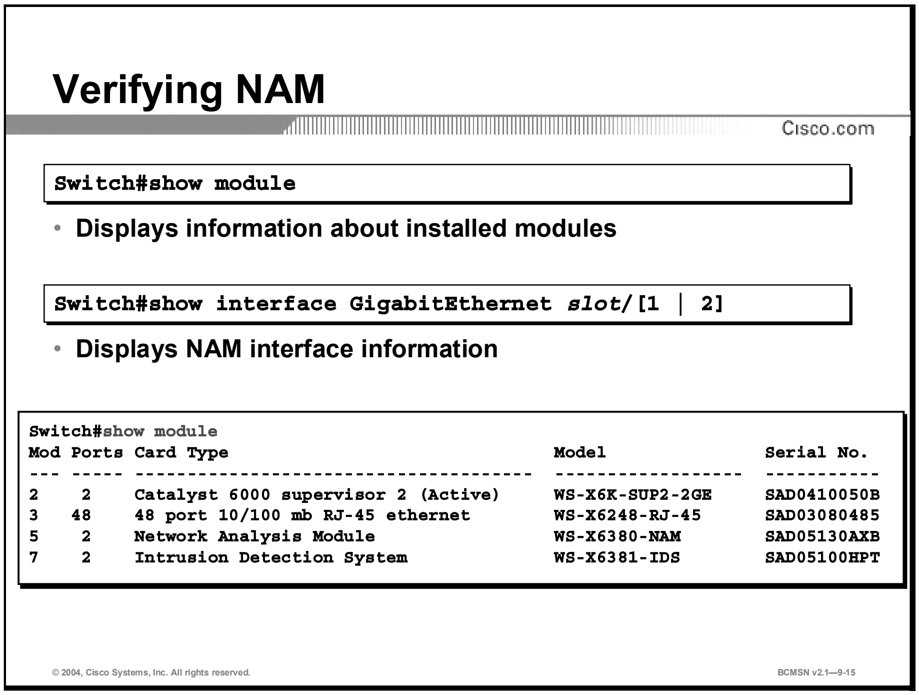

Verifying NAM

Cisco.com

```
Switch#show module
```

- Displays information about installed modules

```
Switch#show interface GigabitEthernet slot/[1 | 2]
```

- Displays NAM interface information

```
Switch#show module
Mod Ports Card Type                              Model              Serial No.
--- ----- -------------------------------------- ------------------ -----------
2    2    Catalyst 6000 supervisor 2 (Active)    WS-X6K-SUP2-2GE    SAD0410050B
3   48    48 port 10/100 mb RJ-45 ethernet       WS-X6248-RJ-45     SAD03080485
5    2    Network Analysis Module                WS-X6380-NAM       SAD05130AXB
7    2    Intrusion Detection System             WS-X6381-IDS       SAD05100HPT
```

© 2004, Cisco Systems, Inc. All rights reserved. BCMSN v2.1—9-15

To verify that the switch acknowledges the new NAM and has brought it online, enter the **show module** command. You can also use the **show interface GigabitEthernet** *slot*/[**1** | **2**] command for the same purpose.

Example: show module Command Output

This example shows the output of the **show module** command:

```
Switch#show module
Mod Ports Card Type                                    Model
Serial No.
--- ----- -------------------------------------      --------
----------      -----------
2    2    Catalyst 6000 supervisor 2 (Active)         WS-X6K-
SUP2-2GE        SAD0410050B
3   48    48 port 10/100 mb RJ-45 ethernet            WS-
X6248-RJ-45         SAD03080485
5    2    Network Analysis Module                     WS-
X6380-NAM           SAD05130AXB
7    2    Intrusion Detection System                  WS-
X6381-IDS           SAD05100HPT
```

The following list describes several potential problems that you might encounter with a NAM, and their possible solutions:

- **Symptom:** The NAM cannot enable the HTTP server.
  Possible Cause: The NAM could not determine the fully qualified domain name of the server.
  Recommended Action: Reboot the NAM.
- **Symptom:** The user cannot connect to the server.
  Possible Cause: The initial configuration is incorrect or not configured.
  Recommended Action: Reconfigure the NAM.
- **Symptom:** The user cannot connect to the NAM Traffic Analyzer application.
  Possible Cause: The configuration for the HTTP server is not correct.
  Recommended Action: Check the NAM configuration for the HTTP server.

Summary

This topic summarizes the key points discussed in this lesson.

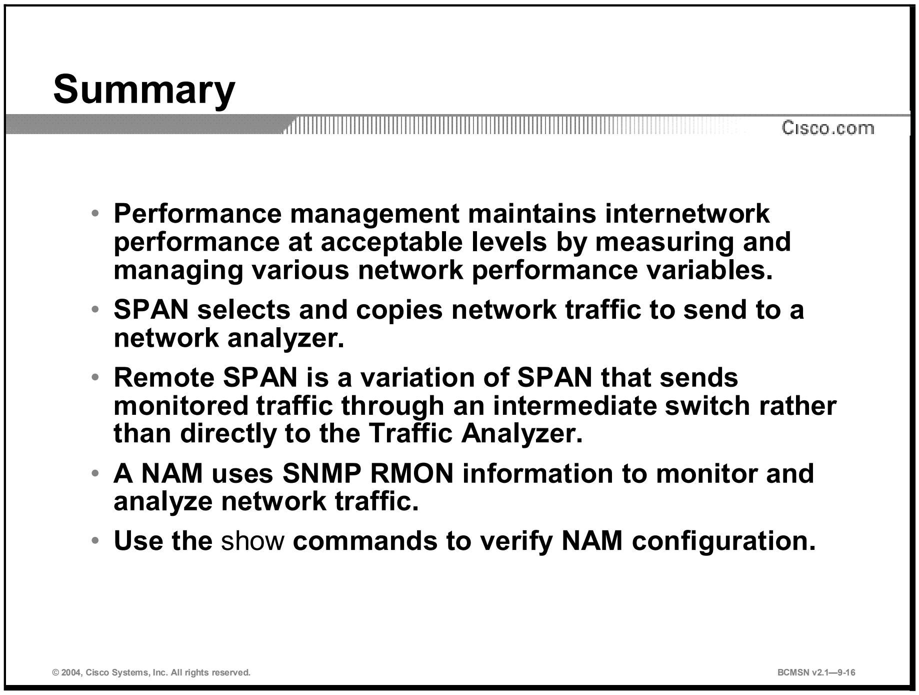

Summary

- **Performance management maintains internetwork performance at acceptable levels by measuring and managing various network performance variables.**
- **SPAN selects and copies network traffic to send to a network analyzer.**
- **Remote SPAN is a variation of SPAN that sends monitored traffic through an intermediate switch rather than directly to the Traffic Analyzer.**
- **A NAM uses SNMP RMON information to monitor and analyze network traffic.**
- **Use the** show **commands to verify NAM configuration.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—9-16

References

For additional information, refer to this resource:

- Your Cisco IOS documentation

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which four performance management issues are critical for a multilayer switched network? (Choose four.)

A) user performance

B) capacity planning

C) application performance

D) inventory and accounting

E) configuration management

F) proactive fault management

Q2) Which command correctly configures a SPAN interface to monitor only ingress traffic?

A) **monitor session 1 source interface fastethernet 5/1 rx**

B) **monitor session 1 source interface fastethernet 5/1 tx**

C) **monitor session 1 destination interface fastethernet 5/1**

D) **monitor session 1 source interface fastethernet 5/1 both**

Q3) Which command correctly configures an RSPAN source from a remote switch destined for a local port on an intermediate device?

A) **monitor session 2 source remote vlan 901**

B) **monitor session 2 source filter vlan 2 – 3**

C) **monitor session 2 destination remote vlan 901**

D) **monitor session 2 source remote interface gigabitethernet 1/2**

Q4) From which two sources can a NAM analyze traffic? (Choose two.)

A) TrafficDirector

B) a SPAN source

C) an RMON application

D) a NetFlow Data Export

E) an SNMP management application

Q5) For which problem is rebooting the NAM a potential solution?

A) NAM cannot enable the HTTP server

B) user cannot connect to the server

C) the NAM cannot receive SPAN traffic

D) user cannot connect to the NAM Traffic Analyzer application

Quiz Answer Key

Q1) A, B, C, F

Relates to: Techniques to Enhance Performance

Q2) A

Relates to: Monitoring Performance with SPAN and VSPAN

Q3) A

Relates to: Monitoring Performance with RSPAN

Q4) B, D

Relates to: Network Analysis Modules

Q5) A

Relates to: Verifying Network Analysis Modules

Securing Multilayer Switched Networks

Overview

Securing the campus network involves the management of physical devices through passwords and access lists as well as securing the physical access to the network through port security. Traffic management control is done by applying access lists to interfaces and routing tables.

Relevance

A network is an extremely important and valuable resource. Ensuring access security helps protect not only your network resources but also a wide variety of company assets.

Objectives

Upon completing this lesson, you will be able to:

- Explain basic security concepts for the multilayer switched network
- Configure authentication, authorization, and accounting on Catalyst switches
- Configure port security and port-based authentication with 802.1x
- Verify the network access security configuration
- Configure VLAN access lists
- Verify the VLAN access list security configuration

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Security in the Multilayer Switched Network
- Configuring AAA
- Configuring Network Access Security
- Verifying Network Access Security
- Configuring Security Using Access Lists
- Configuring PVLANs
- Summary
- Quiz

Security in the Multilayer Switched Network

Network security vulnerabilities include loss of privacy, data theft, impersonation, and loss of integrity. Cisco recommends tasks that you should complete to secure your switched network from attack. This topic introduces security in a multilayer switched network.

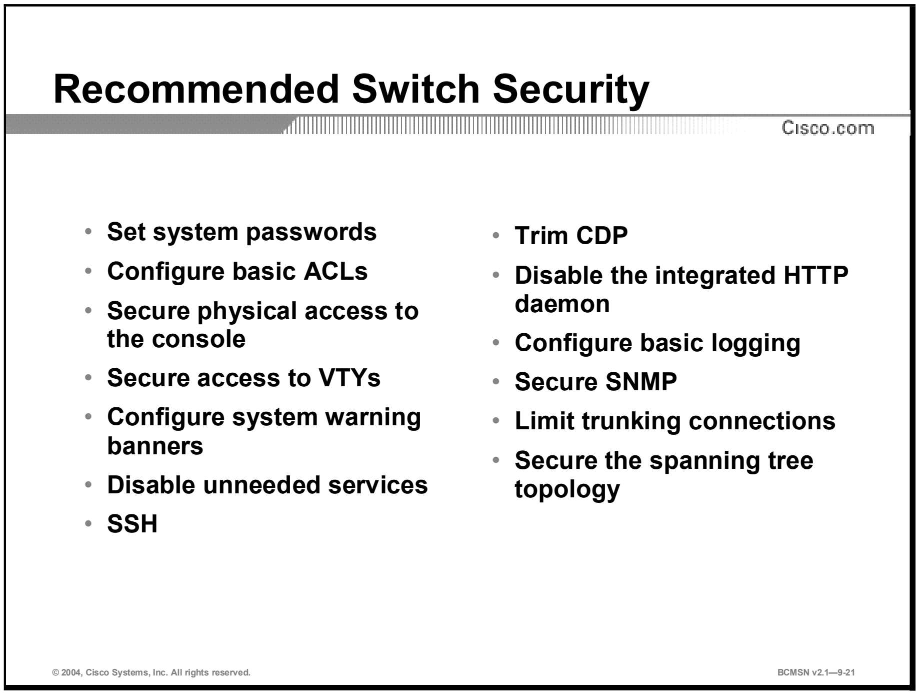

Recommended Switch Security

Cisco.com

- Set system passwords
- Configure basic ACLs
- Secure physical access to the console
- Secure access to VTYs
- Configure system warning banners
- Disable unneeded services
- SSH
- Trim CDP
- Disable the integrated HTTP daemon
- Configure basic logging
- Secure SNMP
- Limit trunking connections
- Secure the spanning tree topology

BCMSN v2.1—9-21

You should implement a basic security configuration on every installed Cisco IOS device. The primary focus is to apply minimal security to mitigate user negligence in regards the network. These techniques are intended to be simple and easy to understand for those interested in implementing a minimum level of security on IOS switches.

Implementing a basic set of security configurations is simple and extremely straightforward. To secure a device, you should do the following:

- **Set system passwords:** Use the **enable secret** command to set the password that grants enabled access to the IOS system. Because the **enable secret** command simply implements a Message Digest 5 (MD5) hash on the configured password, that password still remains vulnerable to dictionary attacks. Therefore, standard practices in selecting a feasible password apply. Try to pick passwords that contain both letters and numbers as well as special characters, for example, $pecia1$ instead of “specials,” where the “s” has been replaced by “$” and the “l” with “1”.
- **Configure basic ACLs:** Create basic ACLs to limit management and remote access traffic. In the following example, a basic access list allows remote access via Telnet from the management network subnet of 192.168.1.0/24 and the specific single host with the IP address of 192.168.5.177.

```
access-list 5 permit 192.168.1.0 0.0.0.255
access-list 5 permit 192.168.5.177 0.0.0.0
```

- **Secure access to the console:** Console access requires a minimum level of security both physically and logically. An individual who gains console access to a system will gain the ability to recover or reset the system-enable password, giving that person the ability to bypass all other security implemented on that system. Consequently, it is imperative to secure access to the console.
- **Secure access to VTYs:** The minimum recommended security to implement for secure access to virtual terminal lines (VTYs) is as follows:
  - Apply the basic ACL for inband access to all VTYs.
  - Configure a line password for all configured VTYs.
  - If the installed IOS image permits, use Secure Shell Protocol (SSH) to access the device remotely, instead of Telnet.
- **Configure system warning banners:** For both legal and administrative purposes, configuring a system warning banner to display prior to login is a convenient and effective way of reinforcing security and general usage policies. By clearly stating the ownership, usage, access, and protection policies prior to a login, future potential prosecution becomes more solidly backed.
- **Disable unneeded services:** By default, Cisco devices implement multiple TCP and UDP servers to facilitate management and integration into existing environments. For most installations these services are typically not required, and disabling them can greatly reduce overall security exposure. These commands will disable the services not typically used:

  ```
  no service tcp-small-servers
  no service udp-small-servers
  no service finger
  no service config
  ```
- **Secure Shell Protocol (SSH):** The protocol and application provide a secure, remote connection to a router. Two versions of SSH are available, SSH Version 1 and SSH Version 2. SSH Version 1 is implemented in IOS software. It encrypts all traffic, including passwords, between a remote console and a network router across a Telnet session. Because SSH sends no traffic in the clear, network administrators can conduct remote access sessions that casual observers will not be able to view. The SSH server in IOS software will work with publicly and commercially available SSH clients.
- **Cisco Discovery Protocol (CDP):** Security concerns regarding CDP remain controversial primarily due to a lack of understanding of its operation. CDP is a link-layer protocol that provides information regarding the properties of an adjacent Layer 2 device. CDP is often useful within a LAN as a troubleshooting and network management tool. Some Cisco management applications like CiscoWorks Campus Manager rely on CDP for accurate network topology discovery. CDP does not reveal security-specific information, but it is possible for an attacker to exploit this information in a reconnaissance attack, whereby an attacker gains knowledge of a network for the purpose of launching other types of attacks. A more practical method of implementing CDP is to follow a few very simple guidelines.
  - If CDP is not required, or the device is located in an unsecure environment, disable CDP globally on the device.
  - If it is required, disable CDP on a per-interface basis on ports connected to untrusted networks. Because CDP is a link-level protocol, it is not transient across a network (unless a Layer 2 tunneling mechanism is in place). Limit it to run only between trusted devices, disabling it everywhere else.

- **Disable the integrated HTTP daemon:** Although IOS software provides an integrated HTTP server for management, it is highly recommended to disable it to reduce overall exposure. If HTTP is absolutely required, use basic ACLs to isolate access from only trusted subnets.
- **Configure basic logging:** To assist and simplify both problem troubleshooting and security investigations, monitor switch subsystem information received from the logging facility. View the output in the on-system logging buffer memory. To render the on-system logging useful, increase the default buffer size.
- **Secure SNMP:** Whenever possible, avoid using SNMP read-write features. SNMP v2c authentication consists of simple text strings communicated between devices in clear, unencrypted text. In most cases, a read-only community string may be configured. In doing so, applying the basic access list to mask access to SNMP to trusted hosts only is an absolutely necessary precaution.
- **Limit trunking connections:** By default, Catalyst switches running IOS software are configured to automatically negotiate trunking capabilities. This situation poses a serious hazard to the infrastructure. It allows the possibility of an unsecured third party to be introduced into the infrastructure, as part of the infrastructure. Potential attacks include interception of traffic, redirection of traffic, denial of service (DoS), and more. To avoid this risk, disable automatic negotiation of trunking, and manually enable it on links that will require it.
- **Secure the spanning tree topology:** It is important to protect the Spanning Tree Protocol (STP) process of the switches composing the infrastructure. Inadvertent or malicious introduction of STP BPDUs could potentially overwhelm a device or pose a DoS attack. The first step in stabilizing a spanning tree installation is to positively identify the intended root bridge in the design, and to hard set the STP bridge priority of that bridge to an acceptable root value. Do the same for the designated backup root bridge. These actions will protect against inadvertent shifts in STP due to an uncontrolled introduction of a new switch.

  In addition to taking these steps, on some platforms the BPDU guard feature may be available. If this feature is available for your platform, enable it on access ports in conjunction with the PortFast feature to protect the network from unwanted BPDU traffic injection. Upon receipt of a BPDU, the feature will automatically disable the port.

Configuring AAA

Authentication, authorization, and accounting (AAA) network security services provide the primary framework through which you set up access control on a switch. This topic explains how to configure AAA on Catalyst switches.

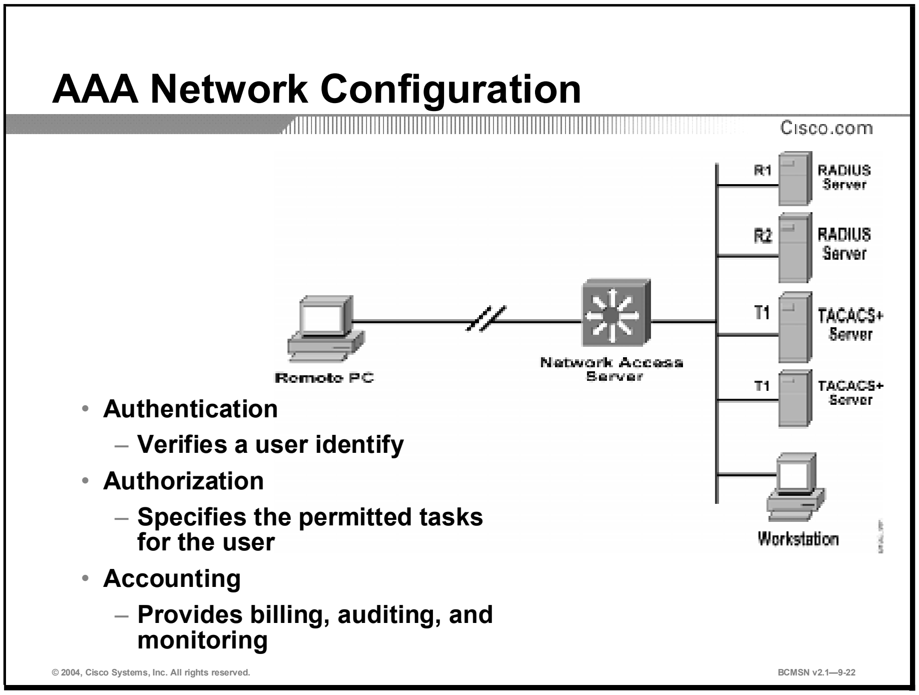

AAA Network Configuration

Cisco.com

- Authentication
  - Verifies a user identify
- Authorization
  - Specifies the permitted tasks for the user
- Accounting
  - Provides billing, auditing, and monitoring

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—9-22

AAA is an architectural framework for configuring a set of three independent security functions in a consistent manner. AAA provides a modular way of performing these services:

- **Authentication:** Provides the method of identifying users, including login and password dialog, challenge and response, messaging support, and, depending on the security protocol you select, encryption.

  Authentication is the way in which a user is identified prior to being allowed access to the network and network services. You configure AAA authentication by defining a named list of authentication methods and then applying that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they will be performed; it must be applied to a specific interface before any of the defined authentication methods will be performed. The only exception is the default method list (which is named “default”). The default method list is automatically applied to all interfaces if no other method list is defined. A defined method list overrides the default method list.

  All authentication methods, except for local, line password, and enable authentication, must be defined through AAA.

- **Authorization:** Provides the method for remote access control, including one-time authorization or authorization for each service, per-user account list and profile, user group support, and support of IP, Internetwork Packet Exchange (IPX), AppleTalk Remote Access (ARA), and Telnet.

AAA authorization works by assembling a set of attributes that describe what the user is authorized to perform, such as access to different parts of the network. These attributes are compared to the information contained in a database for a given user, and the result is returned to AAA to determine the actual capabilities and restrictions of the user. The database can be located locally on the multilayer switch, or it can be hosted remotely on a RADIUS or TACACS+ security server. Remote security servers, such as RADIUS and TACACS+, authorize users for specific rights by associating attribute-value pairs, which define those rights with the appropriate user. All authorization methods must be defined through AAA.

As with authentication, you configure AAA authorization by defining a named list of authorization methods, and then applying that list to various interfaces.

- **Accounting:** Provides the method for collecting and sending security server information used for billing, auditing, and reporting, such as user identities, start and stop times, executed commands (such as PPP), number of packets, and number of bytes. Security experts can use the information gained from accounting to audit and improve security.

In many circumstances, AAA uses protocols such as RADIUS, TACACS+, or 802.1x to administer its security functions. If your switch is acting as a network access server, AAA is the means through which a switch establishes communication between your network access server and your RADIUS, TACACS+, or 802.1x security server.

AAA is designed to enable you to dynamically configure the type of authentication and authorization that you want on a per-line (per-user) or per-service (for example, IP, IPX, or virtual private dial-up network [VPDN]) basis. You define the type of authentication and authorization that you want by creating method lists, and then apply those method lists to specific services or interfaces.

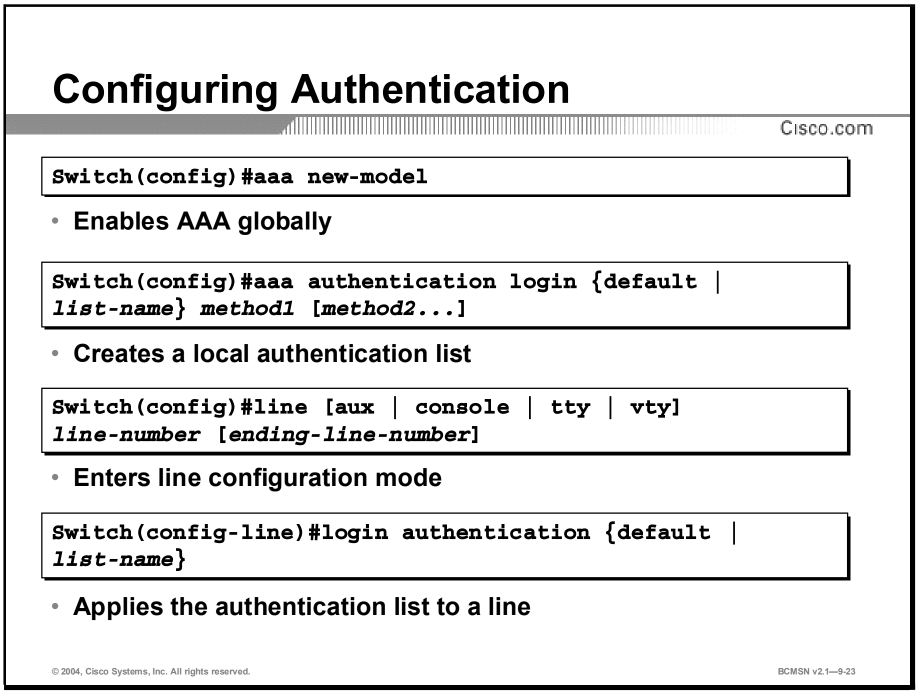

Configuring Authentication

Cisco.com

```
Switch(config)#aaa new-model
```

- Enables AAA globally

```
Switch(config)#aaa authentication login {default |
list-name} method1 [method2...]
```

- Creates a local authentication list

```
Switch(config)#line [aux | console | tty | vty]
line-number [ending-line-number]
```

- Enters line configuration mode

```
Switch(config-line)#login authentication {default |
list-name}
```

- Applies the authentication list to a line

BCMSN v2.1—9-23

The AAA security services facilitate a variety of login authentication methods. Use the **aaa authentication login** command to enable AAA authentication. With the **aaa authentication login** command, you create one or more lists of authentication methods that are tried at login. These lists are applied using the **login authentication** line configuration command.

To configure login authentication by using AAA, use the commands in this table, beginning in global configuration mode.

Step	Description
1.	Enable AAA globally. Switch(config)#**aaa new-model**
2.	Create a local authentication list. Switch(config)#**aaa authentication login** {**default** | *list-name*} *method1* [*method2...*]
3.	Enter line configuration mode for the lines to which you want to apply the authentication list. Switch(config)#**line** [**aux** | **console** | **tty** | **vty**] *line-number* [*ending-line-number*]
4.	Apply the authentication list to a line or set of lines. Switch(config-line)#**login authentication** {**default** | *list-name*}

Example: Configuring Authentication

The following example creates an authentication list called “myway” that uses TACACS+ as the first authentication method and local authentication as the second. The authentication list is then applied to a line.

```
Switch(config)#aaa authentication login myway tacacs+ local
Switch(config)#line con 0
Switch(config-line)#login authentication myway
```

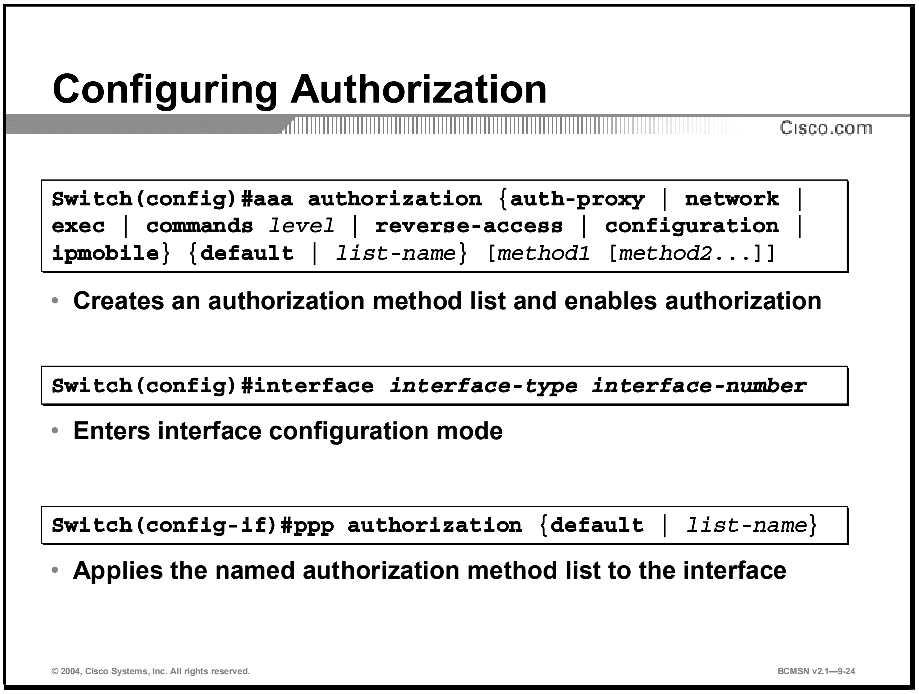

Configuring Authorization

Cisco.com

```
Switch(config)#aaa authorization {auth-proxy | network |
exec | commands level | reverse-access | configuration |
ipmobile} {default | list-name} [method1 [method2...]]
```

- **Creates an authorization method list and enables authorization**

```
Switch(config)#interface interface-type interface-number
```

- **Enters interface configuration mode**

```
Switch(config-if)#ppp authorization {default | list-name}
```

- **Applies the named authorization method list to the interface**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—9-24

AAA authorization enables you to limit the services available to a user. When AAA authorization is enabled, the multilayer switch uses information retrieved from the user profile, which is located either in the local user database on the switch or on the security server, to configure the user session. When this task is done, the user will be granted access to a requested service only if the information in the user profile allows it.

Just as with AAA authentication, you create method lists to define the ways that authorization will be performed and the sequence in which these methods will be performed. Method lists are specific to the authorization type requested:

- **Auth-proxy:** Applies specific security policies on a per-user basis.
- **Commands:** Applies to the EXEC mode commands that a user issues. Command authorization attempts authorization for all EXEC mode commands, including global configuration commands, associated with a specific privilege level.
- **EXEC:** Applies to the attributes associated with a user EXEC terminal session.
- **Network:** Applies to network connections. These connections can include a PPP, Serial Line Internet Protocol (SLIP), or AppleTalk Remote Access Protocol (ARAP) connection.
- **Reverse access:** Applies to reverse Telnet sessions.

When you create a named method list, you are defining a particular list of authorization methods for the indicated authorization type.

AAA supports five different methods of authorization:

- **TACACS+:** The network access server exchanges authorization information with the TACACS+ security daemon. TACACS+ authorization defines specific rights for users by associating attribute-value pairs, which are stored in a database on the TACACS+ security server, with the appropriate user.
- **If-Authenticated:** The user is allowed to access the requested function, provided that the user has been authenticated successfully.
- **None:** The network access server does not request authorization information; authorization is not performed over this line or interface.
- **Local:** The router or access server consults its local database, as defined by the **username** command, for example, to authorize specific rights for users. Only a limited set of functions can be controlled via the local database.
- **RADIUS:** The network access server requests authorization information from a RADIUS security server. RADIUS authorization defines specific rights for users by associating attributes.

To configure AAA authorization using named method lists, use these commands, beginning in global configuration mode:

```
Switch(config)#aaa authorization {auth-proxy | network | exec
| commands level | reverse-access | configuration | ipmobile}
{default | list-name} [method1 [method2...]]
Switch(config)#line [aux | console | tty | vty] line-number
[ending-line-number]
Switch(config-line)#authorization {arap | commands level |
exec | reverse-access} {default | list-name}
OR
Switch(config)# interface interface-type interface-number
Switch(config-if)#ppp authorization {default | list-name}
```

To have the multilayer switch request authorization information via a TACACS+ security server, use the **aaa authorization** command with the **group tacacs+** value for the *method* variable.

To allow users to have access to the functions that they request as long as they have been authenticated, use the **aaa authorization** command with the **if-authenticated** method keyword. If you select this method, all requested functions are automatically granted to authenticated users.

To select local authorization, which means that the router or access server consults its local user database to determine the functions that a user is permitted to use, use the **aaa authorization** command with the **local** method keyword. The functions associated with local authorization are defined by using the **username** global configuration command.

To have the network access server request authorization via a RADIUS security server, use the **radius** method keyword.

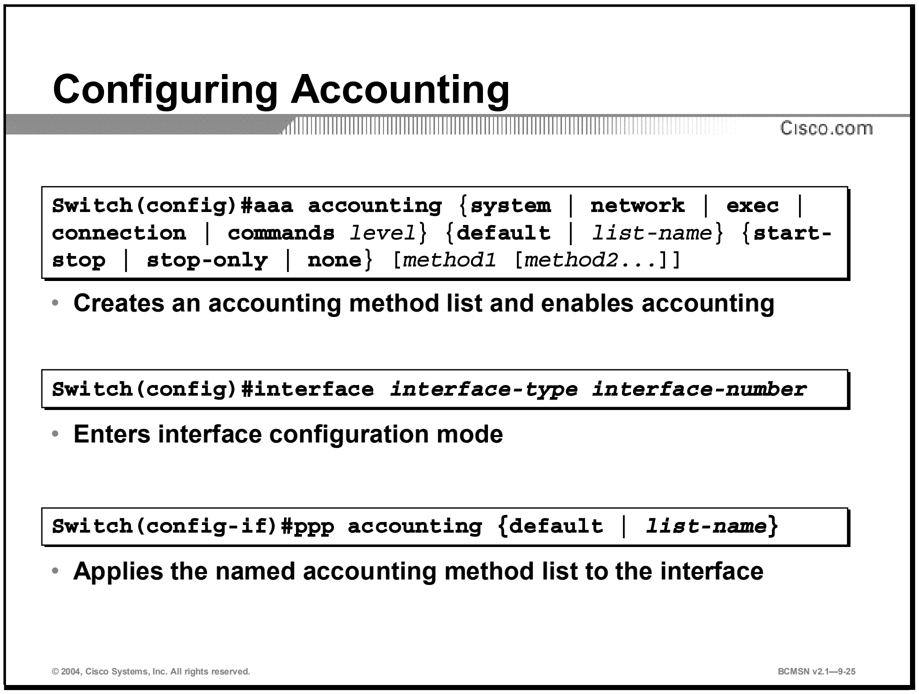

Configuring Accounting

Cisco.com

```
Switch(config)#aaa accounting {system | network | exec |
connection | commands level} {default | list-name} {start-
stop | stop-only | none} [method1 [method2...]]
```

- **Creates an accounting method list and enables accounting**

```
Switch(config)#interface interface-type interface-number
```

- **Enters interface configuration mode**

```
Switch(config-if)#ppp accounting {default | list-name}
```

- **Applies the named accounting method list to the interface**

BCMSN v2.1—9-25

AAA supports six different accounting types:

- **Network accounting:** Provides information for all PPP, SLIP, or ARAP sessions, including packet and byte counts
- **Connection accounting:** Provides information about all outbound connections made from the network, such as Telnet and remote login (rlogin)
- **EXEC accounting:** Provides information about user EXEC terminal sessions (user shells) on the network access server, including username, date, start and stop times, the access server IP address, and (for dial-in users) the telephone number the call originated from
- **System accounting:** Provides information about all system-level events (for example, when the system reboots or when accounting is turned on or off)
- **Command accounting:** Provides information about the EXEC shell commands for a specified privilege level that are being executed on a network access server
- **Resource accounting:** Provides start and stop record support for calls that have passed user authentication

To configure AAA accounting using named method lists, use the commands in this table, beginning in global configuration mode.

Step	Description
1.	Create an accounting method list and enable accounting. `Switch(config)#`**`aaa accounting`** `{`**`system`** `|` **`network`** `|` **`exec`** `|` **`connection`** `|` **`commands`** *`level`*`}` `{`**`default`** `|` *`list-name`*`}` `{`**`start-stop`** `|` **`stop-only`** `|` **`none`**`}` `[`*`method1`* `[`*`method2...`*`]]`

2.	Enter the line configuration mode or the interface to which you want to apply the accounting method list.
	`Switch(config)#`**`line`** `[`**`aux`** `|` **`console`** `|` **`tty`** `|` **`vty`**`]` *`line-number`* `[`*`ending-line-number`*`]`
	or
	`Switch(config)#`**`interface`** *`interface-type interface-number`*

3.	Apply the accounting method list to a line or interface.
	`Switch(config-line)#`**`accounting`** `{`**`arap`** `|` **`commands level`** `|` **`connection`** `|` **`exec`**`}` `{`**`default`** `|` *`list-name`*`}`
	or
	`Switch(config-if)#`**`ppp accounting`** `{`**`default`** `|` *`list-name`*`}`

Configuring Network Access Security

Network access security is provided by port security and port-based authentication (802.1x). This topic explains how to configure port security and port-based authentication with 802.1x.

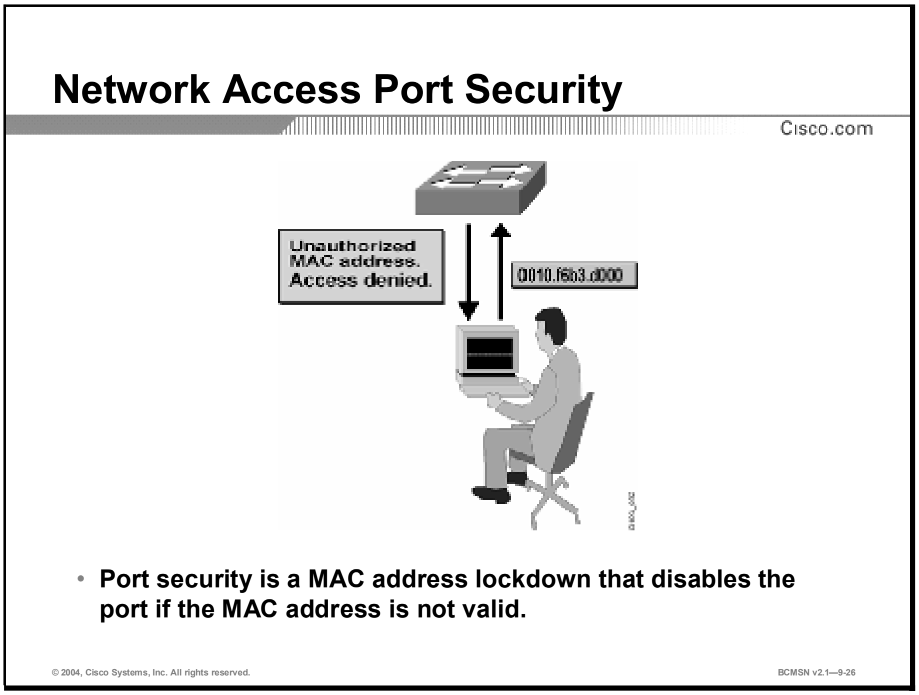

Network Access Port Security

Cisco.com

Unauthorized MAC address. Access denied.
0010.f6b3.d000

- **Port security is a MAC address lockdown that disables the port if the MAC address is not valid.**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—9-26

Port security is a feature of the Cisco Catalyst switches that allows the switch to block input from a port when the MAC address of a station attempting to access the port is different from the configured MAC address. This blocking feature is referred to as a MAC address lockdown.

When a port receives a frame, the port compares the source address of the frame to the secure source address that was originally learned by the port. If the addresses do not match, the port is disabled and the LED for the port turns amber.

Note Port security cannot be applied to trunk ports where addresses might change frequently. Not all hardware supports port security. Check your documentation or http://www.cisco.com to see if your hardware supports this feature.

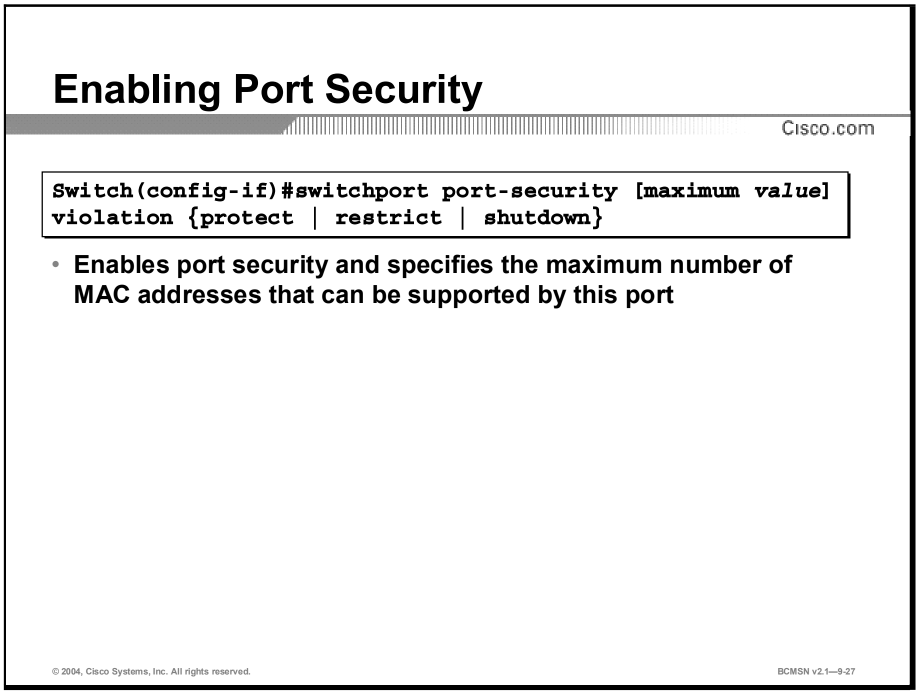

Enabling Port Security

Cisco.com

```
Switch(config-if)#switchport port-security [maximum value]
violation {protect | restrict | shutdown}
```

- **Enables port security and specifies the maximum number of MAC addresses that can be supported by this port**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—9-27

By default, the switch allows all MAC addresses to access the network. It relies on other types of security, such as file server operating systems and applications, to provide for network security. Port security allows a network administrator to configure a set of allowed devices or MAC addresses to provide additional security. If port security is enabled, only the MAC addresses that are explicitly allowed can use the port. A MAC address can be allowed in two ways:

- **Static assignment of the MAC address:** The network administrator can code the MAC address when port security is assigned. This is the more secure of the two methods, but is difficult to manage.
- **Dynamic learning of the MAC address**: If the MAC address is not specified, the port can learn the secure MAC address. The first MAC address seen on the port becomes the secure MAC address.

Use this command to enable port security:

```
Switch(config-if)#switchport port-security [maximum value]
violation {protect | restrict | shutdown}
```

The **maximum** *value* option allows the network administrator to define the maximum number of MAC addresses that can be supported by this port. The maximum number can range from 1 to 132. The default value is 132. The **violation** option specifies what action to take when a violation occurs:

- **protect:** When the number of secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value.
- **restrict:** A port security violation restricts data and causes the SecurityViolation counter to increment.
- **shutdown:** The interface is error-disabled when a security violation occurs.

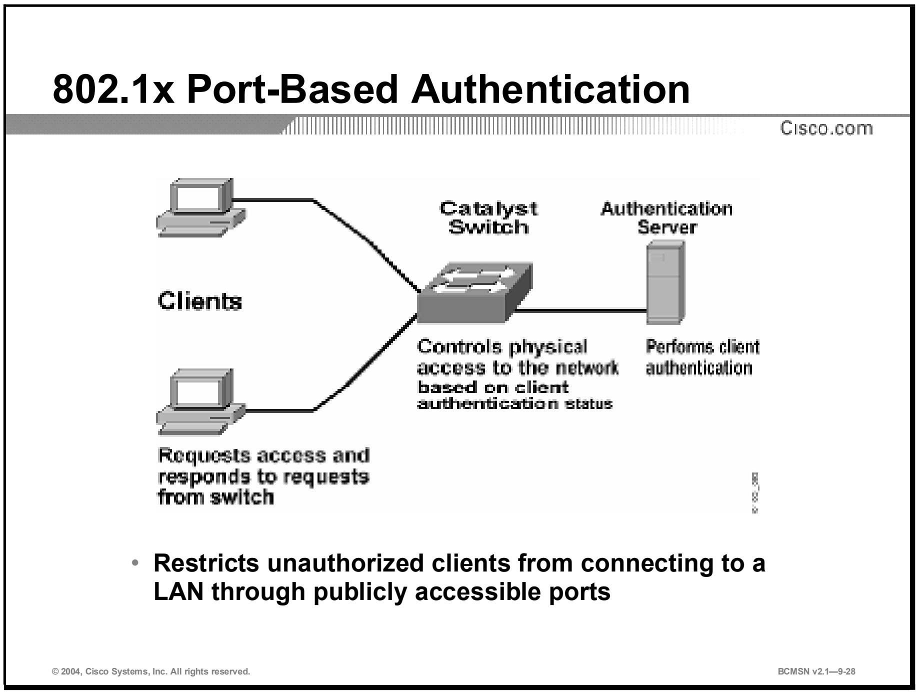

802.1x Port-Based Authentication

Cisco.com

Catalyst Switch

Authentication Server

Clients

Controls physical access to the network based on client authentication status

Performs client authentication

Requests access and responds to requests from switch

- **Restricts unauthorized clients from connecting to a LAN through publicly accessible ports**

BCMSN v2.1—9-28

The IEEE 802.1x standard defines a port-based access control and authentication protocol that restricts unauthorized workstations from connecting to a LAN through publicly accessible ports. The authentication server authenticates each workstation connected to a switch port before making available any services offered by the switch or the LAN.

Until the workstation is authenticated, 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the workstation is connected. After authentication is successful, normal traffic can pass through the port.

With 802.1x port-based authentication, the devices in the network have specific roles as follows:

- **Client:** The device (workstation) that requests access to the LAN and switch services and responds to requests from the switch. The workstation must be running 802.1x-compliant client software such as that offered in the Microsoft Windows XP operating system. (The port that the client is attached to is the supplicant [client] in the IEEE 802.1x specification.)
- **Authentication server:** Performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. The RADIUS security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server.
- **Switch (also called the authenticator):** Controls the physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client (supplicant) and the authentication server, requesting identity information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch uses a RADIUS software agent, which is responsible for encapsulating and decapsulating the EAP frames and interacting with the authentication server.

The switch port state determines whether or not the client is granted access to the network. The port starts in the unauthorized state. While in this state, the port disallows all ingress and egress traffic except for 802.1x protocol packets. When a client is successfully authenticated, the port transitions to the authorized state, allowing all traffic for the client to flow normally.

If the switch requests the client identity (authenticator initiation) and the client does not support 802.1x, the port remains in the unauthorized state and the client is not granted access to the network.

In contrast, when an 802.1x-enabled client connects to a port and the client initiates the authentication process (supplicant initiation) by sending the EAPOL-start frame to a switch not running the 802.1x protocol, no response is received, and the client begins sending frames as if the port is in the authorized state.

You control the port authorization state by using the **dot1x port-control** interface configuration command and these keywords:

- **force-authorized:** Disables 802.1x port-based authentication and causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1x-based authentication of the client. This is the default setting.
- **force-unauthorized:** Causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
- **auto:** Enables 802.1x port-based authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port transitions from down to up (authenticator initiation) or when an EAPOL-start frame is received (supplicant initiation). The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. The switch uniquely identifies each client attempting to access the network by using the client MAC address.

If the client is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails, and network access is not granted.

When a client logs off, it sends an EAPOL-logoff message, causing the switch port to transition to the unauthorized state.

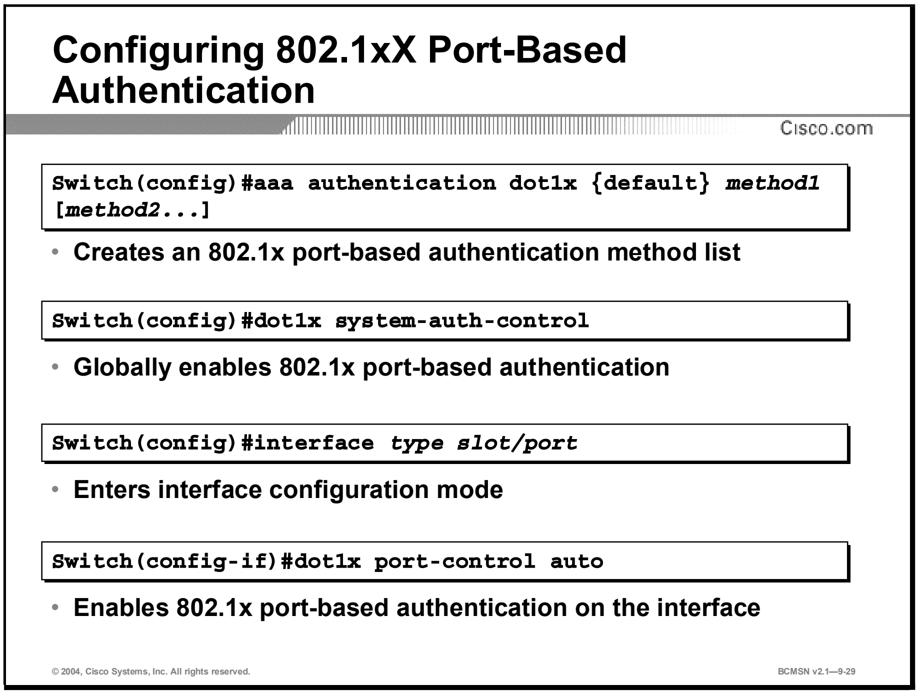

Configuring 802.1xX Port-Based Authentication

Cisco.com

```
Switch(config)#aaa authentication dot1x {default} method1
[method2...]
```

- **Creates an 802.1x port-based authentication method list**

```
Switch(config)#dot1x system-auth-control
```

- **Globally enables 802.1x port-based authentication**

```
Switch(config)#interface type slot/port
```

- **Enters interface configuration mode**

```
Switch(config-if)#dot1x port-control auto
```

- **Enables 802.1x port-based authentication on the interface**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—9-29

To configure 802.1x port-based authentication, perform the tasks presented in this table.

Step	Description
1.	Enable AAA. `Switch(config)#aaa new-model`
2.	Create an 802.1x port-based authentication method list. `Switch(config)#aaa authentication dot1x {default} method1 [method2...]`
3.	Globally enable 802.1x port-based authentication. `Switch(config)#dot1x system-auth-control`
4.	Enter interface configuration mode and specify the interface to be enabled for 802.1x port-based authentication. `Switch(config)#interface type slot/port`
5.	Enable 802.1x port-based authentication on the interface. `Switch(config-if)#dot1x port-control auto`
6.	Return to privileged EXEC mode. `Switch(config)#end`

Example

The example shows how to enable AAA and 802.1x on Fast Ethernet port 5/1:

```
Switch#configure terminal
Switch(config)#aaa new-model
Switch(config)#aaa authentication dot1x default group radius
Switch(config)#dot1x system-auth-control
Switch(config)#interface fastethernet 5/1
Switch(config-if)#dot1x port-control auto
Switch(config-if)#end
```

Verifying Network Access Security

Use **show** commands to verify the configuration of port security. This topic explains how to verify port security.

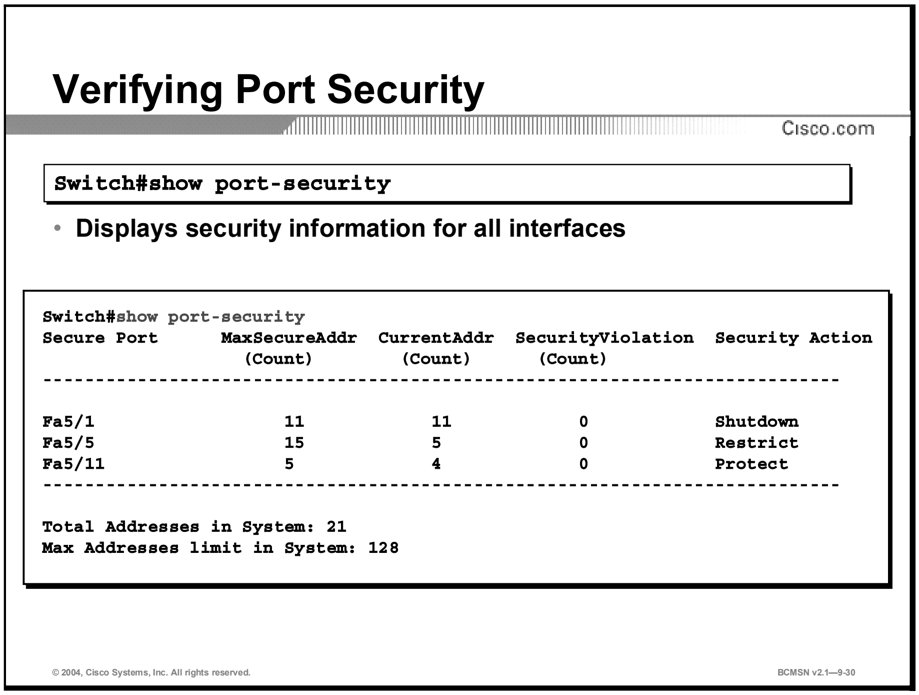

Verifying Port Security

```
Switch#show port-security
```

- **Displays security information for all interfaces**

```
Switch#show port-security
Secure Port      MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
                   (Count)        (Count)       (Count)
---------------------------------------------------------------------------

Fa5/1                 11            11            0            Shutdown
Fa5/5                 15            5             0            Restrict
Fa5/11                5             4             0            Protect
---------------------------------------------------------------------------

Total Addresses in System: 21
Max Addresses limit in System: 128
```

BCMSN v2.1—9-30

Verify port security with the **show port-security** command.

Example: show port-security Command Output

This example displays output from the **show port-security** command when you do not enter an interface:

```
Switch#show port-security
Secure Port      MaxSecureAddr  CurrentAddr  SecurityViolation
Security Action
                   (Count)        (Count)       (Count)
--------------------------------------------------------------
-------

Fa5/1                11             11            0
Shutdown
Fa5/5                15             5             0
Restrict
Fa5/11               5              4             0
Protect
--------------------------------------------------------------
-------

Total Addresses in System: 21
Max Addresses limit in System: 128
```

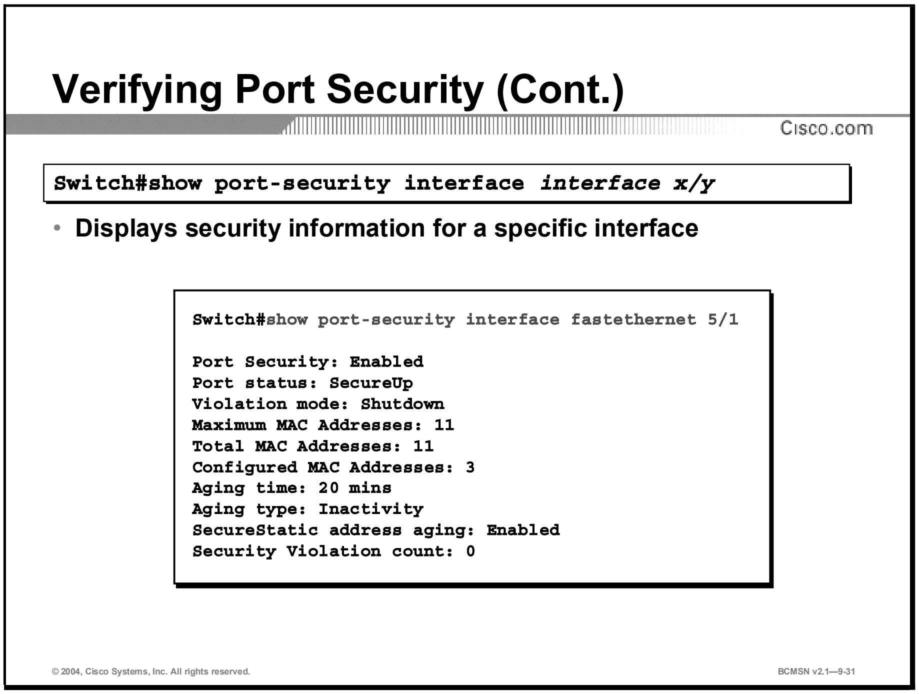

Verifying Port Security (Cont.)

Cisco.com

```
Switch#show port-security interface interface x/y
```

- Displays security information for a specific interface

```
Switch#show port-security interface fastethernet 5/1

Port Security: Enabled
Port status: SecureUp
Violation mode: Shutdown
Maximum MAC Addresses: 11
Total MAC Addresses: 11
Configured MAC Addresses: 3
Aging time: 20 mins
Aging type: Inactivity
SecureStatic address aging: Enabled
Security Violation count: 0
```

BCMSN v2.1—9-31

Use the **interface** argument to restrict the output to a specific interface.

Example: show port-security Command for a Specific Interface

This example displays output from the **show port-security** command for a specified interface:

```
Switch#show port-security interface fastethernet 5/1

Port Security: Enabled
Port status: SecureUp
Violation mode: Shutdown
Maximum MAC Addresses: 11
Total MAC Addresses: 11
Configured MAC Addresses: 3
Aging time: 20 mins
Aging type: Inactivity
SecureStatic address aging: Enabled
Security Violation count: 0
```

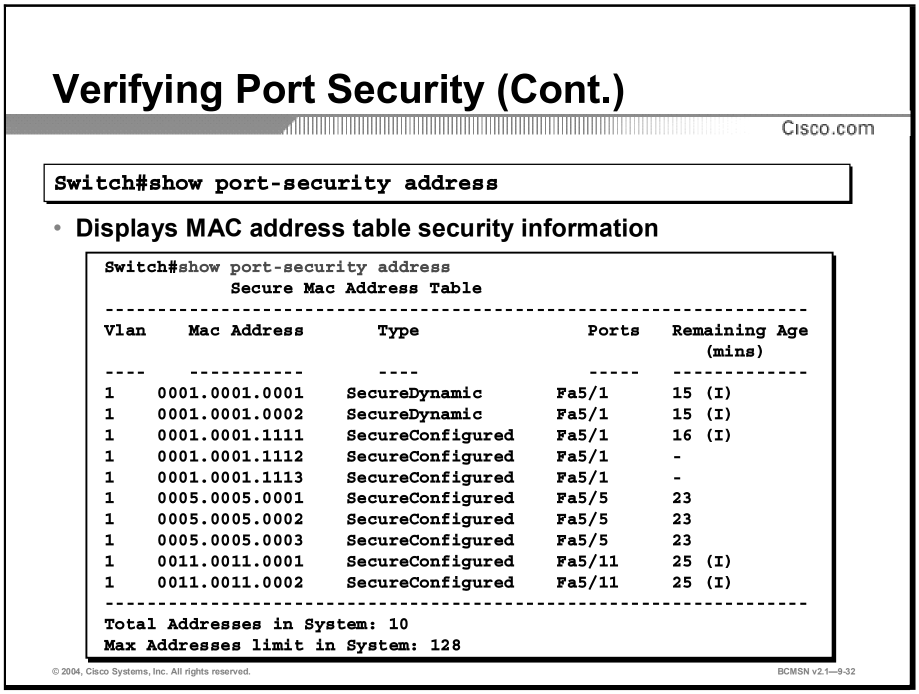

Verifying Port Security (Cont.)

Cisco.com

```
Switch#show port-security address
```

- **Displays MAC address table security information**

```
Switch#show port-security address
          Secure Mac Address Table
-------------------------------------------------------------------
Vlan    Mac Address       Type              Ports   Remaining Age
                                                       (mins)
----    -----------       ----              -----   -------------
1    0001.0001.0001    SecureDynamic       Fa5/1      15 (I)
1    0001.0001.0002    SecureDynamic       Fa5/1      15 (I)
1    0001.0001.1111    SecureConfigured    Fa5/1      16 (I)
1    0001.0001.1112    SecureConfigured    Fa5/1      -
1    0001.0001.1113    SecureConfigured    Fa5/1      -
1    0005.0005.0001    SecureConfigured    Fa5/5      23
1    0005.0005.0002    SecureConfigured    Fa5/5      23
1    0005.0005.0003    SecureConfigured    Fa5/5      23
1    0011.0011.0001    SecureConfigured    Fa5/11     25 (I)
1    0011.0011.0002    SecureConfigured    Fa5/11     25 (I)
-------------------------------------------------------------------
Total Addresses in System: 10
Max Addresses limit in System: 128
```

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—9-32

Use the **address** argument to display MAC address table security information.

Example: Displaying MAC Address Table Security Information

This example displays output from the **show port-security address** privileged EXEC command:

```
Switch#show port-security address
            Secure Mac Address Table
--------------------------------------------------------------
-----
Vlan    Mac Address       Type                Ports
Remaining Age

(mins)
----    -----------       ----                -----    --------
-----
1    0001.0001.0001    SecureDynamic       Fa5/1      15 (I)
1    0001.0001.0002    SecureDynamic       Fa5/1      15 (I)
1    0001.0001.1111    SecureConfigured    Fa5/1      16 (I)
1    0001.0001.1112    SecureConfigured    Fa5/1      -
1    0001.0001.1113    SecureConfigured    Fa5/1      -
1    0005.0005.0001    SecureConfigured    Fa5/5      23
1    0005.0005.0002    SecureConfigured    Fa5/5      23
1    0005.0005.0003    SecureConfigured    Fa5/5      23
1    0011.0011.0001    SecureConfigured    Fa5/11     25 (I)
1    0011.0011.0002    SecureConfigured    Fa5/11     25 (I)
--------------------------------------------------------------
-----
Total Addresses in System: 10
Max Addresses limit in System: 128
```

Configuring Security Using Access Lists

ACLs are useful for controlling access in a multilayer switched network. This topic explains how to configure security with ACLs.

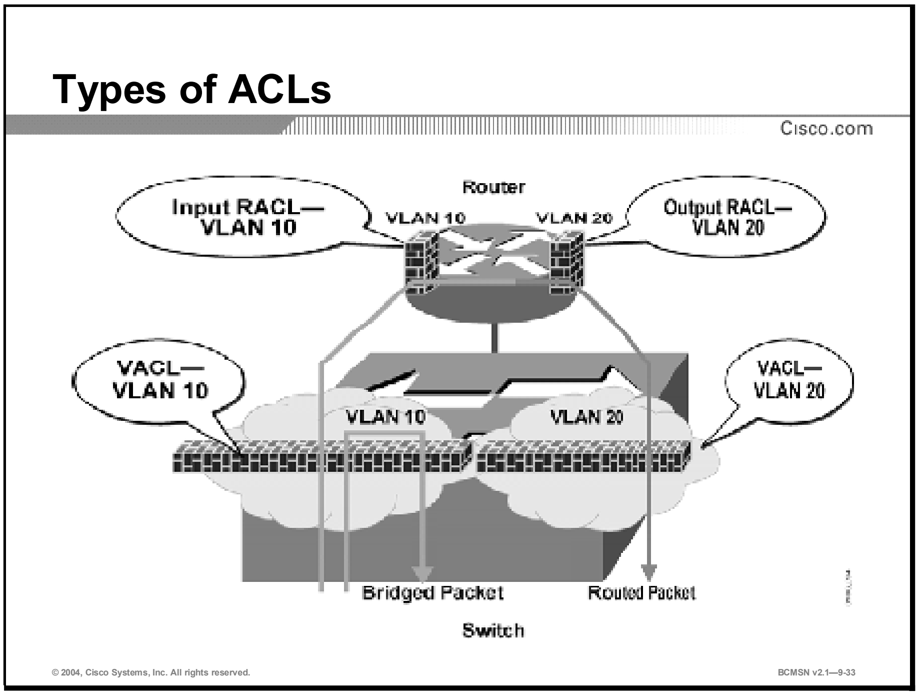

Types of ACLs

Cisco.com

BCMSN v2.1—9-33

Cisco multilayer switches support three types of ACLs:

- **Router access control lists (RACLs):** Supported in the ternary content addressable memory (TCAM) hardware on Cisco multilayer switches
- **Quality of service (QoS) access control lists:** Supported in the TCAM hardware on Cisco multilayer switches
- **VLAN access control lists (VACLs):** Supported in software on Cisco multilayer switches

Catalyst switches support four ACL lookups per packet: input and output security ACL and input and output QoS ACL.

Catalyst switches use two methods of performing a merge: order independent and order dependent. With order independent merge, ACLs are transformed from a series of order dependent actions to a set of order independent masks and patterns. The resulting access control entry can be very large. The merge is processor- and memory-intensive.

Order dependent merge is a recent improvement on some Catalyst switches in which ACLs retain their order dependent aspect. The computation is much faster and is less processor-intensive.

RACLs are supported in hardware through IP standard ACLs and IP extended ACLs, with permit and deny actions. ACL processing is an intrinsic part of the packet-forwarding process. ACL entries are programmed in hardware. Lookups occur in the pipeline whether ACLs are configured or not. With RACLs, access list statistics and logging are not supported.

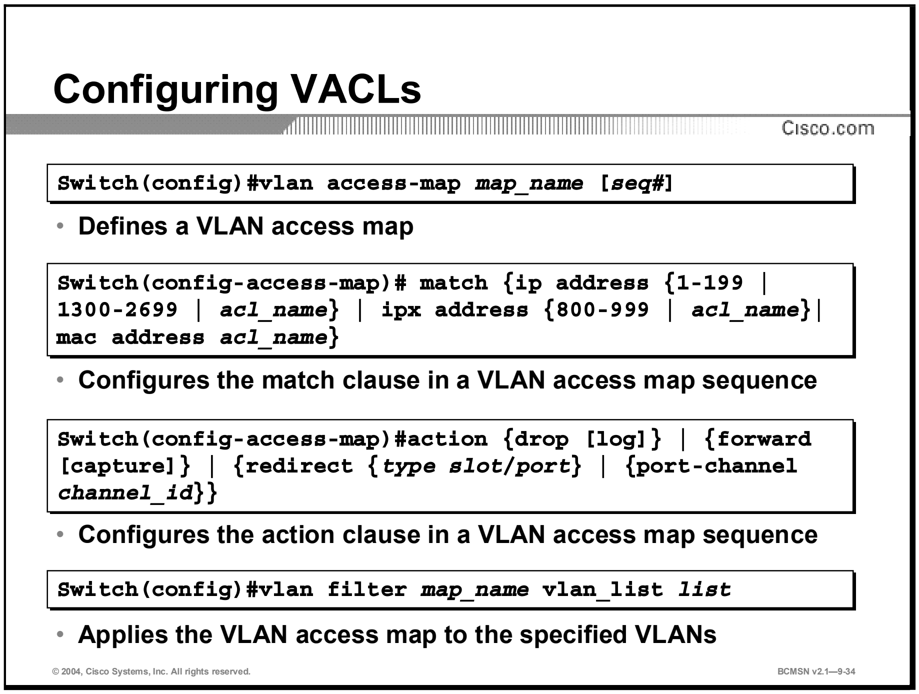

Configuring VACLs

Cisco.com

```
Switch(config)#vlan access-map map_name [seq#]
```

- Defines a VLAN access map

```
Switch(config-access-map)# match {ip address {1-199 |
1300-2699 | acl_name} | ipx address {800-999 | acl_name}|
mac address acl_name}
```

- Configures the match clause in a VLAN access map sequence

```
Switch(config-access-map)#action {drop [log]} | {forward
[capture]} | {redirect {type slot/port} | {port-channel
channel_id}}
```

- Configures the action clause in a VLAN access map sequence

```
Switch(config)#vlan filter map_name vlan_list list
```

- Applies the VLAN access map to the specified VLANs

BCMSN v2.1—9-34

VACLs (also called VLAN access maps in IOS software) apply to all traffic on the VLAN. They filter based on Ethertype and MAC address traffic.

VACLs follow route-map conventions, where map sequences are checked in order.

When a matching permit access control entry (ACE) is encountered, the switch takes the action. When a matching deny ACE is encountered, the switch checks the next ACL in the sequence or checks the next sequence.

Three VACL actions are permitted:

- **Permit** (with capture, Catalyst 6500 only)
- **Redirect** (Catalyst 6500 only)
- **Deny** (with logging, Catalyst 6500 only)

The VACL capture option copies traffic to specified capture ports. VACL ACEs installed in hardware are merged with RACLs and other features.

Two features are supported only on the Catalyst 6500:

- **VACL capture:** Forwarded packets are captured on capture ports. The capture option is only on permit ACEs. The capture port can be an IDS monitor port or any Ethernet port. The capture port must be in an output VLAN for Layer 3-switched traffic.
- **VACL redirect:** Matching packets are redirected to specified ports. You can configure up to five redirect ports. Redirect ports must be in a VLAN where VACL is applied.

To configure VACLs, complete these steps:

<table>
<tr><th>Step</th><th>Description</th></tr>
<tr><td>1.</td><td>Define a VLAN access map.
Switch(config)#<b>vlan access-map</b> <i>map_name</i> [<i>seq#</i>]</td></tr>
<tr><td>2.</td><td>Configure a match clause.
Switch(config-access-map)# <b>match</b> {<b>ip address</b> {<b>1-199</b> | <b>1300-2699</b> | <i>acl_name</i>} | <b>ipx address</b> {<b>800-999</b> | <i>acl_name</i>}| <b>mac address</b> <i>acl_name</i>}</td></tr>
<tr><td>3.</td><td>Configure an action clause.
Switch(config-access-map)#action {drop [log]} | {forward [capture]} | {redirect {{<b>fastethernet</b> | <b>gigabitethernet</b> | <b>tengigabitethernet</b>} <i>slot/port</i>} | {<b>port-channel</b> <i>channel_id</i>}}</td></tr>
<tr><td>4.</td><td>Apply a map to VLANs.
Switch(config)#<b>vlan filter</b> <i>map_name</i> <b>vlan_list</b> <i>list</i></td></tr>
<tr><td>5.</td><td>Verify the VACL configuration.
Switch#<b>show vlan access-map map_name</b>
Switch#<b>show vlan filter</b> [<b>access-map</b> <i>map_name</i> | <i>vlan_id</i>]</td></tr>
</table>

Configuring PVLANs

PVLANs provide Layer 2 isolation between ports within the same PVLAN. This topic explains how to configure PVLANs.

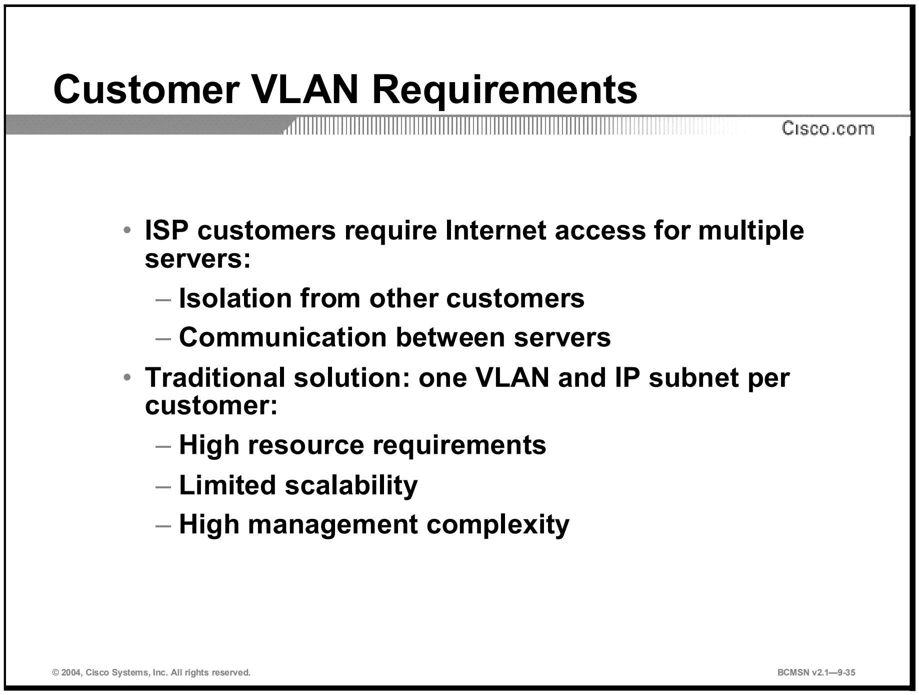

Customer VLAN Requirements

- **ISP customers require Internet access for multiple servers:**
  - **Isolation from other customers**
  - **Communication between servers**
- **Traditional solution: one VLAN and IP subnet per customer:**
  - **High resource requirements**
  - **Limited scalability**
  - **High management complexity**

BCMSN v2.1—9-35

Service provider customers usually want to connect multiple servers to the Internet, isolating their own traffic from other customer traffic while maintaining communication between their own servers.

The traditional solution to this requirement is for the Internet service provider (ISP) to provide one VLAN per customer, with each VLAN having its own IP subnet. A Layer 3 device aggregates the subnets to provide interconnectivity between VLANs and to route traffic to the Internet.

Problems with the traditional solution include the following:

- Supporting a separate VLAN per customer requires a high number of interfaces on the service provider network devices.
- The solution does not scale well, because spanning tree becomes more complicated with the addition of multiple VLANs.
- Maintaining multiple VLANs means maintaining multiple ACLs, increasing network management complexity.

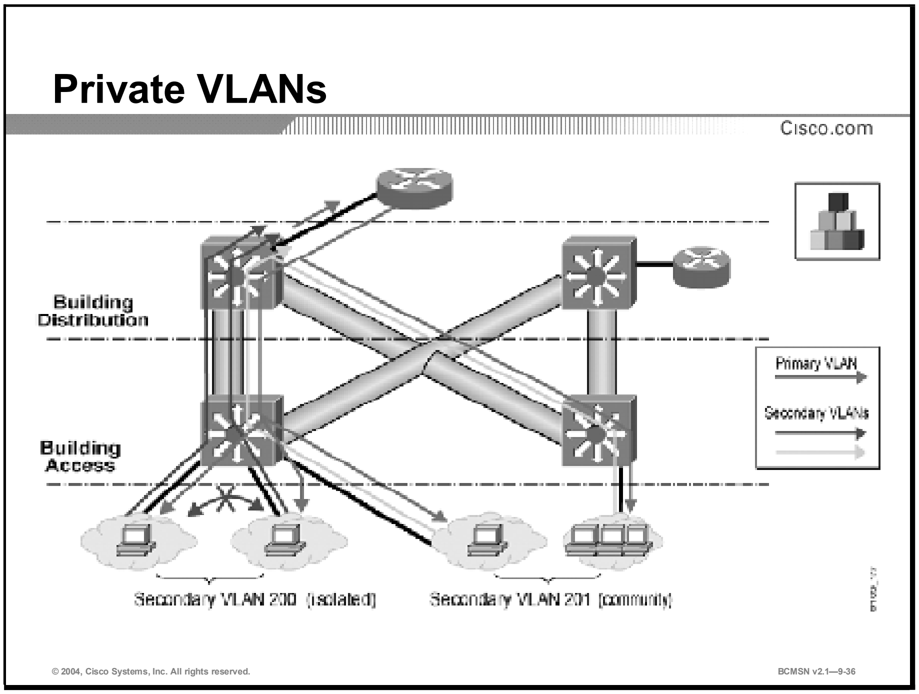

Private VLANs

Cisco.com

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—9-36

PVLANs provide Layer 2 isolation between ports within the same PVLAN. This isolation eliminates the need for a separate VLAN per customer, as well as the requirement of a separate IP subnet per customer. A range of addresses from a single IP network can be assigned to each customer.

The PVLAN solution provides these advantages:

- The number of VLANs is reduced, because multiple customers can share a single VLAN.
- Only one IP address is needed to route traffic from all customers.
- Only one IP subnet is needed, because it encompasses the entire PVLAN.
- Traffic is carried only over the primary VLAN, meaning that just the interface IP address of the primary VLAN needs to be advertised.

Note PVLANs are fully implemented on the Catalyst 6500 or equivalent.

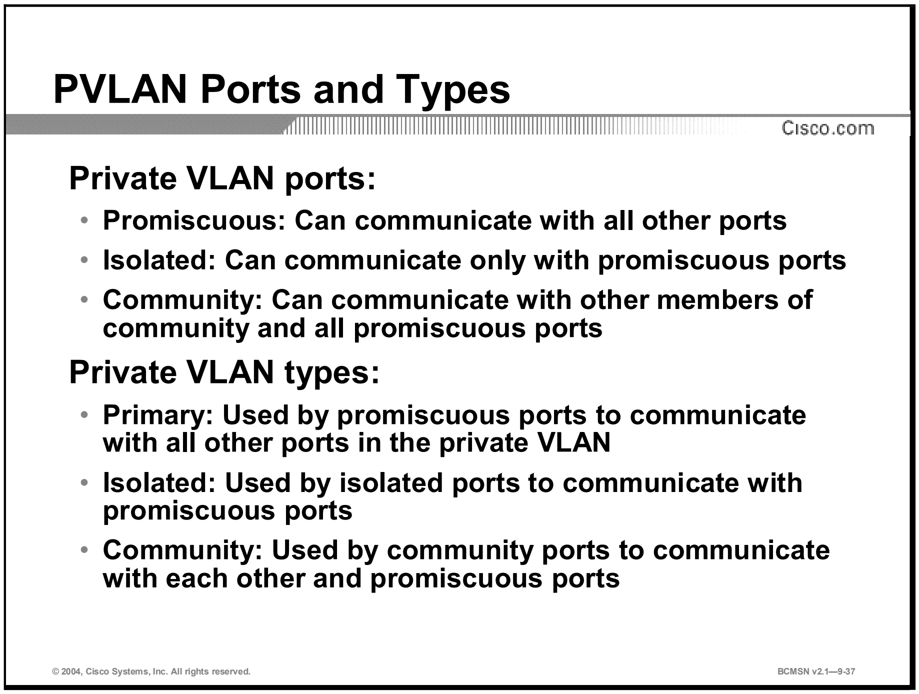

PVLAN Ports and Types

Cisco.com

Private VLAN ports:

- **Promiscuous: Can communicate with all other ports**
- **Isolated: Can communicate only with promiscuous ports**
- **Community: Can communicate with other members of community and all promiscuous ports**

Private VLAN types:

- **Primary: Used by promiscuous ports to communicate with all other ports in the private VLAN**
- **Isolated: Used by isolated ports to communicate with promiscuous ports**
- **Community: Used by community ports to communicate with each other and promiscuous ports**

BCMSN v2.1—9-37

There are three types of PVLAN ports:

- **Promiscuous:** A promiscuous port can communicate with all interfaces, including the community and isolated ports within a PVLAN.
- **Isolated:** An isolated port has complete Layer 2 separation from other ports within the same PVLAN except for the promiscuous port. PVLANs block all traffic to isolated ports except traffic from promiscuous ports. Traffic received from an isolated port is forwarded only to promiscuous ports.
- **Community:** Community ports communicate among themselves and with their promiscuous ports. These interfaces are isolated at Layer 2 from all other interfaces in other communities or in isolated ports within their PVLAN.

Note Because trunks can support the VLANs carrying traffic between isolated, community, and promiscuous ports, isolated and community port traffic might enter or leave the switch through a trunk interface.

PVLAN ports are associated with a set of supporting VLANs that are used to create the PVLAN structure. A PVLAN uses VLANs in three ways:

- **As a primary VLAN:** Carries traffic from promiscuous ports to isolated, community, and other promiscuous ports in the same primary VLAN.
- **As an isolated VLAN:** Carries traffic from isolated ports to a promiscuous port.
- **As a community VLAN:** Carries traffic between community ports and to promiscuous ports. You can configure multiple community VLANs in a PVLAN.

Isolated and community VLANs are called secondary VLANs. You can extend PVLANs across multiple devices by trunking the primary, isolated, and community VLANs to other devices that support PVLANs.

In a switched environment, you can assign an individual PVLAN and associated IP subnet to each individual or common group of end stations. The end stations need to communicate with a default gateway only to gain access outside the PVLAN. With end stations in a PVLAN, you can do the following:

- Designate selected ports connected to end stations. For example, you can designate interfaces connected to servers as isolated to prevent any communication at Layer 2. Or, if the end stations are servers, this configuration prevents Layer 2 communication between the servers.
- Designate the interfaces to which the default gateway(s) and selected end stations are attached as promiscuous ports to allow access to all end stations. Selected end stations can be backup servers or server load balancers, or a network appliance that load-balances TCP/IP traffic across multiple servers.
- Reduce VLAN and IP subnet consumption, because you can prevent traffic between end stations even though they are in the same VLAN and IP subnet.

Note A promiscuous port can service only one primary VLAN. A promiscuous port can service one isolated or many community VLANs.

With a promiscuous port, you can connect a wide range of devices as access points to a PVLAN. For example, you can connect a promiscuous port to the server port to connect an isolated VLAN or a number of community VLANs to the server. A load balancer can load-balance the servers present in the isolated or community VLANs, or you can use a promiscuous port to monitor or back up all the PVLAN servers from an administration workstation.

Follow these guidelines to configure PVLANs:

- Set VTP to transparent mode. After you configure a PVLAN, you cannot change the VTP mode to client or server.
- Do not include VLAN1 or VLANs 1002-1005 in the PVLAN configuration.
- Use only the PVLAN configuration commands to assign ports to primary, isolated, or community VLANs. Layer 2 interfaces assigned to the VLANs that you configure as primary, isolated, or community VLANs are inactive while the VLAN is part of the PVLAN configuration. Layer 2 trunk interfaces remain in the STP forwarding state.
- Configure Layer 3 VLAN interfaces only for primary VLANs. Layer 3 VLAN interfaces for isolated and community VLANs are inactive while the VLAN is configured as an isolated or community VLAN.
- Do not configure PVLAN ports as EtherChannel. While a port is part of the PVLAN configuration, any EtherChannel configuration for it is inactive.
- Do not configure a destination SPAN port as a PVLAN port. While a port is part of the PVLAN configuration, any destination SPAN configuration for it is inactive.
- You can configure a PVLAN port as a SPAN source port.
- You can use VSPAN on primary, isolated, and community VLANs, or use SPAN on only one VLAN to separately monitor egress or ingress traffic.
- A primary VLAN can have one isolated VLAN and multiple community VLANs associated with it.
- An isolated or community VLAN can have only one primary VLAN associated with it.

- Enable PortFast and BPDU guard on isolated and community ports to prevent spanning tree loops due to misconfigurations. When enabled, STP applies the BPDU guard feature to all PortFast-configured Layer 2 LAN ports.
- If you delete a VLAN used in the PVLAN configuration, the PVLAN ports associated with the VLAN become inactive.
- PVLAN ports do not have to be on the same network device as long as the devices are trunk-connected and the primary and secondary VLANs have not been removed from the trunk.
- VTP does not support PVLANs. You must configure PVLANs on each device where you want PVLAN ports.
- To maintain the security of your PVLAN configuration and avoid other use of the VLANs configured as PVLANs, configure PVLANs on all intermediate devices, even devices that have no PVLAN ports.
- Cisco recommends that you prune the PVLANs from the trunks on devices that carry no traffic in the PVLANs.
- You can apply different QoS configurations to primary, isolated, and community VLANs.
- To apply IOS output ACLs to all outgoing PVLAN traffic, configure them on the Layer 3 VLAN interface of the primary VLAN.
- IOS ACLs applied to the Layer 3 VLAN interface of a primary VLAN automatically apply to the associated isolated and community VLANs.
- Do not apply IOS ACLs to isolated or community VLANs. IOS ACL configuration applied to isolated and community VLANs is inactive while the VLANs are part of the PVLAN configuration.
- Do not apply dynamic ACEs to primary VLANs. IOS dynamic ACL configuration applied to a primary VLAN is inactive while the VLAN is part of the PVLAN configuration.
- You can stop Layer 3 switching on an isolated VLAN by deleting the mapping of that VLAN with its primary VLAN.

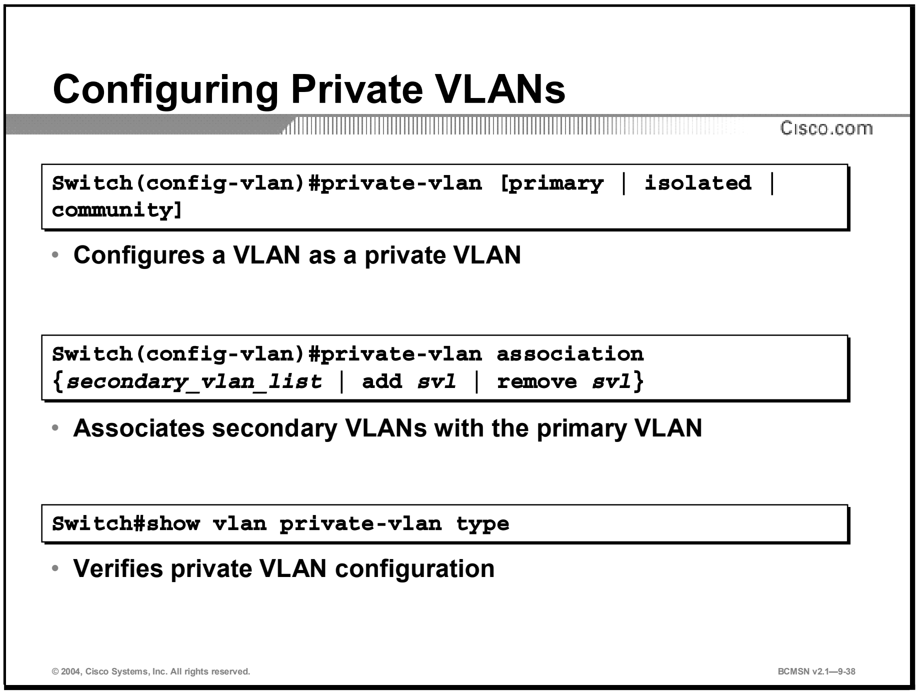

Configuring Private VLANs

Cisco.com

```
Switch(config-vlan)#private-vlan [primary | isolated |
community]
```

- Configures a VLAN as a private VLAN

```
Switch(config-vlan)#private-vlan association
{secondary_vlan_list | add svl | remove svl}
```

- Associates secondary VLANs with the primary VLAN

```
Switch#show vlan private-vlan type
```

- Verifies private VLAN configuration

BCMSN v2.1—9-38

To configure a PVLAN, follow these steps:

Step 1 Set VTP mode to transparent.

Step 2 Create the secondary VLANs.

Note Isolated and community VLANs are secondary VLANs.

Step 3 Create the primary VLAN.

Step 4 Associate the secondary VLAN to the primary VLAN.

Note Only one isolated VLAN can be mapped to a primary VLAN, but more than one community VLAN can be mapped to a primary VLAN.

Step 5 Configure an interface to an isolated or community port.

Step 6 Associate the isolated port or community port to the primary-secondary VLAN pair.

Step 7 Configure an interface as a promiscuous port.

Step 8 Map the promiscuous port to the primary-secondary VLAN pair.

Use these commands to configure a VLAN as a PVLAN:

```
Switch(config)#vlan vlan_ID
Switch(config-vlan)#[no] private-vlan {isolated | primary}
```

Example: PVLAN Configurations

This example shows how to configure VLAN202 as a primary VLAN and verify the configuration:

```
Switch#configure terminal
Switch(config)#vlan 202
Switch(config-vlan)#private-vlan primary
Switch(config-vlan)#end
Switch#show vlan private-vlan type

Primary Secondary Type              Interfaces
------- --------- ----------------- --------------------------
----------------
202               primary
```

This example shows how to configure VLAN440 as an isolated VLAN and verify the configuration:

```
Switch#configure terminal
Switch(config)#vlan 440
Switch(config-vlan)#private-vlan isolated
Switch(config-vlan)#end
Switch#show vlan private-vlan type

Primary Secondary Type              Interfaces
------- --------- ----------------- --------------------------
-------
202               primary
440               isolated
```

To associate secondary VLANs with a primary VLAN, perform this procedure:

```
Switch(config)#vlan primary_vlan_ID
Switch(config-vlan)#[no] private-vlan association
{secondary_vlan_list | add secondary_vlan_list | remove
secondary_vlan_list}
```

When you associate secondary VLANs with a primary VLAN, note the following:

- The *secondary_vlan_list* parameter contains only one isolated VLAN ID.
- Use the **remove** keyword with the *secondary_vlan_list* variable to clear the association between the secondary VLAN and the primary VLAN. The list can contain only one VLAN.
- Use the **no** keyword to clear all associations from the primary VLAN.
- The command does not take effect until you exit VLAN configuration submode.

Example: Associating an Isolated VLAN with a Primary VLAN

This example shows how to associate isolated VLAN440 with primary VLAN202 and verify the configuration:

```
Switch# configure terminal
Switch(config)# vlan 202
Switch(config-vlan)# private-vlan association 440
Switch(config-vlan)# end
Switch# show vlan private-vlan

Primary Secondary Type              Interfaces
------- --------- ----------------- --------------------------
-------
202     440       isolated
```

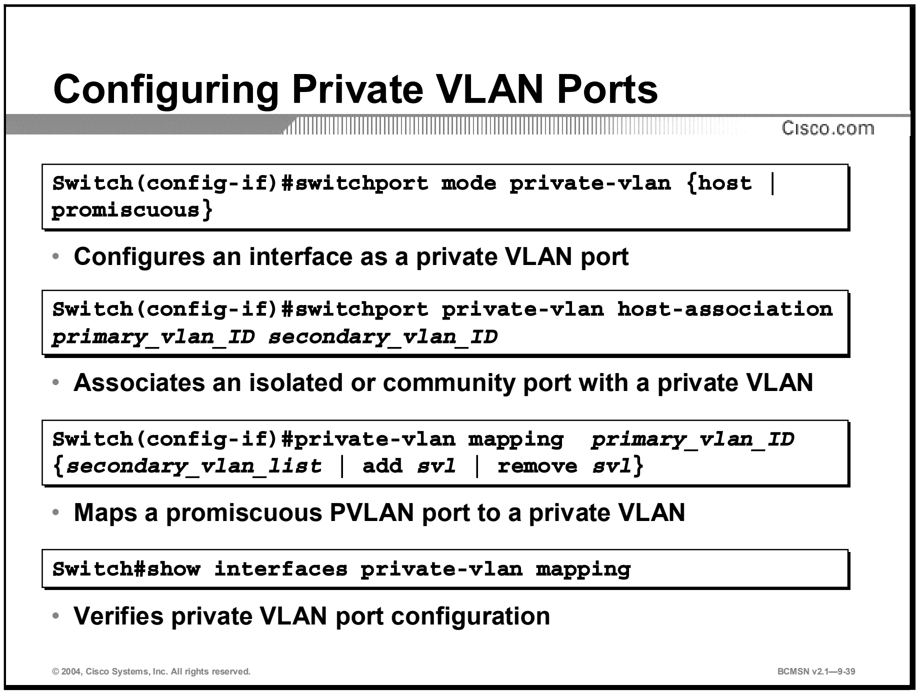

Configuring Private VLAN Ports

Cisco.com

```
Switch(config-if)#switchport mode private-vlan {host |
promiscuous}
```

- Configures an interface as a private VLAN port

```
Switch(config-if)#switchport private-vlan host-association
primary_vlan_ID secondary_vlan_ID
```

- Associates an isolated or community port with a private VLAN

```
Switch(config-if)#private-vlan mapping  primary_vlan_ID
{secondary_vlan_list | add svl | remove svl}
```

- Maps a promiscuous PVLAN port to a private VLAN

```
Switch#show interfaces private-vlan mapping
```

- Verifies private VLAN port configuration

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—9-39

To configure a Layer 2 interface as a PVLAN promiscuous port, perform this procedure:

```
Switch(config)#interface {fastethernet | gigabitethernet}
slot/port
Switch(config-if)#switchport mode private-vlan {host |
promiscuous}
Switch(config-if)#[no] switchport private-vlan mapping
primary_vlan_ID {secondary_vlan_list | add secondary_vlan_list
| remove secondary_vlan_list}
```

When you configure a Layer 2 interface as a PVLAN promiscuous port, note the following:

- The *secondary_vlan_list* parameter cannot contain spaces. It can contain multiple comma-separated items. Each item can be a single PVLAN ID or a hyphenated range of PVLAN IDs.
- Enter a *secondary_vlan_list* or use the **add** keyword with a *secondary_vlan_list* to map the secondary VLANs to the PVLAN promiscuous port.
- Use the **remove** keyword with a *secondary_vlan_list* to clear the mapping between secondary VLANs and the PVLAN promiscuous port.
- Use the **no** keyword to clear all mapping from the PVLAN promiscuous port.

Example: Configuring PVLAN Ports

This example shows how to configure interface FastEthernet 5/2 as a PVLAN promiscuous port, map it to a PVLAN, and verify the configuration:

```
Switch#configure terminal
Switch(config)#interface fastethernet 5/2
Switch(config-if)#switchport mode private-vlan promiscuous
Switch(config-if)#switchport private-vlan mapping 202 440
Switch(config-if)#end
Switch#show interfaces fastethernet 5/2 switchport
Name: Fa5/2
Switchport: Enabled

Administrative Mode: private-vlan promiscuous
Operational Mode: down
Administrative Trunking Encapsulation: negotiate
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative private-vlan host-association: none
((Inactive))

Administrative private-vlan mapping: 202 (VLAN0202) 440
(VLAN0440)

Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
```

To configure a Layer 2 interface as a PVLAN host port, perform this procedure:

```
Switch(config)#interface {fastethernet |  gigabitethernet}
slot/port
Switch(config-if)#switchport mode private-vlan {host |
promiscuous}
Switch(config-if)#[no] switchport private-vlan host-
association primary_vlan_ID secondary_vlan_ID
```

This example shows how to configure interface FastEthernet 5/1 as a PVLAN host port and verify the configuration:

```
Switch#configure terminal
Switch(config)#interface fastethernet 5/1
Switch(config-if)#switchport mode private-vlan host
Switch(config-if)#switchport private-vlan host-association 202
440
Switch(config-if)#end
Switch#show interfaces fastethernet 5/1 switchport
Name: Fa5/1
Switchport: Enabled

Administrative Mode: private-vlan host
Operational Mode: down
Administrative Trunking Encapsulation: negotiate
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)

Administrative private-vlan host-association: 202 (VLAN0202)
Administrative private-vlan mapping: none

Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
```

To permit routing of secondary VLAN ingress traffic, perform this procedure:

```
Switch(config)#interface vlan primary_vlan_ID
Switch(config-if)#[no] private-vlan mapping  primary_vlan_ID
{secondary_vlan_list | add secondary_vlan_list | remove
secondary_vlan_list}
```

When you permit routing on the secondary VLAN ingress traffic, note the following:

- Enter a value for the *secondary_vlan_list* variable or use the **add** keyword with the *secondary_vlan_list* variable to map the secondary VLANs to the primary VLAN.
- Use the **remove** keyword with the *secondary_vlan_list* variable to clear the mapping between secondary VLANs and the primary VLAN.
- Use the **no** keyword to clear all mapping from the primary VLAN.

Example: Permitting Routing of Secondary VLAN Ingress Traffic

This example shows how to permit routing of secondary VLAN ingress traffic from PVLAN440 and verify the configuration:

```
Switch#configure terminal
Switch(config)#interface vlan 202
Switch(config-if)#private-vlan mapping add 440
Switch(config-if)#end
Switch#show interfaces private-vlan mapping
Interface Secondary VLAN Type
--------- -------------- -----------------
vlan202   440            isolated
```

Summary

This topic summarizes the key points discussed in this lesson.

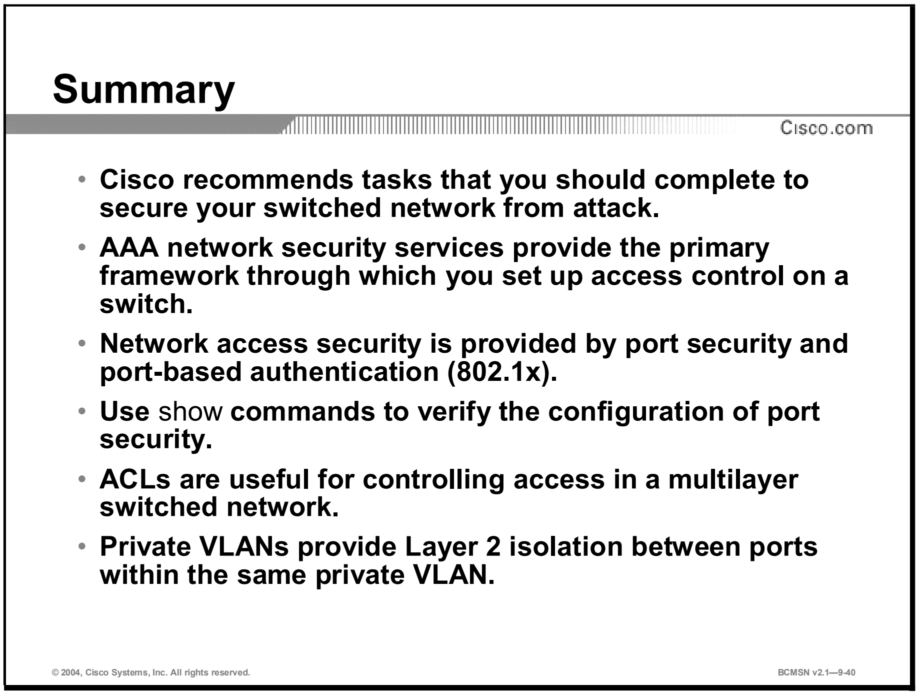

Summary

Cisco.com

- Cisco recommends tasks that you should complete to secure your switched network from attack.
- AAA network security services provide the primary framework through which you set up access control on a switch.
- Network access security is provided by port security and port-based authentication (802.1x).
- Use show commands to verify the configuration of port security.
- ACLs are useful for controlling access in a multilayer switched network.
- Private VLANs provide Layer 2 isolation between ports within the same private VLAN.

BCMSN v2.1—9-40

References

For additional information, refer to this resource:

- "Cisco IOS Security Configuration Guide" at http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_book09186a00800ca5b2.html

Next Step

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 9-1: Optimizing and Securing Multilayer Switched Networks

Quiz

Use the practice items below to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What are three actions that you can take to secure multilayer switches? (Choose three.)

A) set system passwords

B) enable SNMP write functions

C) secure the spanning tree topology

D) configure basic access control lists

E) enable the integrated HTTP daemon

Q2) Which command correctly enables AAA security?

A) **aaa new-model**

B) **ppp authentication chap apple**

C) **aaa group server radius rad2only**

D) **aaa authentication login default group radius**

Q3) What is specified by the "if-authenticated" method?

A) access to the function is always allowed

B) access to the function is allowed if the user was authenticated

C) access to the function is allowed based on the local user database

D) access to the function is allowed based on RADIUS authentication

Q4) Which command correctly enables port security?

A) Switch(config)#**switchport security**

B) Switch(config-if)#**switchport security**

C) Switch(config)#**switchport port-security**

D) Switch(config-if)#**switchport port-security**

Q5) What are the three roles in 802.1x port authentications? (Choose three.)

A) router

B) supplicant

C) authenticator

D) accounting server

E) authentication server

Q6) Which command displays port security information from the MAC address table?

A) **show port-security**

B) **show port-security address**

C) **show port-security mac-table**

D) **show port-security interface fastethernet 5/1**

Q7) Which IOS command applies a VLAN access map to VLANs?

A) **show vlan filter**

B) **vlan access-map** *map_name* [*seq#*]

C) **match {ip | ipx | mac }** *address*

D) **vlan filter** *map_name* **vlan_list** *list*

Q8) Which command correctly configures a PVLAN?

A) Switch(config)#**private-vlan primary**

B) Switch(config)#**private-vlan 202 isolated**

C) Switch(config-vlan)#**private-vlan isolated**

D) Switch(config-vlan)#**private-vlan promiscuous**

Quiz Answer Key

Q1) A, C, D

Relates to: Security in the Multilayer Switched Network

Q2) A

Relates to: Configuring AAA

Q3) B

Relates to: Configuring AAA

Q4) D

Relates to: Configuring Network Access Security

Q5) B, C, E

Relates to: Configuring Network Access Security

Q6) B

Relates to: Practice: Verifying Network Access Security

Q7) D

Relates to: Practice: Configuring Security Using Access Lists

Q8) C

Relates to: Practice: Configuring PVLANs

Lesson Assessments

Overview

Use the lesson assessments here to test what you learned in this module. The correct answers and solutions are found in the Lesson Assessment Answer Key.

Outline

This section includes these assessments:

- Quiz 9-1: Optimizing Multilayer Switched Networks
- Quiz 9-2: Securing Multilayer Switched Networks

Quiz 9-1: Optimizing Multilayer Switched Networks

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Describe techniques to enhance the performance of a multilayer switched network
- Monitor switch ports using SPAN and VLAN-based SPAN
- Monitor switch ports using RSPAN
- Describe the features and operation of network analysis modules on Catalyst switches to improve network traffic management
- Verify and troubleshoot the operation of network analysis modules

Quiz

Answer these questions:

Q1) Performance management is the practice of optimizing _____.

A) network security

B) capacity planning

C) user performance

D) network service response time

Q2) What function can an interface configured as a SPAN destination serve?

A) receive SPAN traffic

B) receive network traffic

C) forward SPAN traffic and other network traffic

D) forward SPAN traffic and receive SPAN traffic

Q3) What is the purpose of the RSPAN VLAN?

A) carries RSPAN source traffic

B) carries RSPAN session traffic

C) carries RSPAN destination traffic

D) carries RSPAN management traffic

Q4) Which command correctly configures a NAM to receive and monitor SPAN traffic?

A) **monitor session 1 source interface gi 8/1**

B) **monitor session 1 source interface gi 8/2**

C) **monitor session 1 destination interface gi 8/1**

D) **monitor session 1 destination interface gi 8/2**

Q5) Which command correctly displays information about the SPAN destination port on the NAM?

A) **show module GigabitEthernet 8/1**

B) **show module GigabitEthernet 8/2**

C) **show interface GigabitEthernet 8/1**

D) **show interface GigabitEthernet 8/2**

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Quiz 9-2: Securing Multilayer Switched Networks

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Explain basic security concepts for the multilayer switched network
- Configure authentication, authorization, and accounting on Catalyst switches
- Configure port security and port-based authentication with 802.1x
- Verify the network access security configuration
- Configure VLAN access lists
- Verify the VLAN access list security configuration

Quiz

Answer these questions:

Q1) For switch security, create basic access lists to limit management and remote access traffic to _____.

A) managers

B) external users

C) trusted subnets

D) all badged employees

Q2) What is accomplished with the command **aaa authentication login local**?

A) specifies authentication with the line password

B) specifies authentication with the local password

C) specifies authentication with the enable password

D) specifies authentication with the case-sensitive local password

Q3) Which command correctly configures default network authorization using TACACS+, RADIUS, and the local user database, in that order?

A) **aaa authorization network default group tacacs+ radius local**

B) **aaa authorization network default group local tacacs+ radius**

C) **aaa authorization network default group tacacs+ local radius**

D) **aaa authentication network default group tacacs+ radius local**

Q4) Which command specifies that an interface be disabled when a port security violation occurs?

A) **switchport port-security violation disable**

B) **switchport port-security violation protect**

C) **switchport port-security violation restrict**

D) **switchport port-security violation shutdown**

Q5) Which command includes the aging time and aging type for an interface in the output?

A) **show port-security**

B) **show port-security address**

C) **show port-security mac-table**

D) **show port-security interface fastethernet 5/1**

Q6) What is the basis for VACL filter traffic?

A) based on capture port

B) based on source address

C) based on router access list configuration

D) based on Ethertype and MAC address traffic

Q7) Which type of PVLAN port is completely separated from other ports except for one?

A) isolated

B) separated

C) community

D) promiscuous

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Lesson Assessment Answer Key

Quiz 9-1: Optimizing Multilayer Switched Networks

Q1) D

Q2) A

Q3) B

Q4) C

Q5) C

Quiz 9-2: Securing Multilayer Switched Networks

Q1) C

Q2) B

Q3) A

Q4) D

Q5) B

Q6) D

Q7) A

Module 10

Understanding Metro Ethernet

Overview

Service providers offering Metro Ethernet services (and enterprise customers using such services) have several choices for connectivity and transport mechanisms. The choice of implementations depends on several factors, including current network infrastructure and desired level of service. Each choice offers its own benefits and drawbacks, as well as having a specific set of implementation requirements.

Upon completing this module, you will be able to:

- Describe Metro Ethernet connectivity and Layer 1 transport options
- Describe Metro Ethernet tunneling options

Outline

The module contains these components:

- Examining Metro Ethernet Connectivity Services and Layer 1 Transport Options
- Examining Metro Ethernet Tunneling
- Lesson Assessments

Examining Metro Ethernet Connectivity Services and Layer 1 Transport Options

Overview

Service providers require networks that are integrated with their legacy infrastructures and are flexible to respond to changing market conditions and diverse customer requirements. These requirements are especially important in metropolitan-area networks (MANs) because of the critical interface between enterprise and service provider networks. Cisco Systems offers a variety of Metro networking solutions with a choice of Layer 1 implementation options.

Relevance

Network administrators and designers alike need to understand the implementation choices available for WANs and MANs to better design and deploy network options.

Objectives

Upon completing this lesson, you will be able to:

- Describe Cisco Metro Ethernet networking solutions
- List the criteria used to evaluate Metro Ethernet connectivity options
- Describe Transparent LAN Service, its benefits and disadvantages
- Describe Directed VLAN Service, its benefits and disadvantages
- Describe the benefits and disadvantages of DWDM as a Layer 1 implementation of Metro Ethernet
- Describe the benefits and disadvantages of SONET as a Layer 1 implementation of Metro Ethernet
- Describe the benefits and disadvantages of CWDM as a Layer 1 implementation of Metro Ethernet

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Metro Ethernet Solutions
- Metro Ethernet Connectivity
- Transparent LAN Service
- Directed VLAN Service
- Metro Ethernet over SONET
- Metro Ethernet over DWDM
- Metro Ethernet over CWDM
- Summary
- Quiz

Metro Ethernet Solutions

Cisco Metro Ethernet offerings include Metro IP, Metro Ethernet switching, and Metro optical transport platforms and technologies. This topic describes Metro Ethernet solutions.

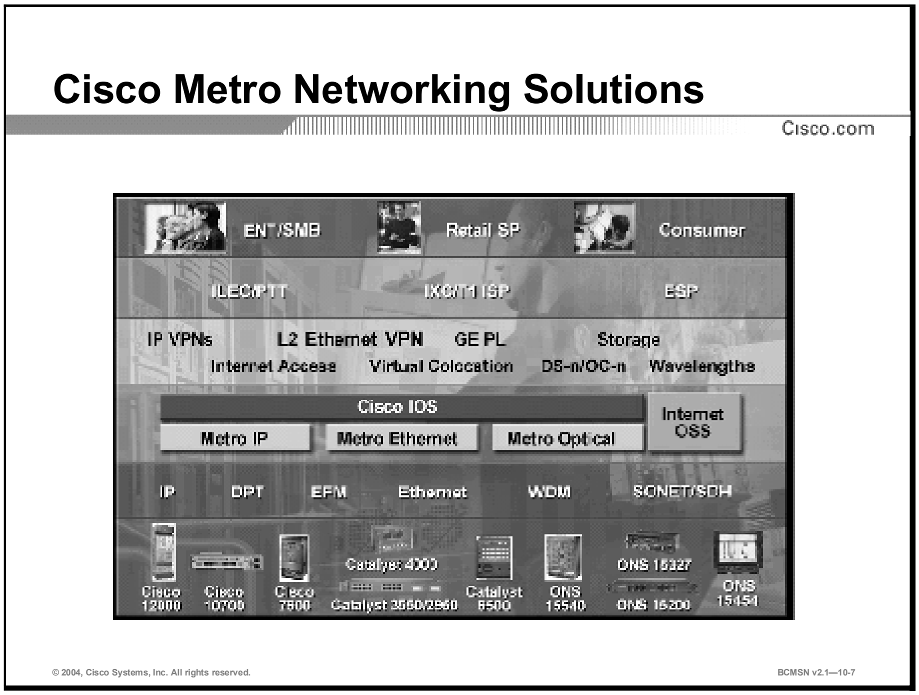

Cisco Metro Networking Solutions

Cisco.com

BCMSN v2.1—10-7

Service providers require networks that are integrated with their legacy infrastructures and are flexible to respond to changing market conditions and diverse customer requirements. These requirements are especially important in MANs because of the critical interface between enterprise and service provider networks. Ethernet has emerged as the leading access medium for a variety of reasons:

- Enables most cost-effective services
- Provides ample bandwidth between the enterprise and the MAN
- Eliminates access bottlenecks
- Allows service providers to deliver multiple data, voice, and video services to enterprise and consumer customers over a high-speed access connection

Cisco Metro solutions deliver a comprehensive multilayer service portfolio for providers to quickly scale their customer base and revenues. This portfolio, which includes Metro IP, Metro Ethernet switching, and Metro optical transport platforms and technologies, can ensure a flexible and efficient foundation for profitable Metro services.

Cisco delivers Metro solutions that enable service providers to grow quickly and profitably with a comprehensive multilayer service portfolio. This service portfolio is enabled through a modular system, which allows service providers deployment flexibility with operational consistency across all Metro environments. The Cisco modular system enables the Metro Ethernet network with both 10 Gigabit Ethernet and 10-Gbps SONET/Synchronous Digital Hierarchy (SDH) OC-192. These benefits are achieved while delivering a high-availability transport and data infrastructure. The figure shows the Cisco portfolio of Metro products. The acronyms and abbreviations used in the figure are:

- Ethernet in the first mile (EFM)
- Ethernet service provider (ESP)
- Dynamic Packet Transport (DPT)
- Gigabit Ethernet Private Line (GE PL)
- Incumbent Local Exchange Carrier/Post, Telephone, and Telegraph (ILEC/PTT)
- Inter-exchange Carrier/Tier 1 Internet Service Provider (IXC/T1 ISP)
- Layer 2 Ethernet Virtual Private Network (VPN)
- Operations Support System (OSS)
- Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH)
- Virtual colocation
- Wavelength division multiplexing (WDM)
- Small and medium business (SMB)
- Enterprise market (ENT)

The Cisco optical platforms shown in the figure include the following:

- **ONS 15540:** Extended services platform with external cross-connect capability (ESPx) integrates data networking, storage area networking (SAN), time-division multiplexing (TDM), SONET, and SDH technologies
- **ONS 15327:** Cisco optical Metro edge and access platform for service providers and enterprise customers
- **ONS 15200:** Cisco metropolitan dense wavelength division multiplexing (DWDM) solution
- **ONS 15454:** Cisco SONET/SDH multiservice provisioning platform providing SONET/SDH transport through OC-192/STM-64, integrated DWDM optical networking, and multiservice interfaces including Ethernet and traditional TDM from DS1 and E1 to OC-192 and STM-64

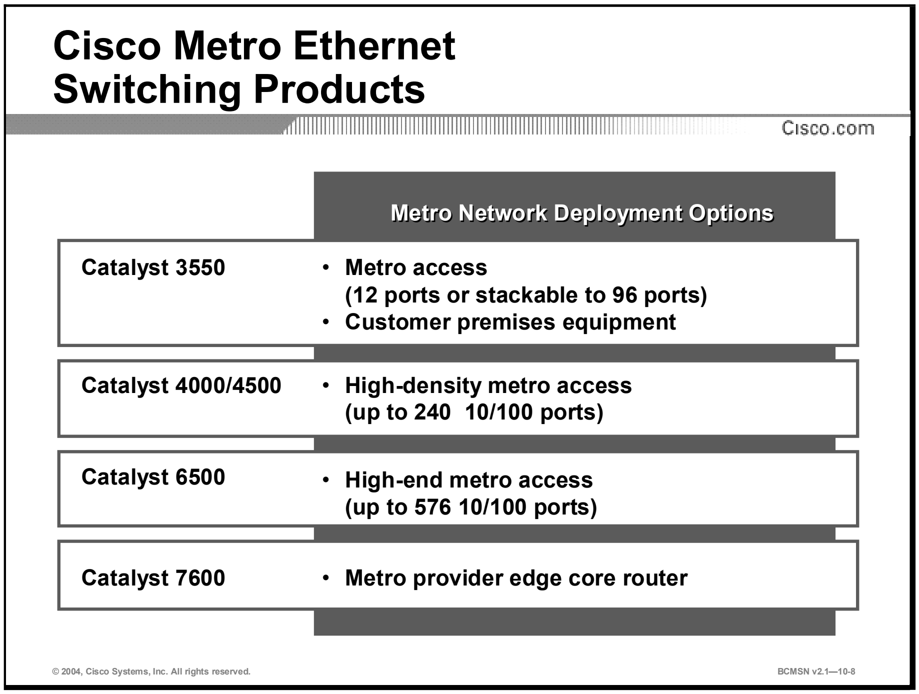

Cisco Metro Ethernet Switching Products

Cisco.com

	Metro Network Deployment Options
Catalyst 3550	• Metro access (12 ports or stackable to 96 ports) • Customer premises equipment
Catalyst 4000/4500	• High-density metro access (up to 240 10/100 ports)
Catalyst 6500	• High-end metro access (up to 576 10/100 ports)
Catalyst 7600	• Metro provider edge core router

BCMSN v2.1—10-8

The table describes the Cisco Metro Ethernet products.

Cisco Product	Metro Network Deployment Options	Key Metro Features
Cisco Catalyst 3550	Metro access (12 ports or stackable to 96 ports) Customer premises equipment (CPE)	■ Bandwidth policing ■ Two queues per port ■ Wire-speed performance ■ Access control lists (ACLs) (Layer 2 and 3) ■ Unidirectional Link Detection (UDLD) Protocol ■ Port security ■ UplinkFast and BackboneFast ■ EtherChannel ■ Jumbo frames ■ 13 million pps
Catalyst 4000/4500	High-density Metro access (up to 240 10/100 ports)	■ Two queues per port with priority queuing (PQ) ■ Post ■ 802.1w Rapid Spanning Tree Protocol (RSTP) ■ UplinkFast and BackboneFast ■ UDLD protocol ■ EtherChannel ■ Jumbo frames

Cisco Product	Metro Network Deployment Options	Key Metro Features
		■ 24 million pps forwarding
Catalyst 6500	High-end Metro access (up to 576 10/100 ports) Layer 2 POP distribution and core switch	■ Bandwidth policing ■ Three queues per port with PQ ■ 802.1w RSTP ■ UplinkFast and BackboneFast ■ UDLD Protocol ■ Multimodule EtherChannel ■ Full switching fabric and Supervisor redundancy ■ Supervisor high availability ■ 210 million pps forwarding
Catalyst 7600	Metro provider edge core router	■ Bandwidth policing ■ 1000 queues per port with PQ ■ 802.1w RSTP ■ UplinkFast and BackboneFast ■ UDLD Protocol ■ Multimodule EtherChannel ■ Full switching fabric and Supervisor redundancy ■ Supervisor high availability ■ Ethernet over Multiprotocol Label Switching (EoMPLS) ■ Multiprotocol Label Switching (MPLS)

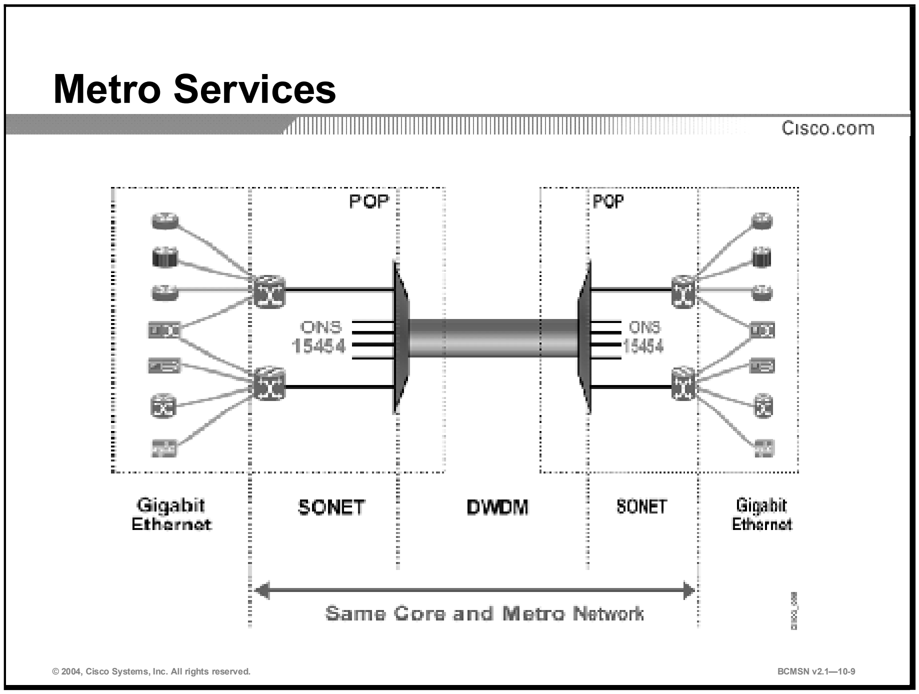

Metro Services

Cisco.com

BCMSN v2.1—10-9

Enterprises can deploy Metro Ethernet for POP connectivity for IP devices along with other services. Gigabit Ethernet speeds offer enterprises maximum cost savings.

All high-speed IP devices and other service switches (Layers 4 to 7) have gigabit connections and still can offer TDM, DS1-3, and OC-192 services.

Metro Ethernet Connectivity

You can transmit Ethernet packets across the WAN using a variety of methods. Which method you use depends on such factors as cost-effectiveness, service levels, transparency, and scalability. This topic describes Metro Ethernet connectivity.

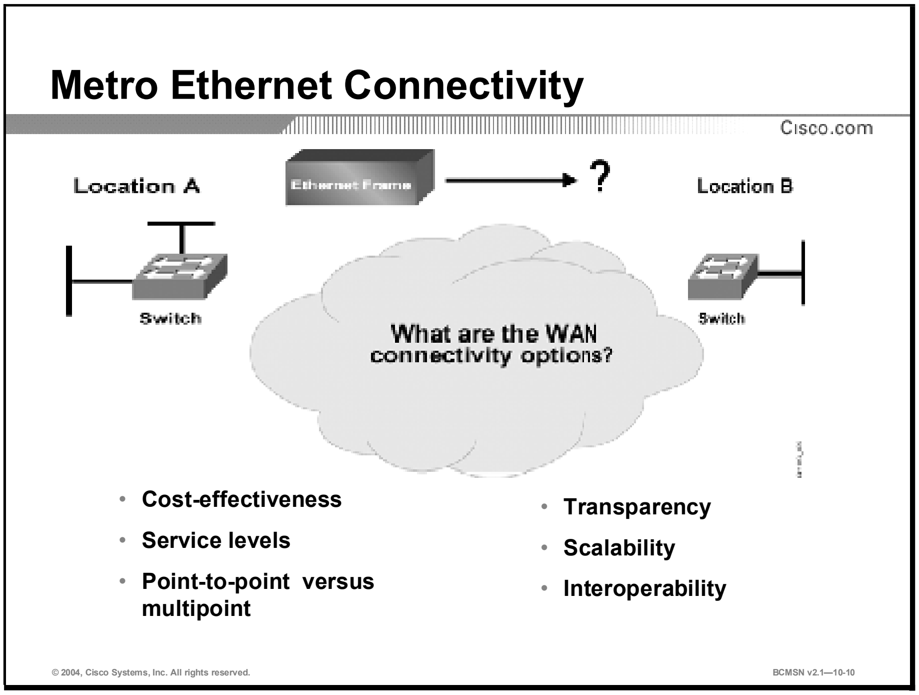

Metro Ethernet Connectivity

Cisco.com

- Cost-effectiveness
- Service levels
- Point-to-point versus multipoint
- Transparency
- Scalability
- Interoperability

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—10-10

Several options are available for sending Ethernet frames across the WAN, ranging from using point-to-point Layer 1 connections such as DWDM to more complex Layer 3 encapsulation protocols that allow Ethernet services to be emulated over packet networks. To evaluate options, you can use these criteria:

- **Cost-effectiveness:** In order for Metro Ethernet access services to live up to their promise of reduced-cost network services, they must be delivered over a network architecture that can be established with relatively low capital expenditure, and maintained with relatively low operation costs.
- **Service levels:** End users have become accustomed to receiving a predefined service level from their service providers for traditional data communication services. The selection of Metro Ethernet connectivity options needs to account for the characteristics of the service levels that can be provided to the end user. These service-level definitions can be as simple as defining the level of service availability on the Metro Ethernet connection, or as complex as defining different classes of service for different types of traffic, including such parameters as latency and committed information rate (CIR).
- **Point-to-point versus multipoint:** The concept of multipoint connectivity has always existed with traditional Layer 2 data communication services. With Metro Ethernet access services, the concept of multipoint connectivity is slightly different from typical Frame Relay-like multipoint connections. Instead of the CPE supporting multiple virtual circuits (VCs) over a single physical interface, there is only a single logical interface from the CPE to the provider edge. At the provider edge, this single logical interface can be mapped into multiple point-to-point circuits to establish a single Layer 2 broadcast domain. With this model of multipoint connectivity, it is possible, but not necessarily desirable, to eliminate

the Layer 3 function on the CPE. For larger multipoint topologies, enterprises should consider a Layer 3 control plane, such as a Multiprotocol Label Switching Virtual Private Network (MPLS VPN).

- **Transparency:** Because Ethernet is the predominant technology used in enterprise networks, Metro Ethernet can accommodate a transparent interface to the enterprise so that it does not need to make any changes to the Layer 2 VLAN configurations within its own networks, even if the connection to the Metro service provider is based on Layer 2 rather than Layer 3 protocols. This capability did not exist with traditional data communication services such as ATM, Frame Relay, or leased lines. Layer 2 transport protocols were different from the Ethernet Layer 2 protocol used in the enterprise LAN, so that a Layer 3 interface from the enterprise to the service provider was a technical requirement.
- **Scalability:** Scalability refers to both the number of end users and the geographic distribution of those customers. As the name suggests, Metro Ethernet access services originated to provide high-speed connectivity within a Metro service area. However, as Metro Ethernet access services evolve, enterprises want an Ethernet service interface that connects the enterprise to the service provider to provide general-purpose WAN connectivity to either private or public networks. For some service providers, it may be acceptable to limit the Metro service definition to a specific metropolitan area. For other service providers, it may be a requirement that the Metro Ethernet access service definition scale to inter-Metro distances. In either case, the control plane must allow new Metro access and aggregation network elements to connect into the core Metro network without major disruption or reconfiguration.

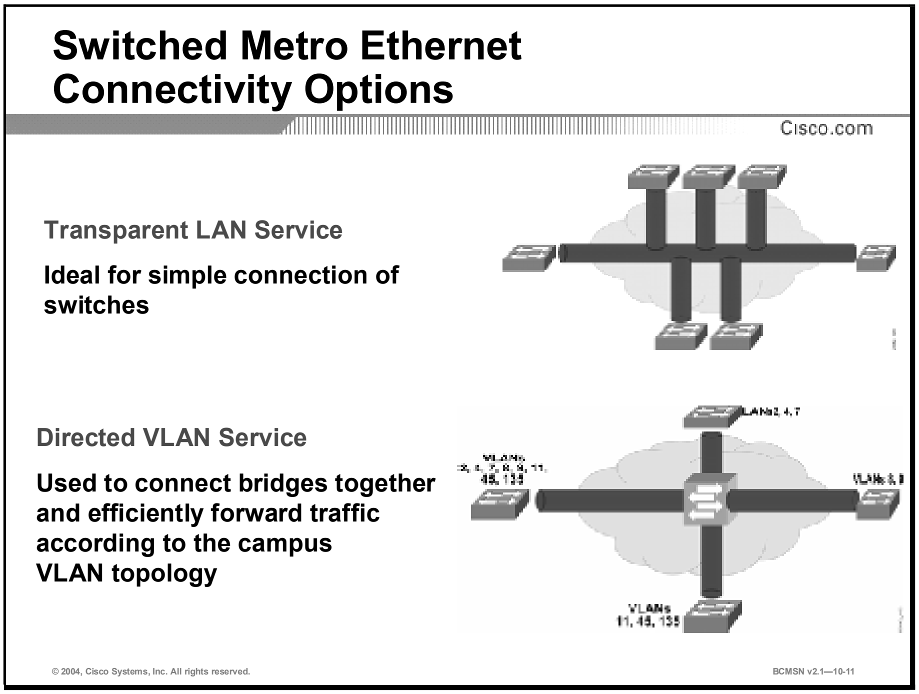

Switched Metro Ethernet Connectivity Options

Cisco.com

Transparent LAN Service

Ideal for simple connection of switches

Directed VLAN Service

Used to connect bridges together and efficiently forward traffic according to the campus VLAN topology

BCMSN v2.1—10-11

Connectivity options determine how traffic is handled as it passes through the center of the network. With a Transparent LAN Service (TLS), the traffic is sent as a multipoint-to-multipoint connection, and all users view all traffic. With a Directed VLAN Service (DVS), the traffic is forwarded according to the campus VLAN topology.

Transparent LAN Service

With a TLS, all of the customer sites are connected together on a single subnet or VLAN. TLS is easy to implement but presents scalability issues and a single failure domain for all customers. This topic describes TLS.

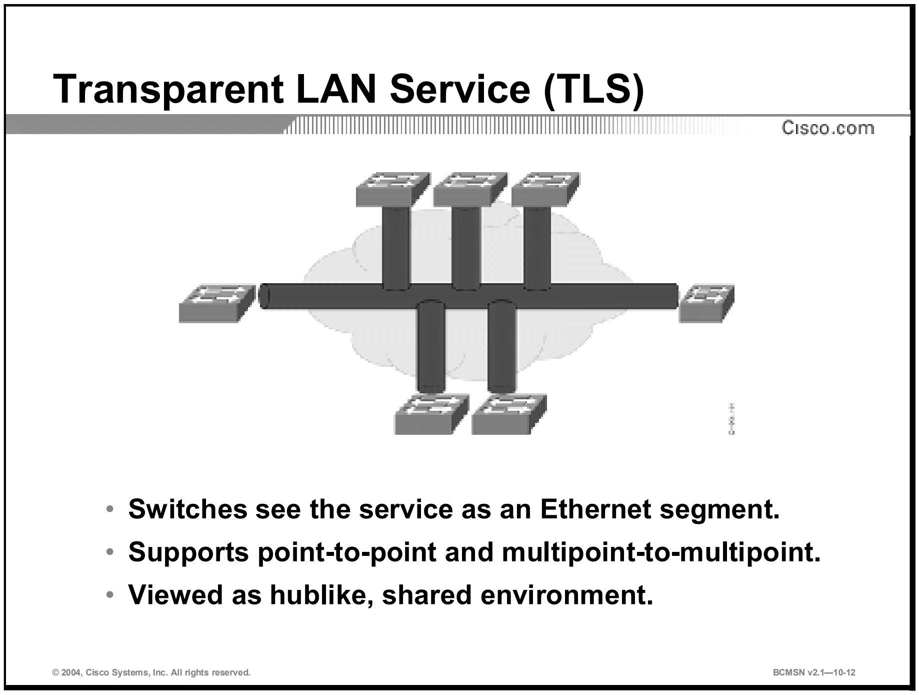

Transparent LAN Service (TLS)

- Switches see the service as an Ethernet segment.
- Supports point-to-point and multipoint-to-multipoint.
- Viewed as hublike, shared environment.

In a TLS, the network looks like an Ethernet hub. The TLS supports point-to-point and point-to-multipoint operation. As far as the enterprise is concerned, all of its routers and multilayer switches are on the same subnet. The service provider can easily provide a TLS solution for the enterprise by maintaining the same VLAN across its entire infrastructure. The figure demonstrates how the network will look from the enterprise perspective.

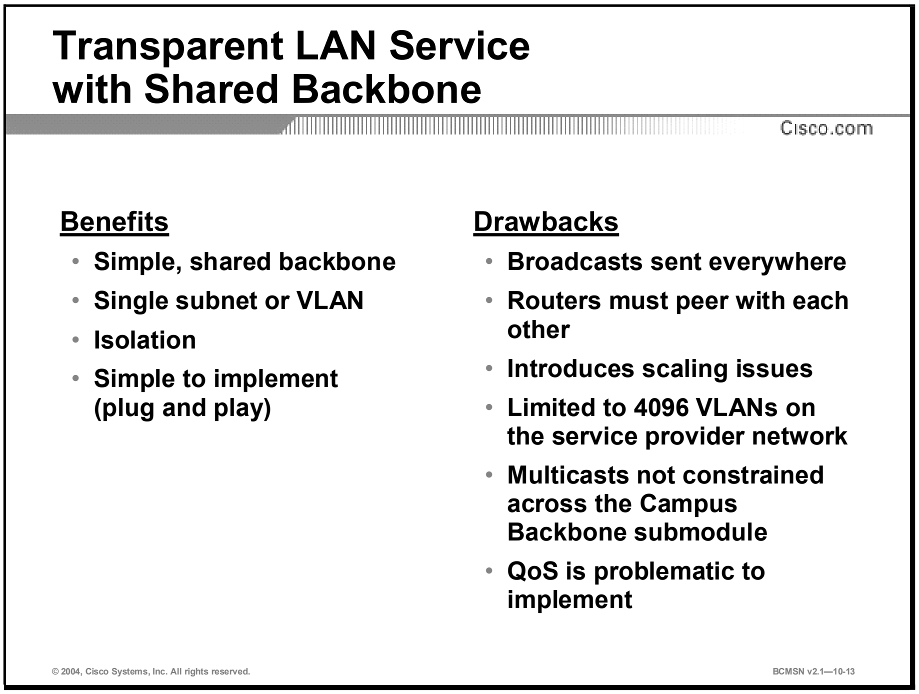

Transparent LAN Service with Shared Backbone

Benefits

- Simple, shared backbone
- Single subnet or VLAN
- Isolation
- Simple to implement (plug and play)

Drawbacks

- Broadcasts sent everywhere
- Routers must peer with each other
- Introduces scaling issues
- Limited to 4096 VLANs on the service provider network
- Multicasts not constrained across the Campus Backbone submodule
- QoS is problematic to implement

BCMSN v2.1—10-13

The key benefit of the TLS is its simplicity of implementation due to its shared backbone. All remote locations can appear on a single subnet or in a single VLAN.

However, this implementation has the following potentially serious performance implications for the enterprise:

- All users are in the same failure domain.
- Broadcasts will be sent to all sites.
- All the routers must peer with each other across the MAN, which causes a scalability problem for routers that need to maintain routing adjacency across the network.
- A VLAN is limited to a maximum of 4096 customer VLANs, which may be insufficient for customers.
- Multicast traffic cannot be constrained as it flows across the core.
- Quality of service (QoS) is difficult to implement.

Directed VLAN Service

VLAN IDs are used with the DVS to select destinations. This topic describes the DVS.

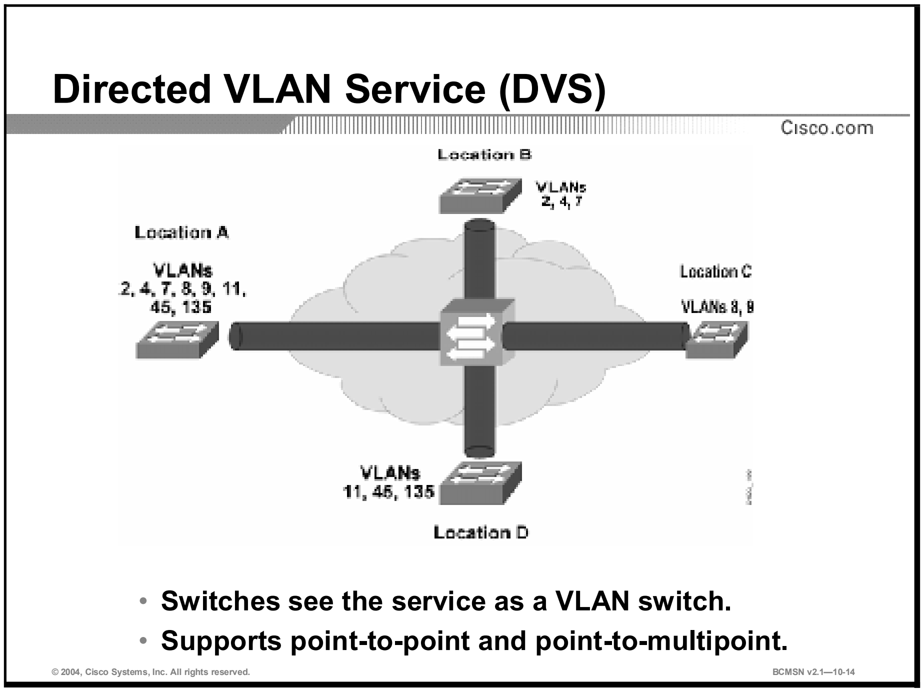

Directed VLAN Service (DVS)

Cisco.com

- **Switches see the service as a VLAN switch.**
- **Supports point-to-point and point-to-multipoint.**

BCMSN v2.1—10-14

In a DVS, VLAN IDs are used to select destinations. The DVS can be point-to-point or point-to-multipoint. Typically, the VLANs of the enterprise will be isolated from the VLANs of the network using additional header information.

The switch in the core sees the VLANs defined at the edge of the network and provides VLAN switching. The core switch must know the enterprise VLAN assignments. This implementation requires Per VLAN Spanning Tree+ (PVST+).

The DVS allows enterprises to support VLANs that can appear in multiple locations across the core. In the figure, VLAN2 appears in locations A and B, while VLAN11 appears in locations A and D. The core switch in the cloud provides VLAN switching capability. Therefore, the core switch must know about VLAN assignments on the remote location switches.

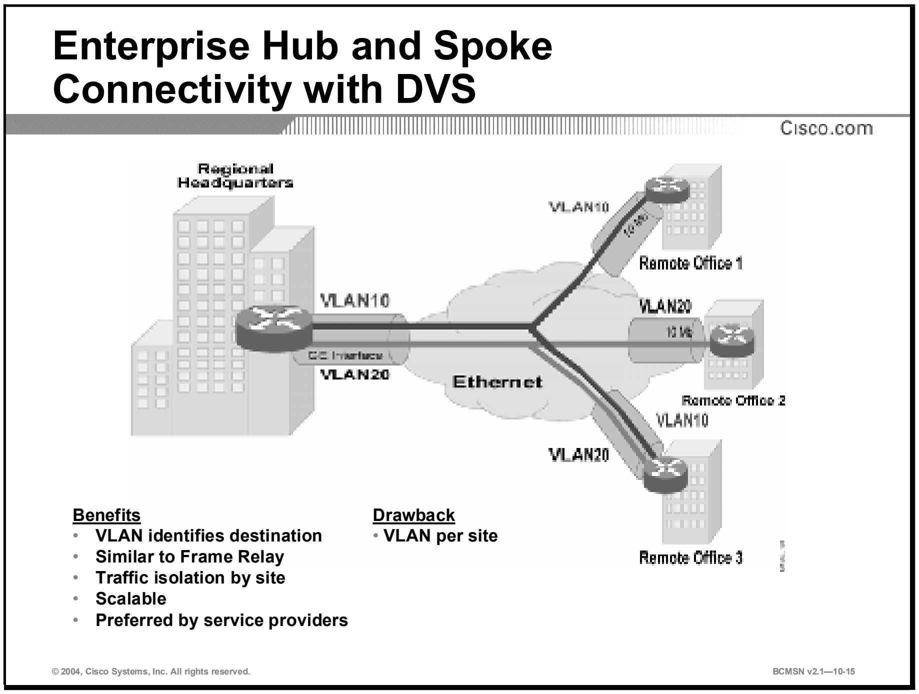

Enterprise Hub and Spoke Connectivity with DVS

Cisco.com

Benefits

- VLAN identifies destination
- Similar to Frame Relay
- Traffic isolation by site
- Scalable
- Preferred by service providers

Drawback

- VLAN per site

BCMSN v2.1—10-15

The figure shows an example of a DVS deployment. The headquarters office is maintaining connectivity with its remote locations by way of selected VLAN IDs.

The VLAN identifier is used to direct the traffic to the appropriate remote locations. In the hub-and-spoke topology, VLAN10 spans the regional headquarters and remote office 1 and 3, while VLAN20 spans the regional headquarters and remote offices 2 and 3.

Metro Ethernet over SONET

Metro Ethernet over SONET takes advantage of the management and fault-tolerance mechanisms of SONET, which is already widely deployed by service providers. This topic discusses Metro Ethernet over SONET.

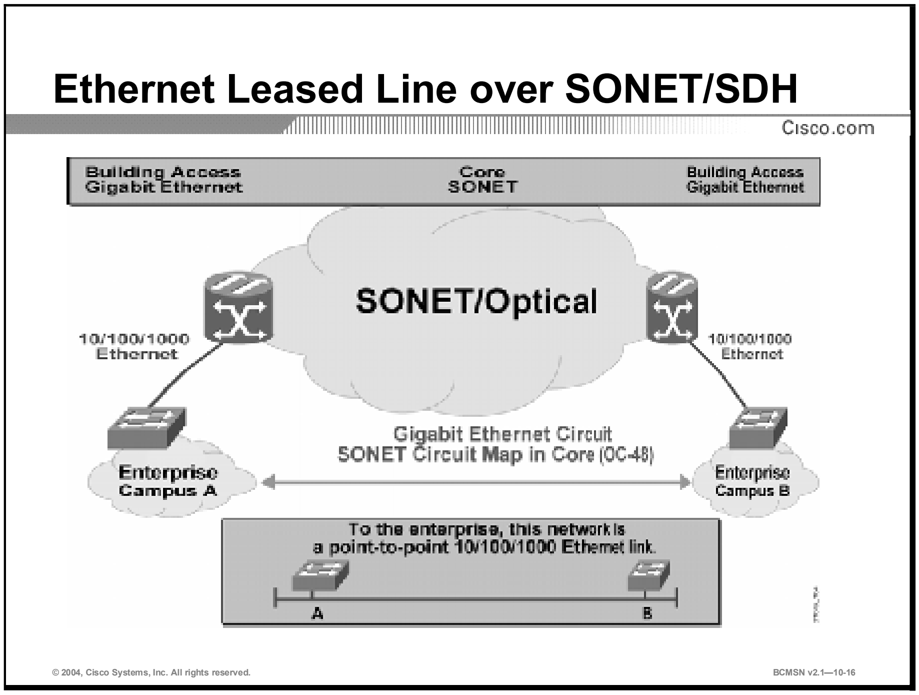

Ethernet Leased Line over SONET/SDH

BCMSN v2.1—10-16

In the figure, the SONET Metro Ethernet network appears to the end user as a point-to-point Gigabit Ethernet link. The Metro Ethernet system has essentially created a long extension cord for the Gigabit Ethernet transmission between enterprise campus A and enterprise campus B.

SONET implementations are established and entrenched in service provider networks today. SONET is used to offer Metro Ethernet solutions that leverage the existing infrastructure including the high-availability features.

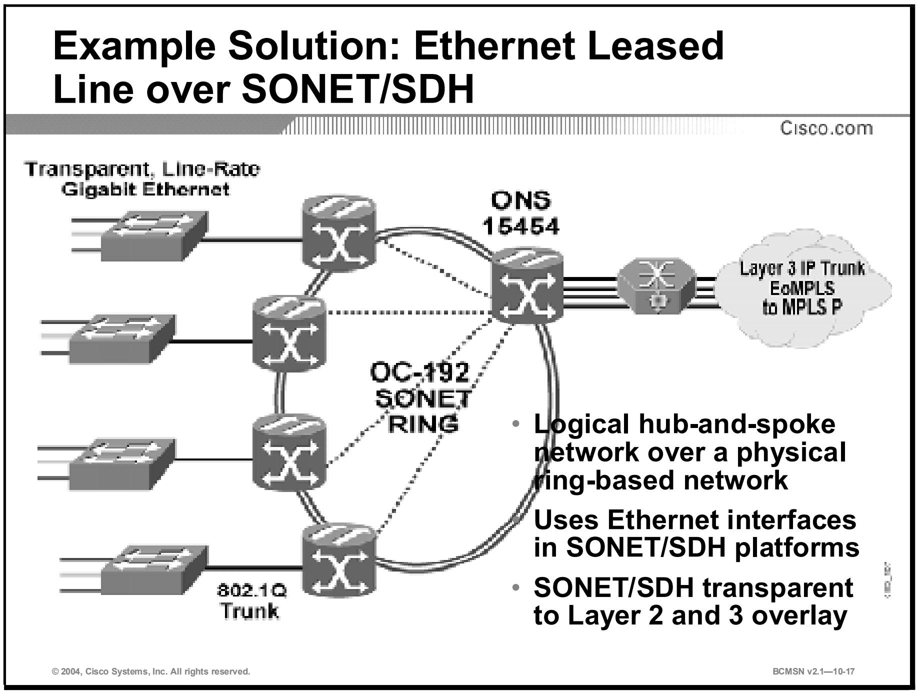

Example Solution: Ethernet Leased Line over SONET/SDH

Cisco.com

Transparent, Line-Rate Gigabit Ethernet
ONS 15454
Layer 3 IP Trunk EoMPLS to MPLS P
OC-192 SONET RING
802.1Q Trunk

- Logical hub-and-spoke network over a physical ring-based network
- Uses Ethernet interfaces in SONET/SDH platforms
- SONET/SDH transparent to Layer 2 and 3 overlay

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—10-17

The figure demonstrates how to configure a Metro Ethernet solution using a SONET ring infrastructure. The Cisco ONS 15454 provides support for both Ethernet and SONET/SDH interfaces, allowing enterprises to build Metro-to-Metro or city-to-city Gigabit Metro Ethernet networks.

The ONS 15454 supports port-based VLANs that group ports into virtual workgroups. These port-based VLANs ensure that Ethernet ports see only traffic from ports within the same VLAN. The different VLANs communicate with each other through a router that spans multiple VLANs by using router ports in each VLAN, or through a router that understands IEEE 802.1Q.

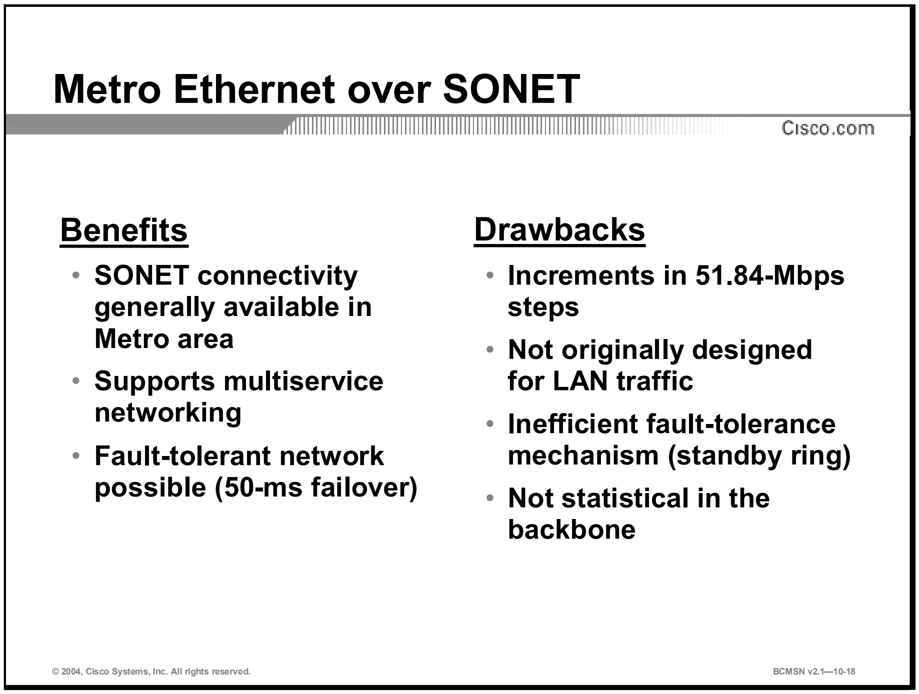

Metro Ethernet over SONET

Benefits

- **SONET connectivity generally available in Metro area**
- **Supports multiservice networking**
- **Fault-tolerant network possible (50-ms failover)**

Drawbacks

- **Increments in 51.84-Mbps steps**
- **Not originally designed for LAN traffic**
- **Inefficient fault-tolerance mechanism (standby ring)**
- **Not statistical in the backbone**

BCMSN v2.1—10-18

A multiservice SONET network is available in most Metro areas and includes a fault-tolerant, 50- ms failover mechanism through protection switching. A Metro Ethernet solution running on a SONET infrastructure benefits from the high degree of management and alarms built into the SONET protocol. These features do not exist in protocols such as Ethernet.

Some disadvantages of using SONET for Metro Ethernet solutions relate to the granularity of bandwidth increments. SONET circuits will typically force increments in steps of at least 51.84 Mbps. Metro Ethernet providers may prefer offering finer granularity such as 1 Mbps, 10 Mbps, or 50 Mbps.

SONET was not originally designed to carry Ethernet-like traffic. SONET was designed to support low-speed voice channels. Therefore, the core equipment can become very costly because most of it is designed to access and switch DS0 voice channels instead of gigabit data channels.

SONET does support a high-availability fault-tolerance mechanism, but this mechanism is sometimes inefficient because the protection circuit is typically idle.

SONET also provides a leased-line type of circuit, which transfers Ethernet frames as bits across the WAN. Unlike ATM, SONET is not able to statistically multiplex traffic across the backbone to increase utilization of this circuit.

Metro Ethernet over DWDM

Metro Ethernet over DWDM is an implementation option that provides gigabit rates with ease of configuration, high scalability, transparency, and optical protection. This topic describes Metro Ethernet over DWDM.

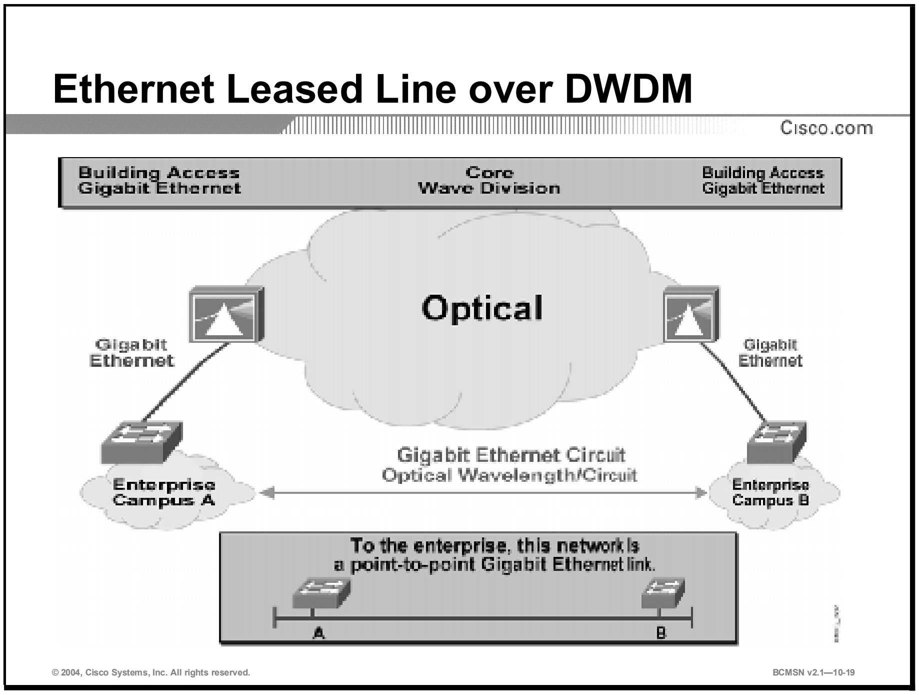

Ethernet Leased Line over DWDM

To the enterprise, this network is a point-to-point Gigabit Ethernet link.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—10-19

The advantages of WDM technology have long been recognized in the long-distance, ultrahigh-bandwidth transport market. In these environments, the laying of additional fibers is an extremely expensive and time-consuming process, leaving WDM-based solutions as the only real answer to the fast growth in bandwidth demand.

In the figure, the DWDM Metro Ethernet network appears to the end user as a point-to-point Gigabit Ethernet link. The Metro Ethernet system has essentially created a long extension cord for the Gigabit Ethernet transmission between enterprise campus A and enterprise campus B.

Ethernet packets are converted bit by bit into an optical stream. Ethernet over DWDM does not require any specific encapsulation. The core network moves data bits transparently from campus to campus.

The two most commonly used wavelength gaps are 1310 and 1550 nm. Transmissions using the 1310-nm wavelength are popular for metropolitan or short-range applications because of their relatively low cost. Because transmissions in the 1550-nm range can travel farther, they are preferred for long-distance transmissions. The increased cost in equipment to process wavelengths in the 1550-nm range is more than outweighed by the reduced necessity to retransmit the signal as it decreases in intensity.

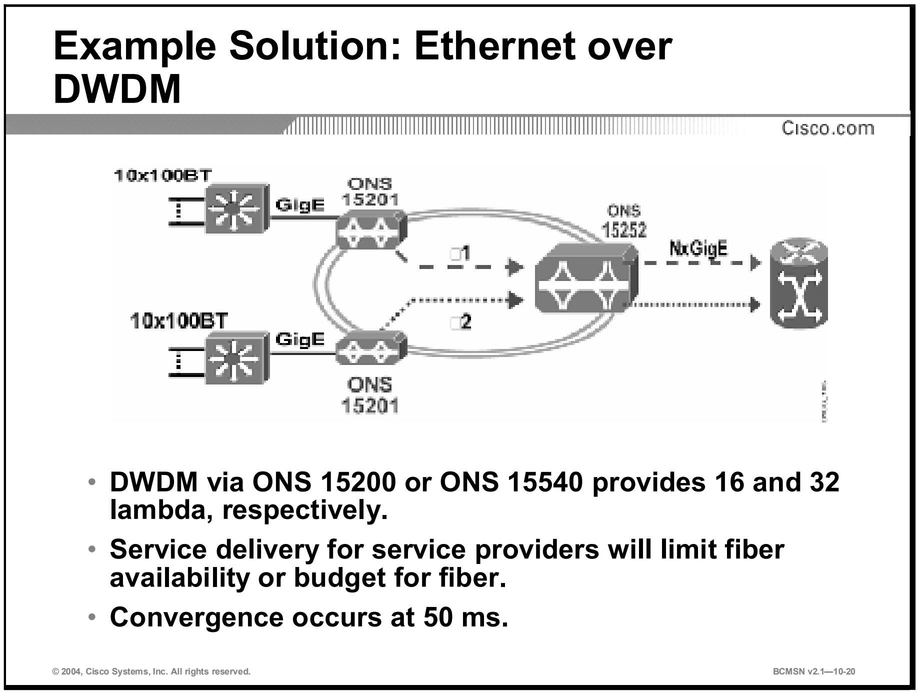

Example Solution: Ethernet over DWDM

Cisco.com

- DWDM via ONS 15200 or ONS 15540 provides 16 and 32 lambda, respectively.
- Service delivery for service providers will limit fiber availability or budget for fiber.
- Convergence occurs at 50 ms.

BCMSN v2.1—10-20

The figure shows an example of an Ethernet-over-DWDM solution using the Cisco ONS 15252. Some of the advantages are as follows:

- Low initial installation cost to provide service to buildings with a low subscriber count (1–10 subnets per building)
- Reduced cost of fiber protection by using fewer Gigabit Ethernet ports on the core network switch
- Good optical characteristics, with more buildings per ring or greater ring circumferences, or both
- Ease of installation and test

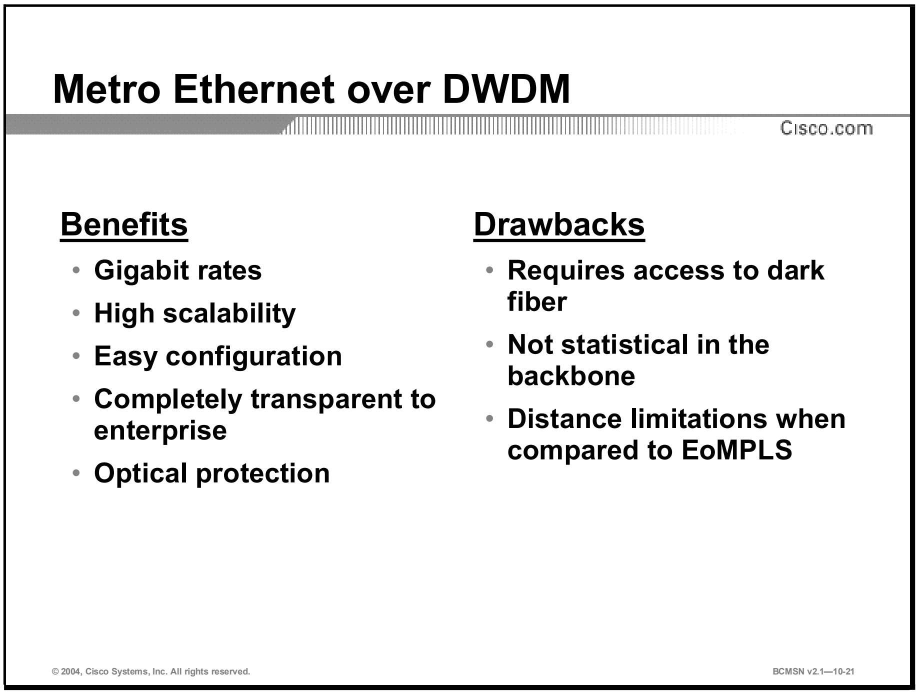

Metro Ethernet over DWDM

Cisco.com

Benefits

- Gigabit rates
- High scalability
- Easy configuration
- Completely transparent to enterprise
- Optical protection

Drawbacks

- Requires access to dark fiber
- Not statistical in the backbone
- Distance limitations when compared to EoMPLS

BCMSN v2.1—10-21

Metro Ethernet over DWDM benefits from the gigabit speeds of DWDM networks. It is totally transparent to the end user, as the bits-in/bits-out service behaves very much like a leased-line or private-line service. No statistical multiplexing is applied to traffic across the core.

Dark fiber refers to unused fiber-optic cable. Sometimes, service providers lay more lines than what they need to limit the costs of having to lay lines repeatedly. Service providers then leave the dark fiber strands to individuals or other companies who want to establish optical connections among their own locations. The fiber is neither controlled by nor connected to the telephone company. Instead, the company or individual provides the necessary components to make it functional.

Metro Ethernet over CWDM

Metro Ethernet over coarse wavelength division multiplexing (CWDM) is a cost-effective implementation choice for a small geographic area. This topic describes Metro Ethernet over CWDM.

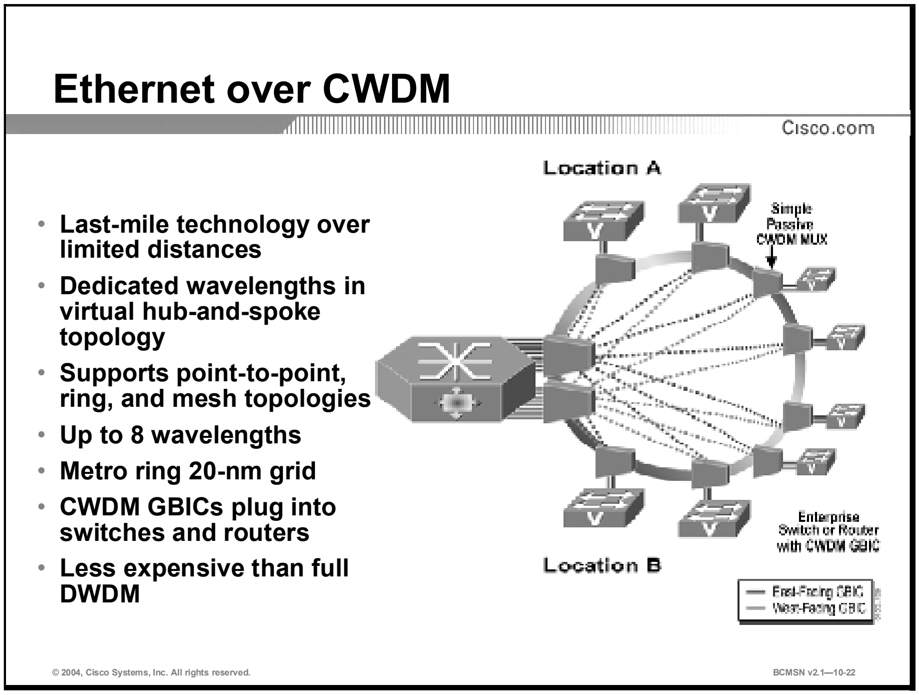

Ethernet over CWDM

Cisco.com

- Last-mile technology over limited distances
- Dedicated wavelengths in virtual hub-and-spoke topology
- Supports point-to-point, ring, and mesh topologies
- Up to 8 wavelengths
- Metro ring 20-nm grid
- CWDM GBICs plug into switches and routers
- Less expensive than full DWDM

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—10-22

CWDM is a last-mile technology because of its limited operations distance. It can support up to eight lambdas and is Gigabit Ethernet-capable. CWDM offers these characteristics:

- Similar to DWDM but at coarser wavelengths (1470 nm, 1490 nm, 1510 nm, 1530 nm, 1550 nm, 1570 nm, 1590 nm, 1610 nm)
- Gigabit Interface Converter (GBIC) form factor
- Single technology to Gigabit Ethernet
- Recommended for Fibre Channel, Enterprise System Connection (ESCON), and so on
- Very cost-effective
- Long drive distance
- Given a 30-dB budget, equals approximately 100 kilometers
- Cannot be optically amplified

A service provider may offer CWDM as a lambda service.

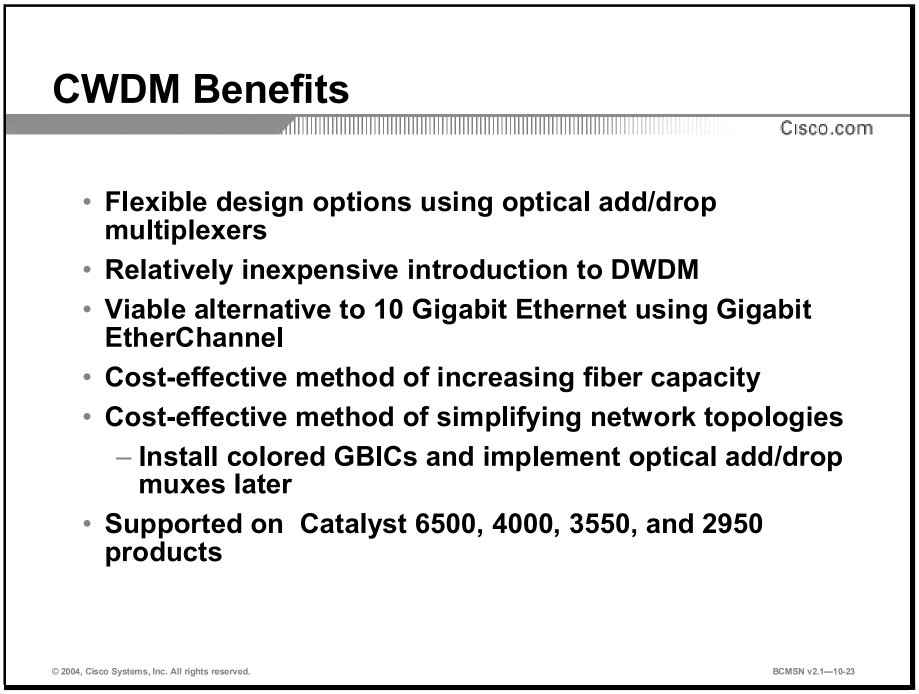

CWDM Benefits

- Flexible design options using optical add/drop multiplexers
- Relatively inexpensive introduction to DWDM
- Viable alternative to 10 Gigabit Ethernet using Gigabit EtherChannel
- Cost-effective method of increasing fiber capacity
- Cost-effective method of simplifying network topologies
  - Install colored GBICs and implement optical add/drop muxes later
- Supported on Catalyst 6500, 4000, 3550, and 2950 products

BCMSN v2.1—10-23

CWDM is a low-cost, access-type technology to be utilized over dark fiber. CWDM cannot be amplified, so it has distance limitations compared to standard DWDM. CWDM is quite effective and economical in small geographic areas.

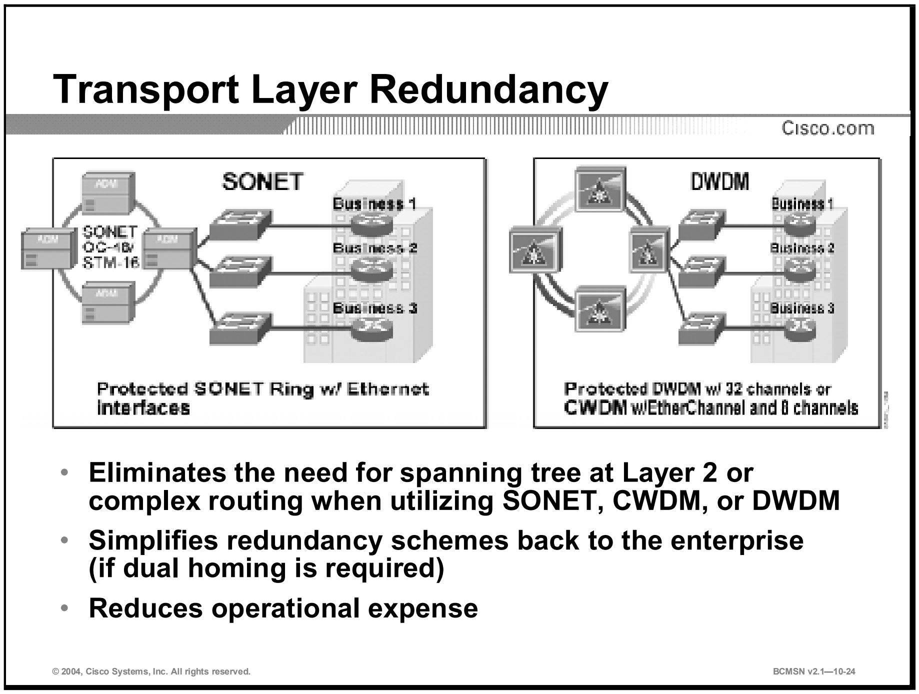

Transport Layer Redundancy

Cisco.com

Protected SONET Ring w/ Ethernet interfaces

Protected DWDM w/ 32 channels or CWDM w/EtherChannel and 8 channels

- **Eliminates the need for spanning tree at Layer 2 or complex routing when utilizing SONET, CWDM, or DWDM**
- **Simplifies redundancy schemes back to the enterprise (if dual homing is required)**
- **Reduces operational expense**

BCMSN v2.1—10-24

The Spanning Tree Protocol (STP) is run across Layer 2 Ethernet networks to allow alternative redundant routes without the occurrence of loops. Complex routing protocols can provide the same function at Layer 3. Enterprises can avoid STP and complex routing across the Metro WAN by leveraging the Layer 1 redundancy mechanisms such as protection switching, which may already be implemented in the SONET network or protected DWDM networks.

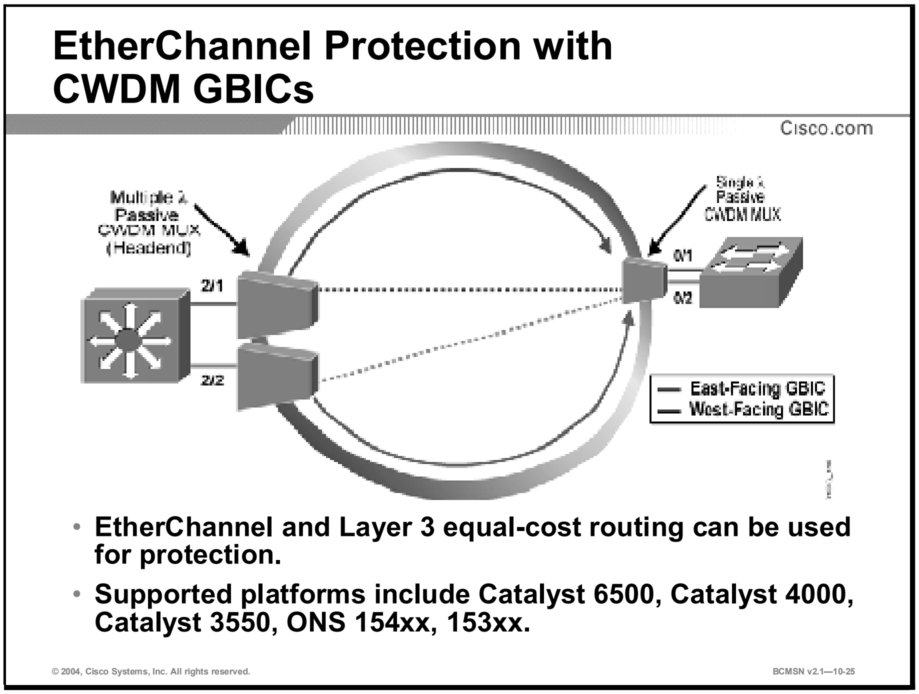

EtherChannel Protection with CWDM GBICs

- **EtherChannel and Layer 3 equal-cost routing can be used for protection.**
- **Supported platforms include Catalyst 6500, Catalyst 4000, Catalyst 3550, ONS 154xx, 153xx.**

BCMSN v2.1—10-25

A single wavelength can carry multiple channels in the form of an Ethernet trunk channel, or enterprises can use multiple wavelengths together as an EtherChannel aggregation channel. Both approaches provide high-bandwidth Layer 1 redundancy schemes.

Summary

This topic summarizes the key points discussed in this lesson.

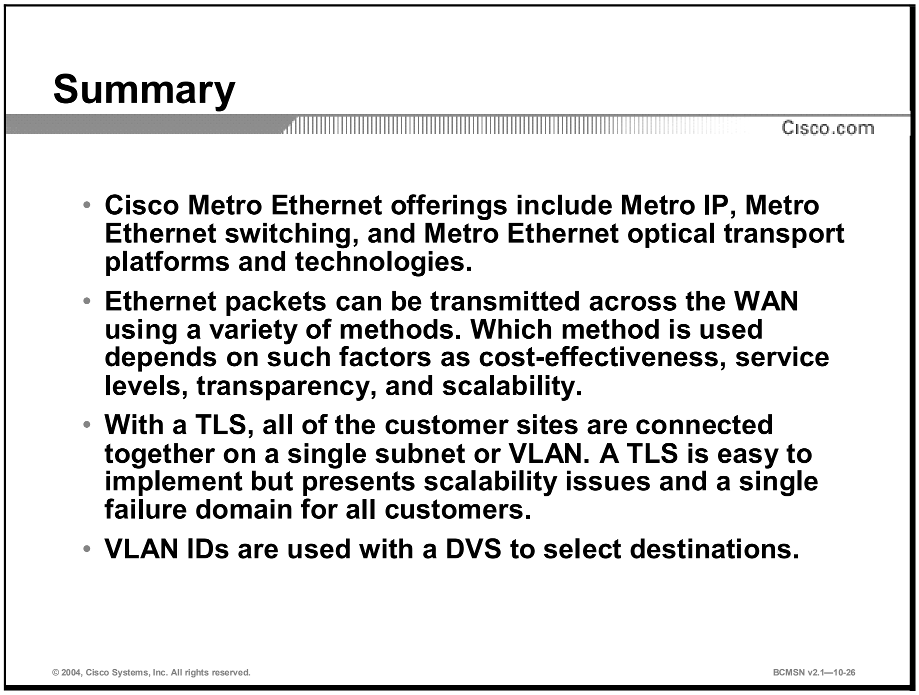

Summary

- Cisco Metro Ethernet offerings include Metro IP, Metro Ethernet switching, and Metro Ethernet optical transport platforms and technologies.
- Ethernet packets can be transmitted across the WAN using a variety of methods. Which method is used depends on such factors as cost-effectiveness, service levels, transparency, and scalability.
- With a TLS, all of the customer sites are connected together on a single subnet or VLAN. A TLS is easy to implement but presents scalability issues and a single failure domain for all customers.
- VLAN IDs are used with a DVS to select destinations.

BCMSN v2.1—10-26

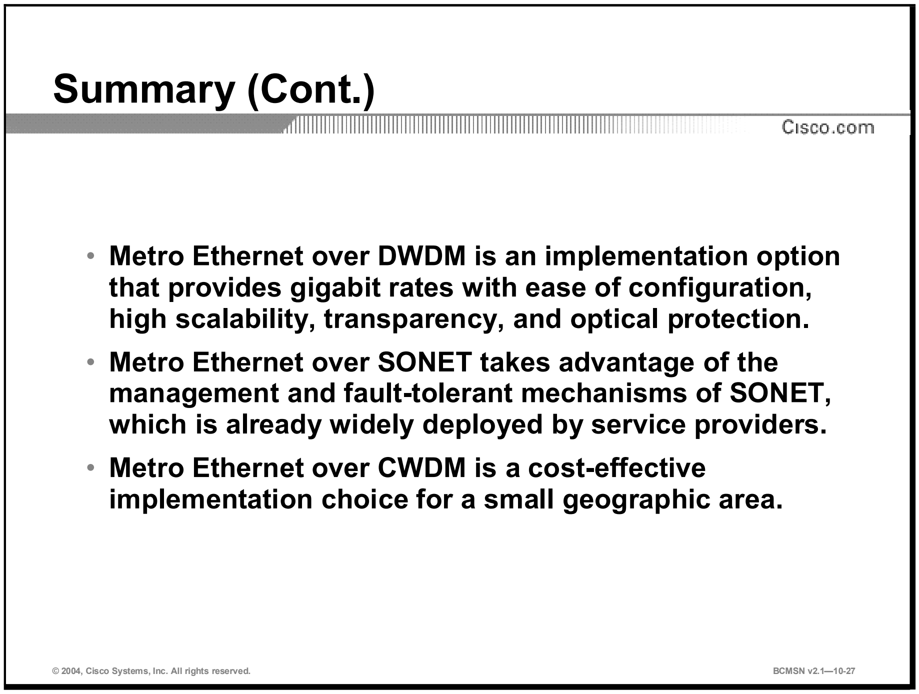

Summary (Cont.)

- Metro Ethernet over DWDM is an implementation option that provides gigabit rates with ease of configuration, high scalability, transparency, and optical protection.
- Metro Ethernet over SONET takes advantage of the management and fault-tolerant mechanisms of SONET, which is already widely deployed by service providers.
- Metro Ethernet over CWDM is a cost-effective implementation choice for a small geographic area.

BCMSN v2.1—10-27

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which Cisco Catalyst switch product provides three queues per port with priority queuing?

A) Catalyst 3550

B) Catalyst 4000

C) Catalyst 6500

D) Catalyst 7600

Q2) Which selection criterion for transmitting Ethernet over the WAN allows for the elimination of Layer 3 functionality on the CPE?

A) scalability

B) service level

C) transparency

D) point-to-point versus point-to-multipoint

Q3) What is the primary benefit of a TLS?

A) scalability

B) ease of implementation

C) single broadcast domain

D) eliminates single failure domain

Q4) Which spanning tree implementation does a DVS require?

A) Rapid Spanning Tree Protocol

B) Multiple Spanning Tree

C) Per-VLAN Spanning Tree+

D) Per-Network Spanning Tree

Q5) Which drawback does Metro Ethernet over SONET share with Metro Ethernet over DWDM?

A) not designed for LAN traffic

B) lack of bandwidth granularity

C) no backbone statistical multiplexing

D) inefficient fault-tolerance mechanism

Q6) What are two advantages of Metro Ethernet over DWDM? (Choose two.)

A) transparency

B) megabit rates

C) distance capabilities

D) ease of configuration

E) statistical multiplexing provisions

Q7) What is one drawback of Metro Ethernet over CWDM compared to Metro Ethernet over DWDM?

A) increased cost

B) distance limitations

C) bandwidth limitations

D) decreased fiber capacity

Quiz Answer Key

Q1) C

Relates to: Metro Ethernet Solutions

Q2) D

Relates to: Metro Ethernet Connectivity

Q3) B

Relates to: Transparent LAN Service

Q4) C

Relates to: Directed VLAN Service

Q5) C

Relates to: Metro Ethernet over SONET

Q6) A, D

Relates to: Metro Ethernet over DWDM

Q7) B

Relates to: Metro Ethernet over CWDM

Examining Metro Ethernet Tunneling

Overview

With a Metro Ethernet implementation, you can encapsulate or unencapsulate the traffic that crosses the service provider core with a Layer 2 or Layer 3 mechanism. The Layer 2 mechanism is 802.1Q-in-Q tunneling, which provides traffic isolation of customer traffic from other customer traffic and the network core. Metro Ethernet Layer 3 transport options include Ethernet over MPLS (EoMPLS) and EoMPLS encapsulation point-to-multipoint. EoMPLS uses an MPLS service provider core network.

Relevance

Both network administrators and network designers need to know what consequences the various encapsulation and transport mechanisms have on network functionality.

Objectives

Upon completing this lesson, you will be able to:

- Describe the benefits and disadvantages of transporting traffic across the service provider network without tunneling
- Describe the benefits and disadvantages of 802.1Q-in-Q tunneling
- Describe Ethernet over MPLS
- List the benefits and disadvantages of EoMPLS as a Layer 3 transport option for Metro Ethernet
- Describe EoMPLS encapsulation point-to-multipoint

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of *Interconnecting Cisco Network Devices* (ICND)

Outline

This lesson includes these topics:

- Overview
- Metro Ethernet Tunneling Options
- Tag Stacking (Q-in-Q Tunneling)
- Metro Ethernet: EoMPLS Encapsulation
- EoMPLS Characteristics
- Metro Ethernet: EoMPLS Encapsulation Point-to-Multipoint
- Summary
- Quiz

Metro Ethernet Tunneling Options

Traffic crossing the service provider core can be transported without Layer 2 or Layer 3 tunneling. Transport without tunneling is easy to implement but does not scale and provides no isolation of enterprise traffic. This topic discusses the Metro Ethernet tunneling options.

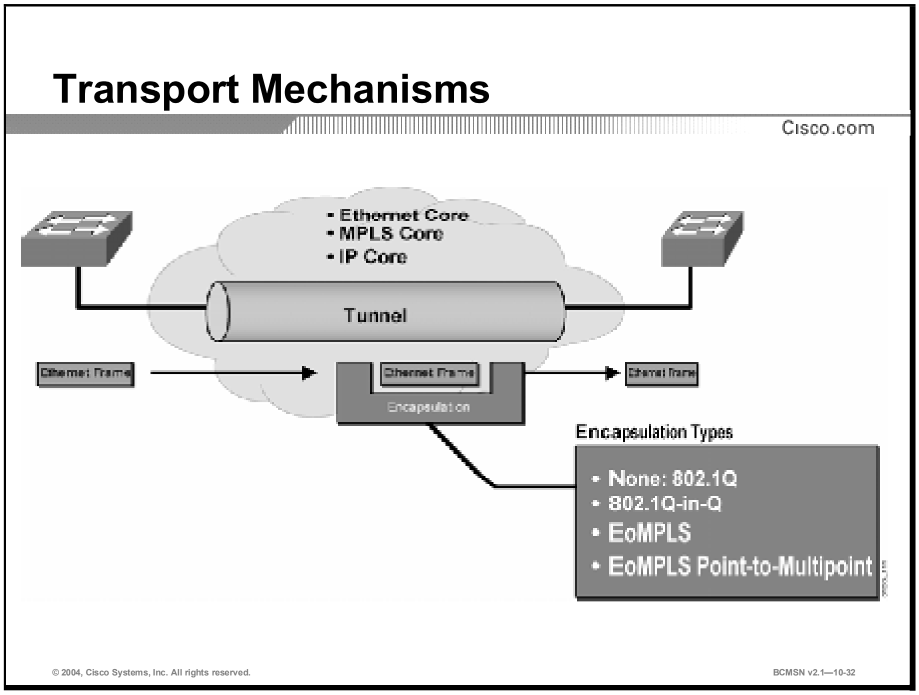

Transport Mechanisms

Cisco.com

- Ethernet Core
- MPLS Core
- IP Core

Encapsulation Types

- None: 802.1Q
- 802.1Q-in-Q
- EoMPLS
- EoMPLS Point-to-Multipoint

BCMSN v2.1—10-32

A service provider network can support these possible encapsulation solutions:

- 802.1Q (no tunneling)
- 802.1Q-in-Q (Layer 2 tunneling)
- EoMPLS (Layer 3 tunneling)
- EoMPLS encapsulation point-to-multipoint (Layer 3 tunneling)

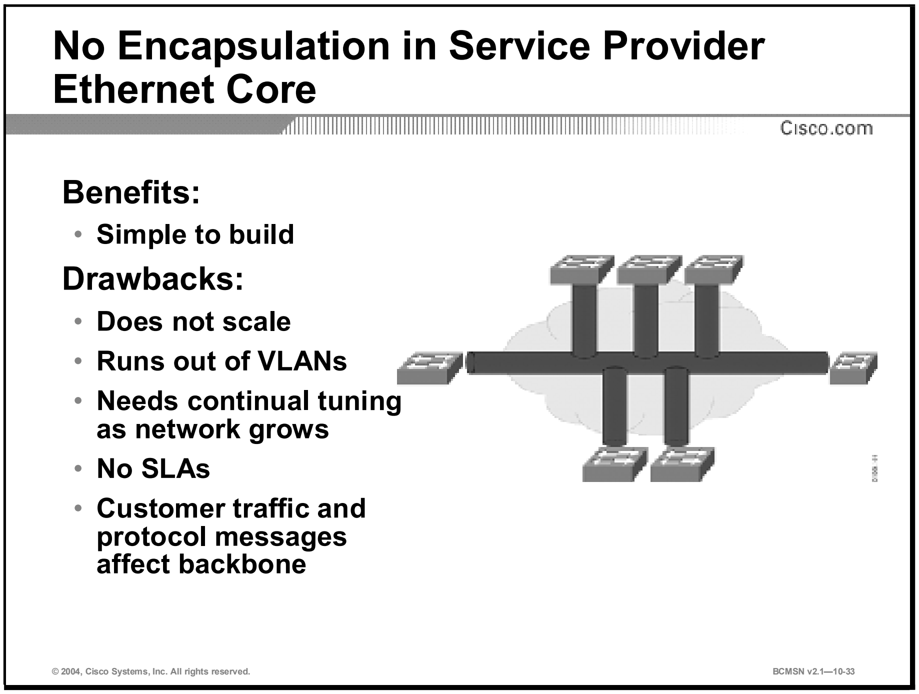

No Encapsulation in Service Provider Ethernet Core

Benefits:

- **Simple to build**

Drawbacks:

- **Does not scale**
- **Runs out of VLANs**
- **Needs continual tuning as network grows**
- **No SLAs**
- **Customer traffic and protocol messages affect backbone**

BCMSN v2.1—10-33

With no encapsulation, user traffic will affect the operation of the core network. This type of networking is simple to build and low in cost. Its connectivity type is the equivalent of a TLS. It might be effective to support the network of a single enterprise. Service providers might use such a low-cost entry point to deliver Ethernet services by using an existing infrastructure, such as with SONET.

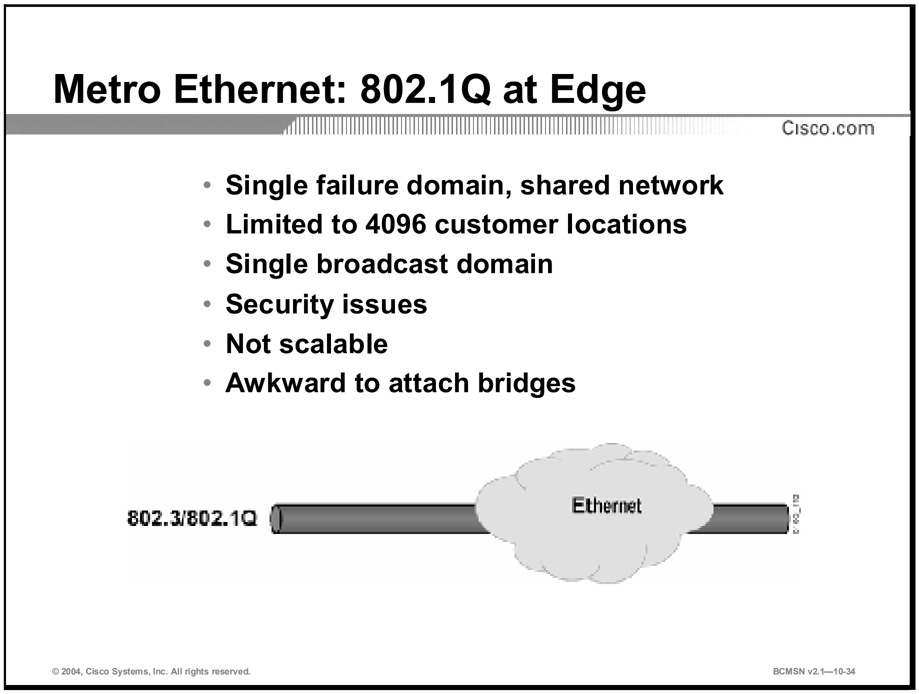

Metro Ethernet: 802.1Q at Edge

- **Single failure domain, shared network**
- **Limited to 4096 customer locations**
- **Single broadcast domain**
- **Security issues**
- **Not scalable**
- **Awkward to attach bridges**

802.3/802.1Q
Ethernet

BCMSN v2.1—10-34

The main issue associated with no encapsulation is congestion in the core from customer traffic. The single failure domain associated with this design is not acceptable in most cases.

The core network design is essentially an Ethernet switch. This method of construction is quite simple; however, it has a number of negative attributes. One of the biggest problems with this implementation is the fact that customer traffic is intermixed and can affect the operation of the backbone network. For example, a customer with a jabbering port could negatively affect other customers on the network.

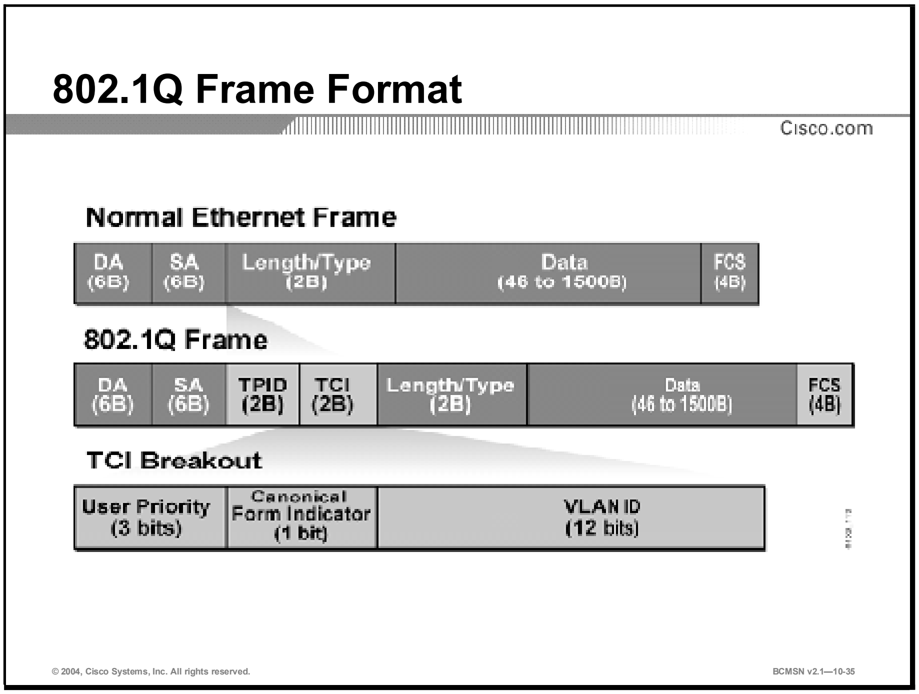

802.1Q Frame Format

Cisco.com

Normal Ethernet Frame

DA (6B)	SA (6B)	Length/Type (2B)	Data (46 to 1500B)	FCS (4B)

802.1Q Frame

DA (6B)	SA (6B)	TPID (2B)	TCI (2B)	Length/Type (2B)	Data (46 to 1500B)	FCS (4B)

TCI Breakout

User Priority (3 bits)	Canonical Form Indicator (1 bit)	VLAN ID (12 bits)

BCMSN v2.1—10-35

The IEEE 802.1Q standard specifies much more than encapsulation types, including spanning tree enhancements, General Attribute Registration Protocol (GARP), and 802.1p QoS tagging.

The 802.1Q frame format preserves the original Ethernet source address and destination address, but switches must now expect to receive baby giant frames, even on access ports where hosts may use tagging to express 802.1p user priority for QoS signaling. The tag is 4 bytes, so 802.1Q Ethernet version 2 frames are 1522 bytes. Finally, 802.1Q supports a numbering space for 4096 VLANs.

All data frames that are transmitted and received are 802.1Q-tagged, except for those on the native VLAN (there is an implicit tag based on the ingress switch port configuration). Frames on the native VLAN are always transmitted untagged and are normally received untagged, but may also be received tagged.

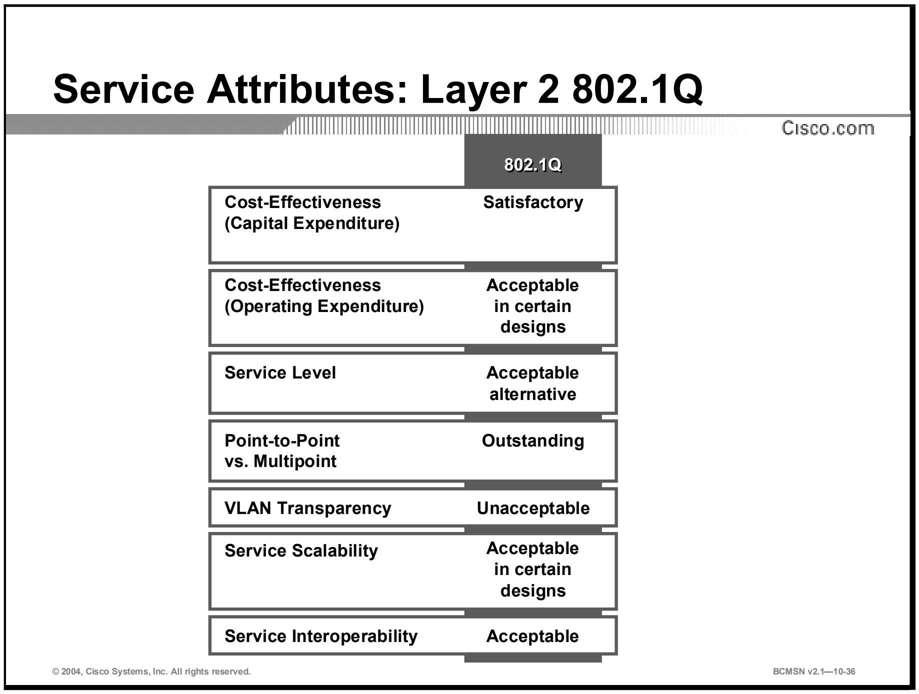

Service Attributes: Layer 2 802.1Q

Cisco.com

	802.1Q
Cost-Effectiveness (Capital Expenditure)	Satisfactory
Cost-Effectiveness (Operating Expenditure)	Acceptable in certain designs
Service Level	Acceptable alternative
Point-to-Point vs. Multipoint	Outstanding
VLAN Transparency	Unacceptable
Service Scalability	Acceptable in certain designs
Service Interoperability	Acceptable

BCMSN v2.1—10-36

The key highlights for 802.1Q are its low cost and point-to-multipoint operation.

The capital expenditures and operating expenditures associated with 802.1Q are acceptable and its point-to-point and multipoint operation offers significant benefits. The main drawback is that the enterprises are sharing a core switch and thus there is no isolation between the traffic of the enterprise and the traffic of network operations. The operations of one enterprise could affect other enterprises as well as the core network itself.

Tag Stacking (Q-in-Q Tunneling)

The Layer 2 encapsulation option for Metro Ethernet is 802.1Q-in-Q tunneling (tag stacking). Tag stacking provides isolation of enterprise traffic through the service provider core by adding an additional tag to enterprise traffic, isolating the traffic of each enterprise to a service provider VLAN. This topic discusses tag stacking.

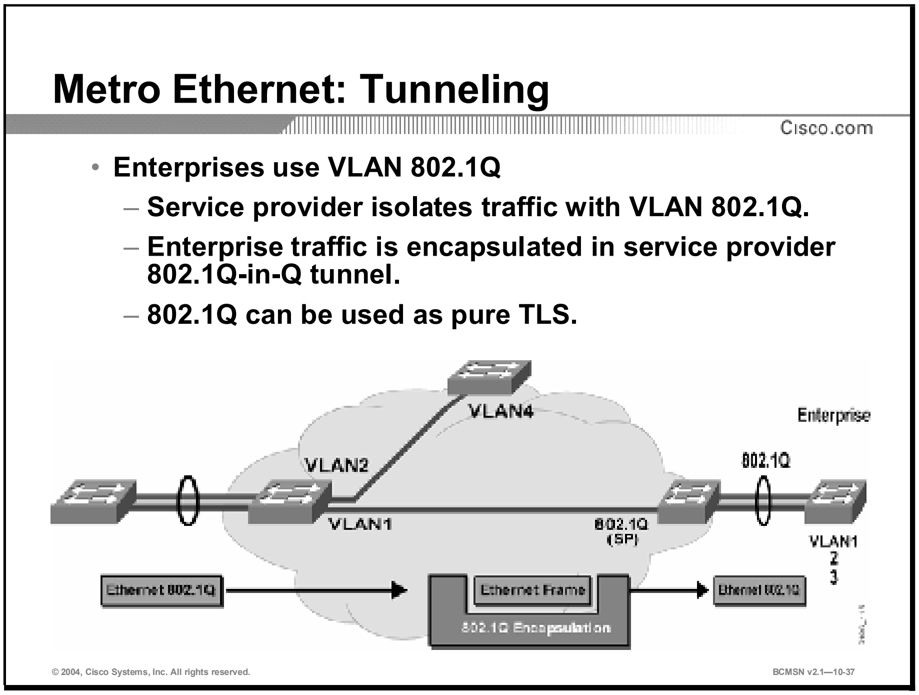

Metro Ethernet: Tunneling

Cisco.com

- Enterprises use VLAN 802.1Q
  - Service provider isolates traffic with VLAN 802.1Q.
  - Enterprise traffic is encapsulated in service provider 802.1Q-in-Q tunnel.
  - 802.1Q can be used as pure TLS.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—10-37

The figure shows an 802.1Q-in-Q tunneling implementation. The Ethernet frame of the enterprise (802.1Q) is encapsulated in the 802.1Q frame of the service provider. This encapsulation provides VLAN isolation for the service provider while allowing the enterprise to utilize the VLAN service.

Tag stacking provides a way to isolate the traffic of one enterprise from other enterprise traffic and the core network. In standard 802.1Q operation, VLAN IDs of the enterprise cannot overlap. With 802.1Q-in-Q tunneling, they can overlap.

One problem with tag stacking is that the unique reference to MAC addresses is lost across the core.

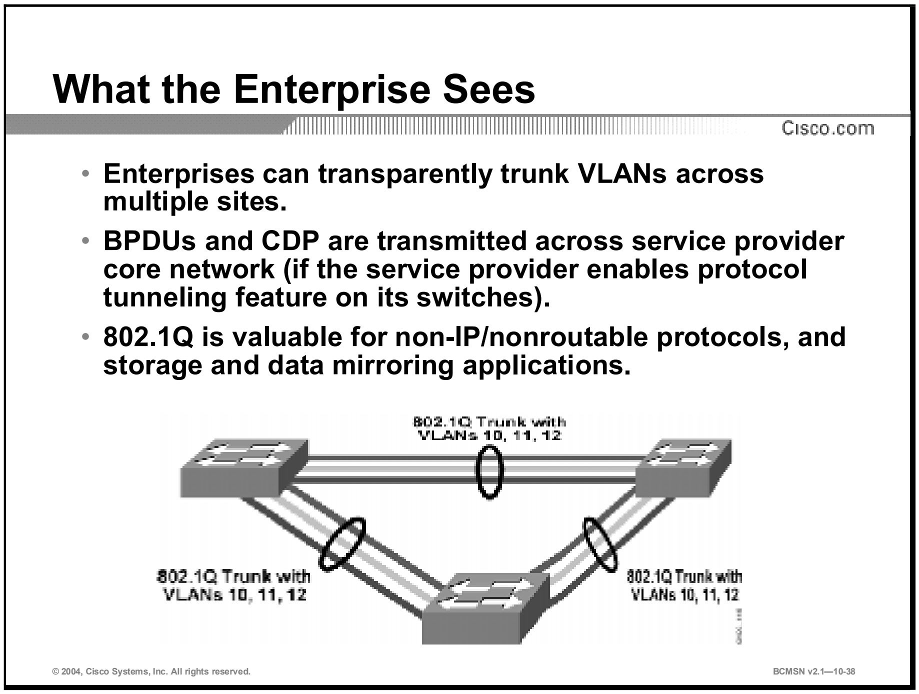

What the Enterprise Sees

Cisco.com

- Enterprises can transparently trunk VLANs across multiple sites.
- BPDUs and CDP are transmitted across service provider core network (if the service provider enables protocol tunneling feature on its switches).
- 802.1Q is valuable for non-IP/nonroutable protocols, and storage and data mirroring applications.

802.1Q Trunk with VLANs 10, 11, 12
802.1Q Trunk with VLANs 10, 11, 12
802.1Q Trunk with VLANs 10, 11, 12

BCMSN v2.1—10-38

The enterprise can transparently trunk VLANs across the core networks. Enterprise bridge protocol data units (BPDUs) will not affect the network 802.1Q environment.

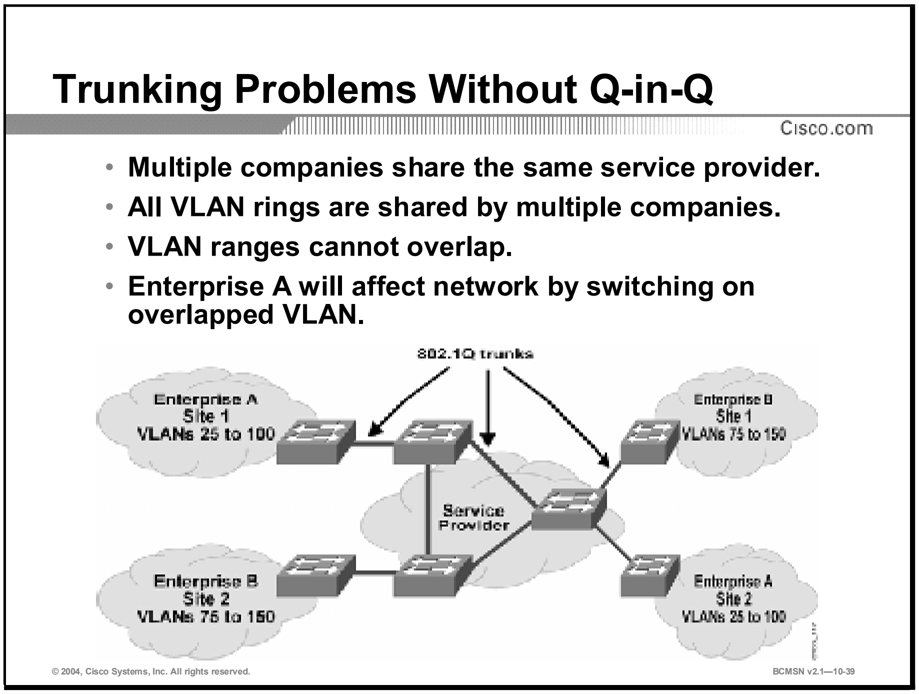

Trunking Problems Without Q-in-Q

Cisco.com

- Multiple companies share the same service provider.
- All VLAN rings are shared by multiple companies.
- VLAN ranges cannot overlap.
- Enterprise A will affect network by switching on overlapped VLAN.

BCMSN v2.1—10-39

Because VLAN identifiers may overlap from location to location, 802.1Q-in-Q tunneling provides a way to identify these VLANs as they relate to their respective sites.

In the figure, enterprise A and enterprise B have overlapping VLAN identifiers. Without 802.1Q-in-Q tunneling, this overlap would present a conflict. With tag stacking, transparency is preserved and overlapping VLAN IDs are not a problem.

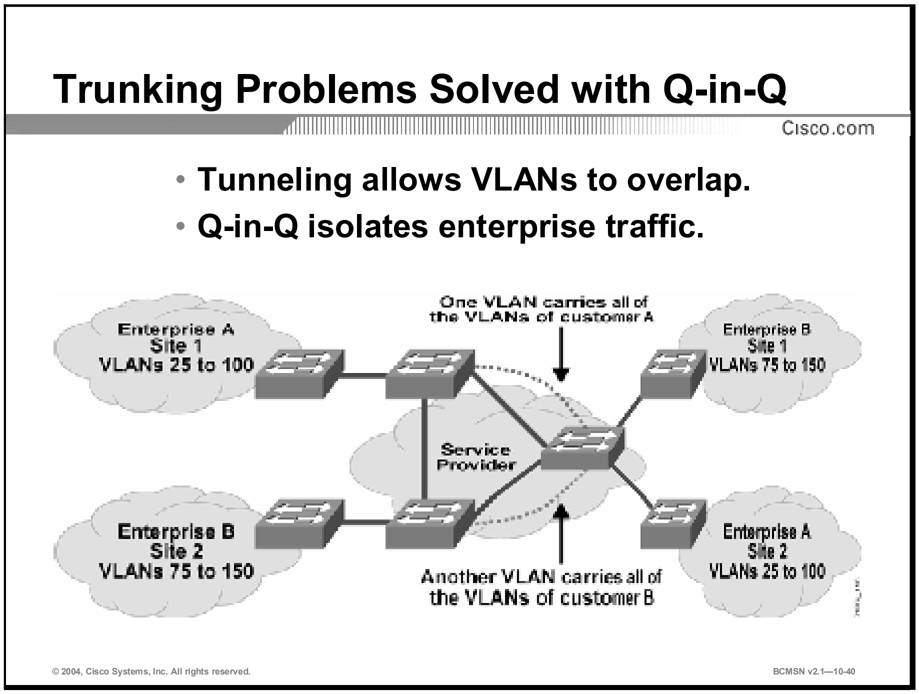

Trunking Problems Solved with Q-in-Q

Cisco.com

- Tunneling allows VLANs to overlap.
- Q-in-Q isolates enterprise traffic.

One VLAN carries all of the VLANs of customer A
Enterprise A Site 1 VLANs 25 to 100
Enterprise B Site 1 VLANs 75 to 150
Service Provider
Enterprise B Site 2 VLANs 75 to 150
Enterprise A Site 2 VLANs 25 to 100
Another VLAN carries all of the VLANs of customer B

BCMSN v2.1—10-40

VLAN identifiers become transparent when the VLANs of a specific site are tunneled inside a new VLAN specific to the service provider network. One VLAN in the service provider core carries all of the VLAN traffic of company A, while another carries the traffic of company B. Thus, the traffic of company A is isolated from traffic of company B through the core network Q-in-Q encapsulation.

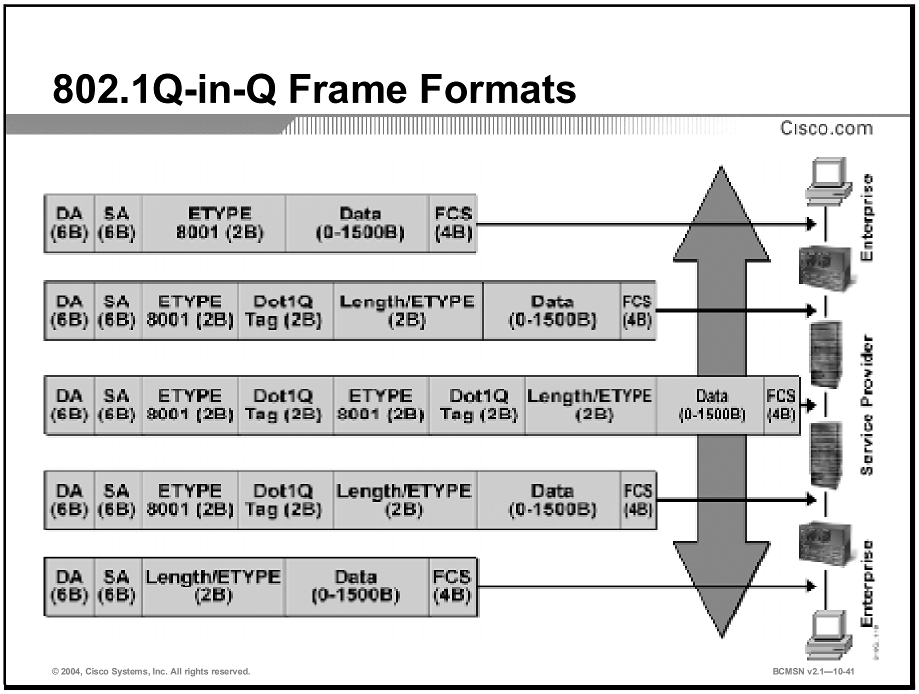

802.1Q-in-Q Frame Formats

Cisco.com

DA	SA	ETYPE	Data	FCS
(6B)	(6B)	8001 (2B)	(0-1500B)	(4B)

DA	SA	ETYPE	Dot1Q	Length/ETYPE	Data	FCS
(6B)	(6B)	8001 (2B)	Tag (2B)	(2B)	(0-1500B)	(4B)

DA	SA	ETYPE	Dot1Q	ETYPE	Dot1Q	Length/ETYPE	Data	FCS
(6B)	(6B)	8001 (2B)	Tag (2B)	8001 (2B)	Tag (2B)	(2B)	(0-1500B)	(4B)

DA	SA	ETYPE	Dot1Q	Length/ETYPE	Data	FCS
(6B)	(6B)	8001 (2B)	Tag (2B)	(2B)	(0-1500B)	(4B)

DA	SA	Length/ETYPE	Data	FCS
(6B)	(6B)	(2B)	(0-1500B)	(4B)

BCMSN v2.1—10-41

As a frame moves from one end of the link to another, dot-1Q tags are added and stripped at multiple points. A tag, identifying the customer VLAN, is added before the frame enters the service provider network, and a second tag, identifying the service provider VLAN, is added as it enters the service provider network. Thus, the frame is double-tagged as it transits the service provider network. The end result is that it contains both the service provider dot-1Q tag along with the original customer dot-1Q tag.

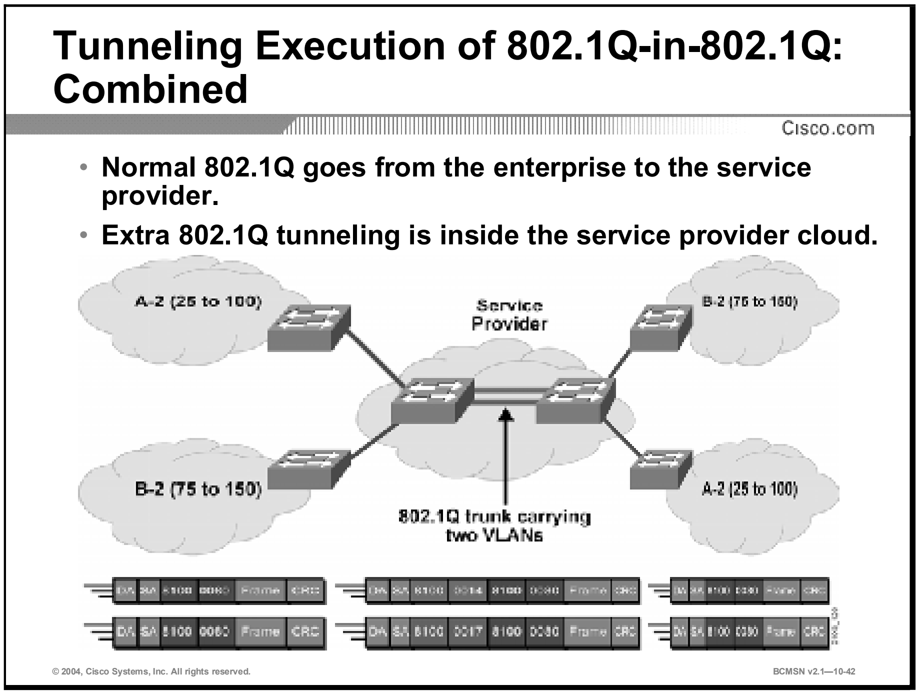

Tunneling Execution of 802.1Q-in-802.1Q: Combined

Cisco.com

- Normal 802.1Q goes from the enterprise to the service provider.
- Extra 802.1Q tunneling is inside the service provider cloud.

A-2 (25 to 100)
Service Provider
B-2 (75 to 150)
B-2 (75 to 150)
A-2 (25 to 100)
802.1Q trunk carrying two VLANs
DA SA 8100 0080 Frame CRC
DA SA 8100 0014 8100 0080 Frame CRC
DA SA 8100 0080 Frame CRC
DA SA 8100 0080 Frame CRC
DA SA 8100 0017 8100 0080 Frame CRC
DA SA 8100 0080 Frame CRC

BCMSN v2.1—10-42

The figure shows an example of tag stacking. Traffic traveling between the two sites of business A-2 is carried in the service provider core on VLAN0014, while traffic between the two sites of business B-2 is carried on the service provider VLAN0017. The traffic is isolated and forwarded to the proper site, even though the same customer VLAN IDs are used by the two customers.

The originating Ethertype fields, Tag Protocol Identifier (TPID) and Tag Control Information (TCI), are retained and an additional TPID and TCI are added to accommodate the Q-in-Q tunnel.

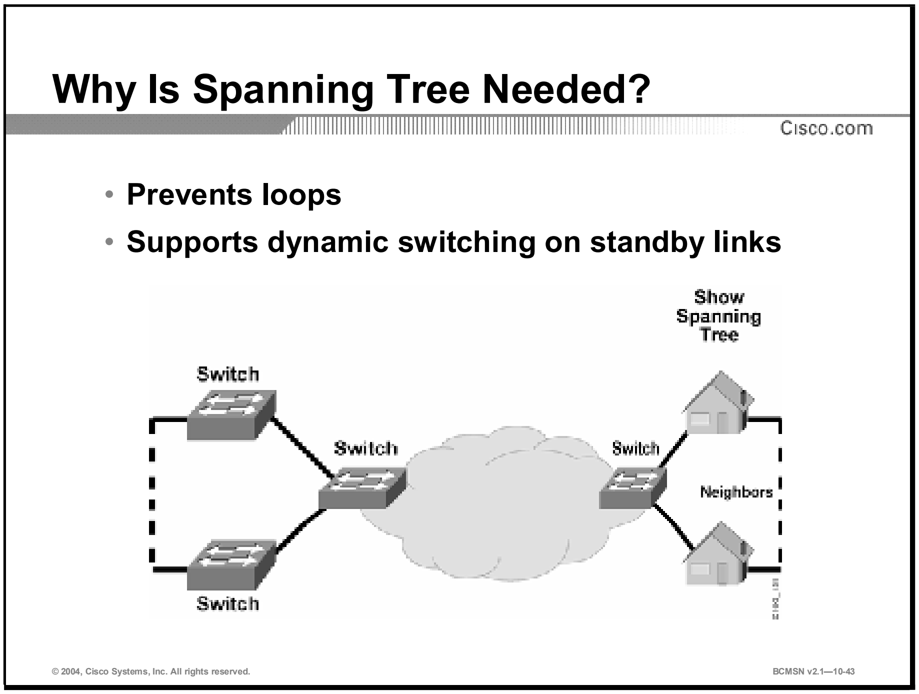

Why Is Spanning Tree Needed?

- Prevents loops
- Supports dynamic switching on standby links

BCMSN v2.1—10-43

The Metro Ethernet Q-in-Q will implement STP to prevent loops. STP supports dynamic switching for standby links. The figure shows possible conditions where looping can occur if two outputs on different switches are connected together.

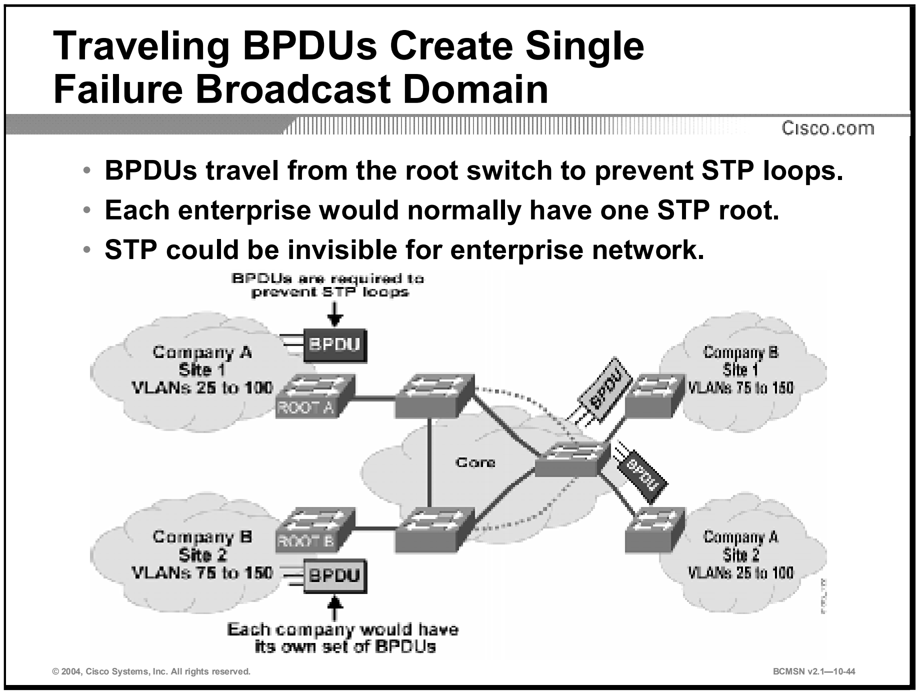

Traveling BPDUs Create Single Failure Broadcast Domain

Cisco.com

- BPDUs travel from the root switch to prevent STP loops.
- Each enterprise would normally have one STP root.
- STP could be invisible for enterprise network.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—10-44

Configuration BPDUs are sent from every port on the root bridge and subsequently flow to all leaf switches to maintain the state of the spanning tree. In steady state, BPDU flow is unidirectional: root ports and blocking ports only receive configuration BPDUs, while designated ports only send configuration BPDUs.

For every BPDU received by a switch from the root, a new BPDU is processed by the switch and sent out containing the root information. In other words, if the root bridge is lost or all paths to the root bridge are lost, then BPDUs are no longer received (until the max_age timer expires and starts re-election).

When the user switches on the Layer 2 tunnel, the customer BPDUs should travel transparently through the service provider core.

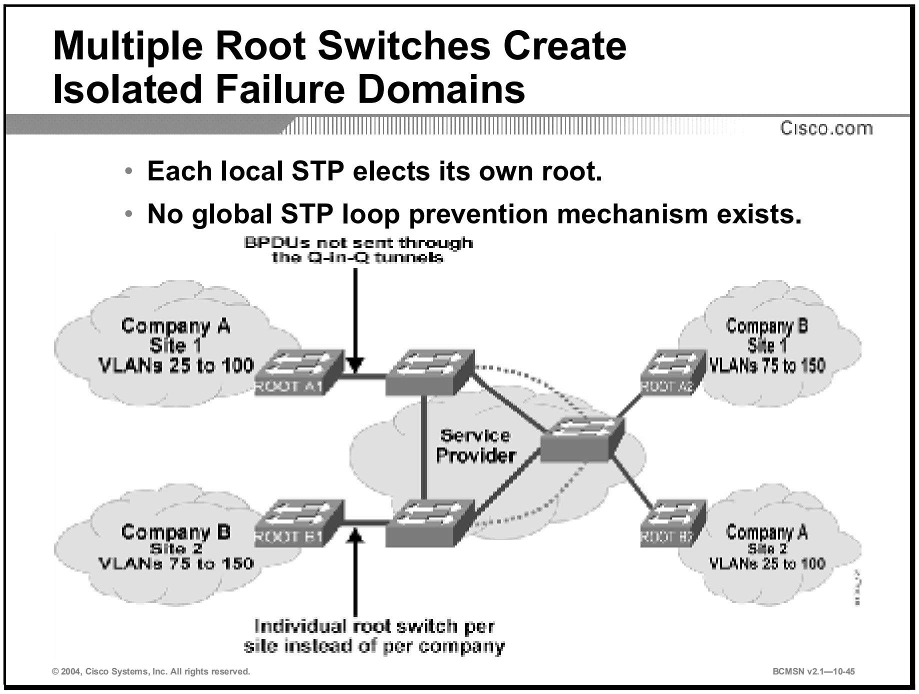

Multiple Root Switches Create Isolated Failure Domains

Cisco.com

- Each local STP elects its own root.
- No global STP loop prevention mechanism exists.

BPDUs not sent through the Q-in-Q tunnels
Company A Site 1 VLANs 25 to 100
ROOT A1
Company B Site 1 VLANs 75 to 150
ROOT A2
Service Provider
Company B Site 2 VLANs 75 to 150
ROOT B1
ROOT B2
Company A Site 2 VLANs 25 to 100
Individual root switch per site instead of per company

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—10-45

Because BPDUs are not sent through the Q-in-Q tunnels, each local spanning tree instance elects its own root. No global STP loop prevention mechanism exists.

The local STPs should elect their root to be the ingress and egress bridge for their network. Each VLAN will have its own root. The figure shows four attachments to the service provider network and, therefore, four roots.

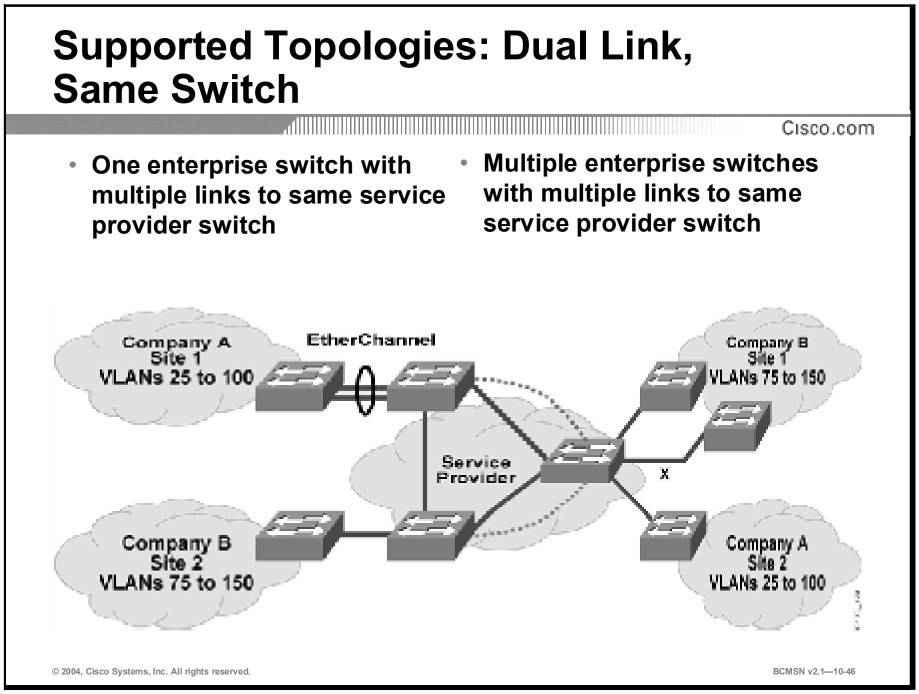

Supported Topologies: Dual Link, Same Switch

Cisco.com

- One enterprise switch with multiple links to same service provider switch
- Multiple enterprise switches with multiple links to same service provider switch

BCMSN v2.1—10-46

Supported topologies include the following:

- One customer switch with multiple links to the same service provider switch
- Multiple customer switches with multiple links to the same service provider switch

Note that in the figure, the site with two links from the same site will experience one link being shut down by STP. Use of EtherChannel will maximize available bandwidth in this scenario.

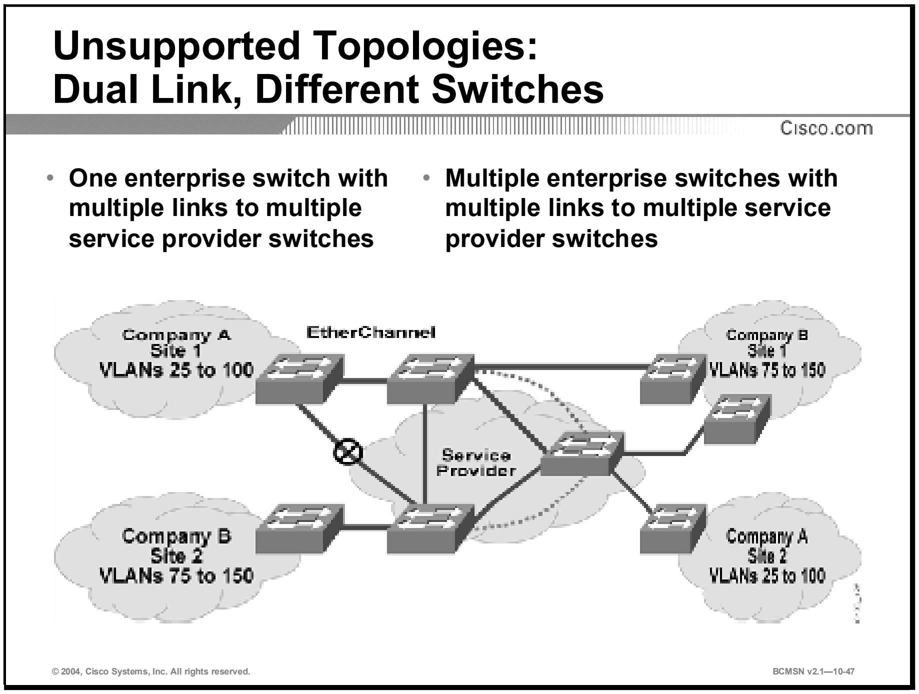

Unsupported Topologies: Dual Link, Different Switches

Cisco.com

- One enterprise switch with multiple links to multiple service provider switches
- Multiple enterprise switches with multiple links to multiple service provider switches

BCMSN v2.1—10-47

Unsupported topologies include the following:

- One customer switch with multiple links to multiple service provider switches
- Multiple customer switches with multiple links to multiple service provider switches

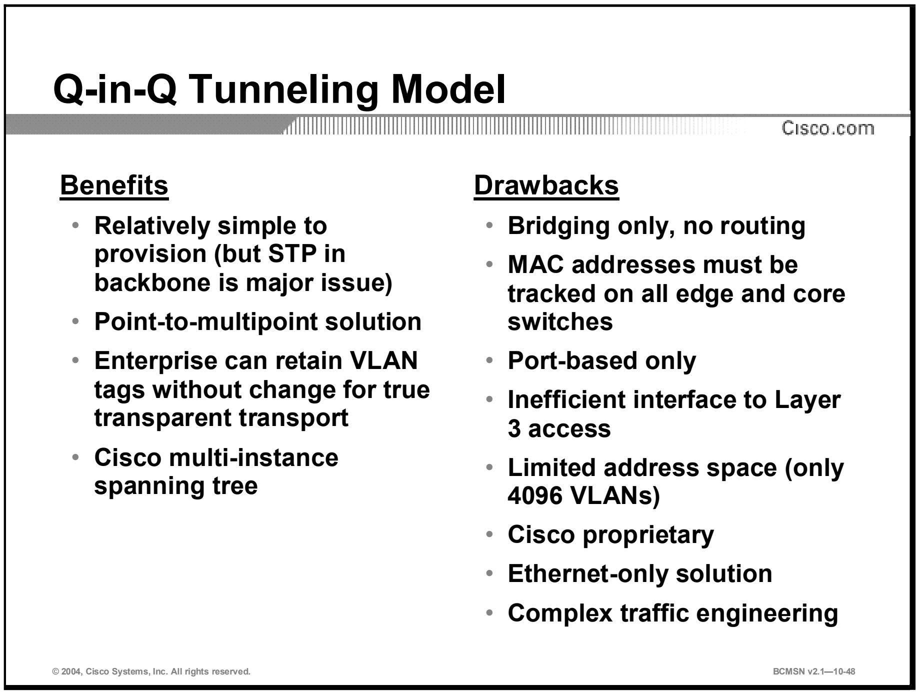

Q-in-Q Tunneling Model

Cisco.com

Benefits

- Relatively simple to provision (but STP in backbone is major issue)
- Point-to-multipoint solution
- Enterprise can retain VLAN tags without change for true transparent transport
- Cisco multi-instance spanning tree

Drawbacks

- Bridging only, no routing
- MAC addresses must be tracked on all edge and core switches
- Port-based only
- Inefficient interface to Layer 3 access
- Limited address space (only 4096 VLANs)
- Cisco proprietary
- Ethernet-only solution
- Complex traffic engineering

BCMSN v2.1—10-48

The primary advantage of Q-in-Q encapsulation is its VLAN transparency. The disadvantages are its Layer 2 operation and its lack of scalability.

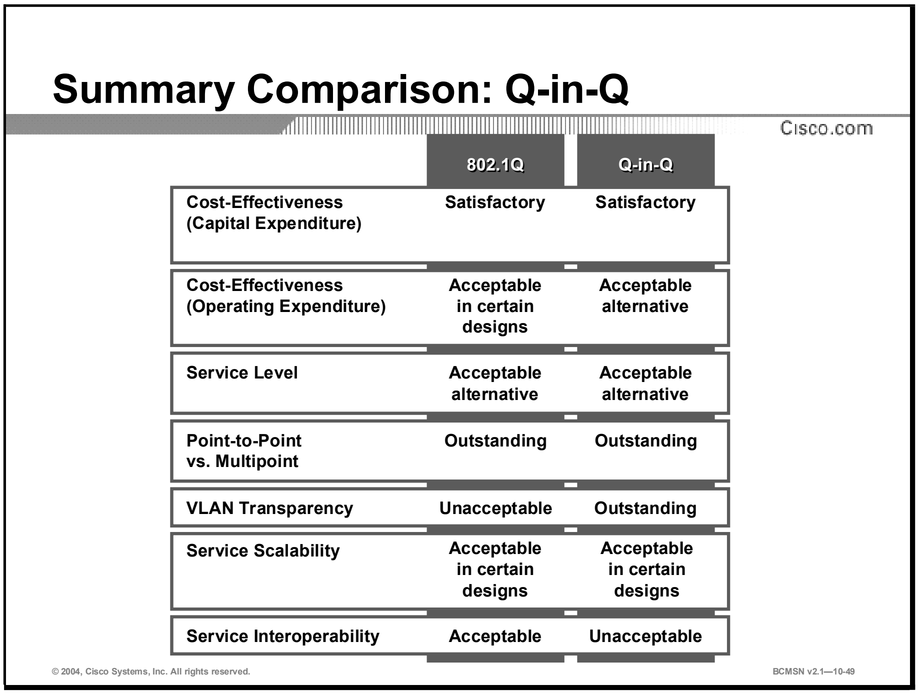

Summary Comparison: Q-in-Q

	802.1Q	Q-in-Q
Cost-Effectiveness (Capital Expenditure)	Satisfactory	Satisfactory
Cost-Effectiveness (Operating Expenditure)	Acceptable in certain designs	Acceptable alternative
Service Level	Acceptable alternative	Acceptable alternative
Point-to-Point vs. Multipoint	Outstanding	Outstanding
VLAN Transparency	Unacceptable	Outstanding
Service Scalability	Acceptable in certain designs	Acceptable in certain designs
Service Interoperability	Acceptable	Unacceptable

BCMSN v2.1—10-49

The primary advantage that 802.1Q-in-Q has over 802.1Q is transparency to customer traffic through encapsulation. A major Q-in-Q drawback is limited interoperability with existing protocols.

Metro Ethernet: EoMPLS Encapsulation

EoMPLS is a tunneling mechanism that maps VLANs through an MPLS core. This solution provides scalability beyond the 4096 VLAN limit (the limit imposed by the size of the VLAN field). This topic explains the operations of EoMPLS.

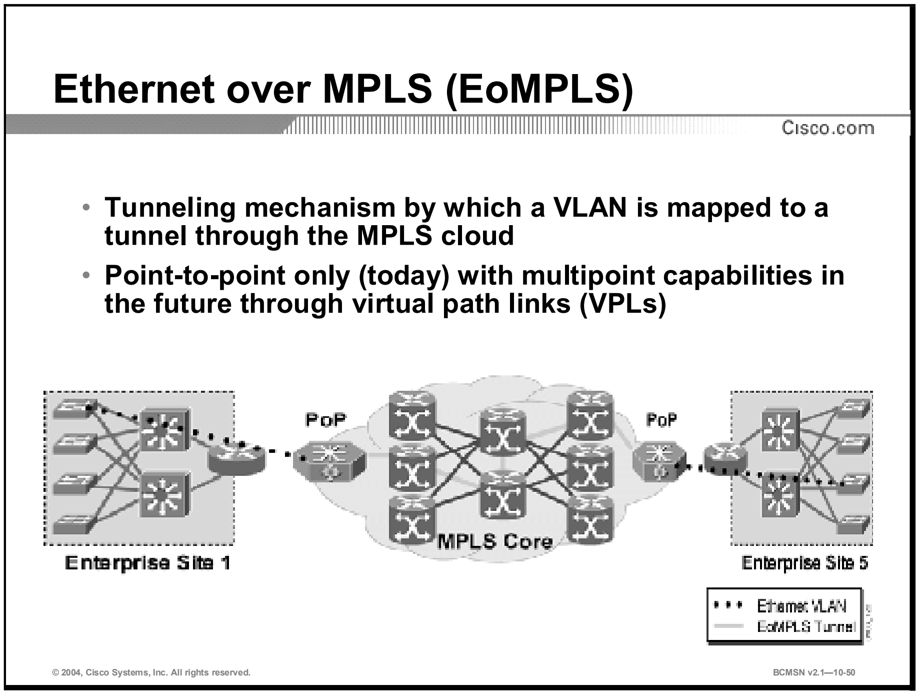

Ethernet over MPLS (EoMPLS)

Cisco.com

- **Tunneling mechanism by which a VLAN is mapped to a tunnel through the MPLS cloud**
- **Point-to-point only (today) with multipoint capabilities in the future through virtual path links (VPLs)**

PoP
PoP
MPLS Core
Enterprise Site 1
Enterprise Site 5
• • • Ethernet VLAN
—— EoMPLS Tunnel

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—10-50

As service providers worldwide have looked to scale their networks, MPLS has emerged as a highly scalable and highly beneficial technology. Many Internet service providers (ISPs), particularly in Europe and Asia, have deployed MPLS to scale their networks. MPLS is beginning to be seen as highly advantageous to Ethernet local exchange carriers (LECs) and Incumbent Local Exchange Carriers (ILECs) who want to provide Ethernet transport services to enterprise customers at Layer 2 while still being able to scale their core architectures. The EoMPLS solution, when used in conjunction with a pure Layer 2 network architecture, has the ability to scale the entire network beyond the 4096 VLAN limitation and to provide the inherent scalability of a Layer 3 network.

The Cisco EoMPLS solution, based on the draft Internet Engineering Task Force (IETF) standard, is an extension of MPLS, which naturally complements the VLAN functionality inherent in Layer 2 architectures. At its simplest, EoMPLS provides a tunneling mechanism for Layer 2 traffic through an MPLS-enabled Layer 3 core. This solution allows the service provider the best of both worlds: the scalability of an MPLS core without having to worry about spanning tree or a Layer 2 transparent service offering.

The service provider uses a small- to mid-range switching product as the access device. The enterprise customer, which is most often using a router as its access device, would be mapped into a particular service provider VLAN, for example, VLAN25. This VLAN is trunked across the fiber plant, and potentially across some additional Layer 2 switches, until it reaches the headend router located in the POP of the service provider. This router is the edge router of the MPLS network, providing access to the MPLS core. The headend router maps VLAN25 to an EoMPLS VC, for example, Tunnel 100. There are two MPLS labels used in EoMPLS; the label inserted by the headend router is the first. To access the MPLS core, a second label is placed on the frame.

The MPLS network uses this label to switch the frame to the appropriate exit point. At the exit point, the second label is stripped off and the first label examined. It is this first label that determines the appropriate egress VLAN on the destination headend router. At this point, the frame can be switched at Layer 2 to the destination customer site. The current IETF standard draft supports only point-to-point tunnels through the MPLS network.

Because many service providers want to offer an alternative service to Frame Relay, it is important to analyze how enterprise customers solve their Frame Relay routing issues. Today, enterprise customers use a headend router, often referred to as the hub, which brings in the physical connections and terminates multiple Frame Relay permanent virtual circuits (PVCs). Each PVC is mapped to a particular spoke, or branch office site, which is a subinterface on the hub router. This subinterface is its own routed subnet, so a hub router connecting to 25 remote sites will have a PVC and a subnet allocated per spoke. This arrangement allows the enterprise to segment its traffic and isolate any failure domains that might exist at the hub or at any of the spokes. While having the hub and the spokes on the same subnet (or VLAN, in the case of Ethernet) may arguably make IP address management easier, the troubleshooting and failure risks for the enterprise increase dramatically. This is why transparent LAN services are not widely used.

In examining problems that enterprise customers face, it is easy to see how EoMPLS can benefit them. Here is an example of how an enterprise customer requiring connectivity from a hub site to multiple spokes and its service provider would solve the problem.

First, the enterprise would hand off an 802.1Q trunk to the service provider; each VLAN on the trunk would be destined for a remote site. (This scenario is analogous to an enterprise handing off a Frame Relay connection with multiple data-link connection identifiers [DLCIs] to the service provider.) The service provider would transport these VLANs, at Layer 2, to the EoMPLS headend router. The router of the service provider would map each of the VLANs of the enterprise to an individual EoMPLS tunnel, one tunnel to each of the remote sites. Thus, although EoMPLS does not natively solve a point-to-multipoint problem, it clearly solves the connectivity problems of the enterprise.

The Cisco EoMPLS technology also provides a solution for service providers looking to interconnect metropolitan networks. As the size of the Metro area increases and as service providers become more interested in connecting their Metro services together, scalability becomes a key concern. While it is possible that the 4096 VLAN limitation may not be reached in a Metro area, it will most certainly be exceeded, many times over, as more Metro areas are connected together. The MPLS solution from Cisco offers the capability of scaling the network core, inter-Metro sites, and inter-POP connectivity. Within a Metro area, the service provider has the ability to use Layer 2 switching to the extent that it will scale, and then to use EoMPLS to bridge the VLAN world to the MPLS world.

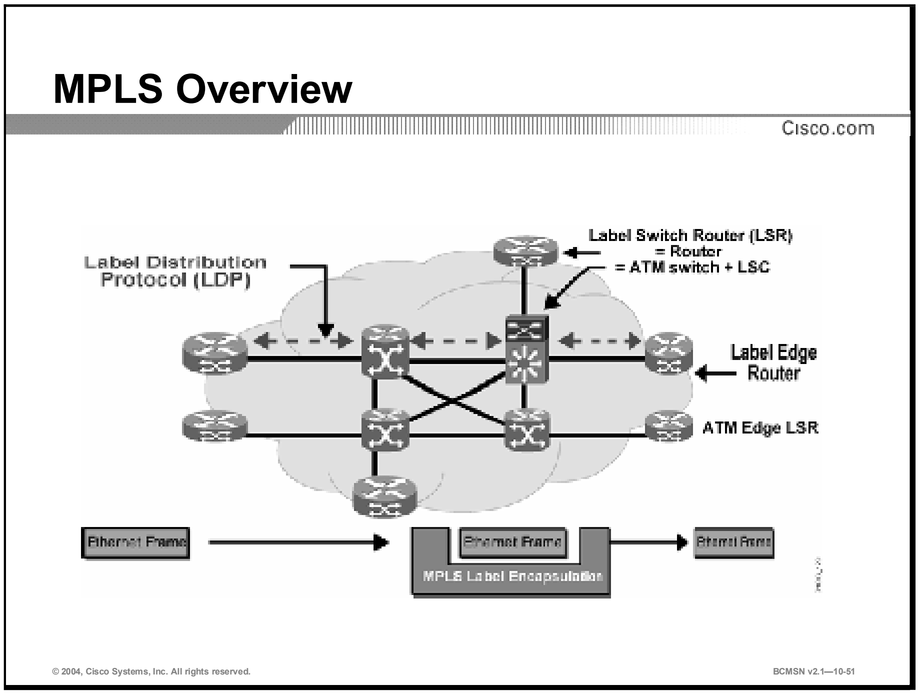

MPLS Overview

Cisco.com

BCMSN v2.1—10-51

Some key MPLS terms are as follows:

- **Label distribution protocol (LDP):** A protocol that communicates labels and their meaning among label switch routers.
- **Label switch router (LSR):** A device that switches labeled frames according to precomputed switching tables. This device can be either a switch or a router.
- **Label switch controller (LSC):** An MPLS-enabled router that controls the operation of an ATM switch in such a way that the two function together as an ATM LSR.
- **Edge label switch router (edge LSR), or label edge router (LER):** The edge device that performs initial frame processing and classification. It applies the first label.

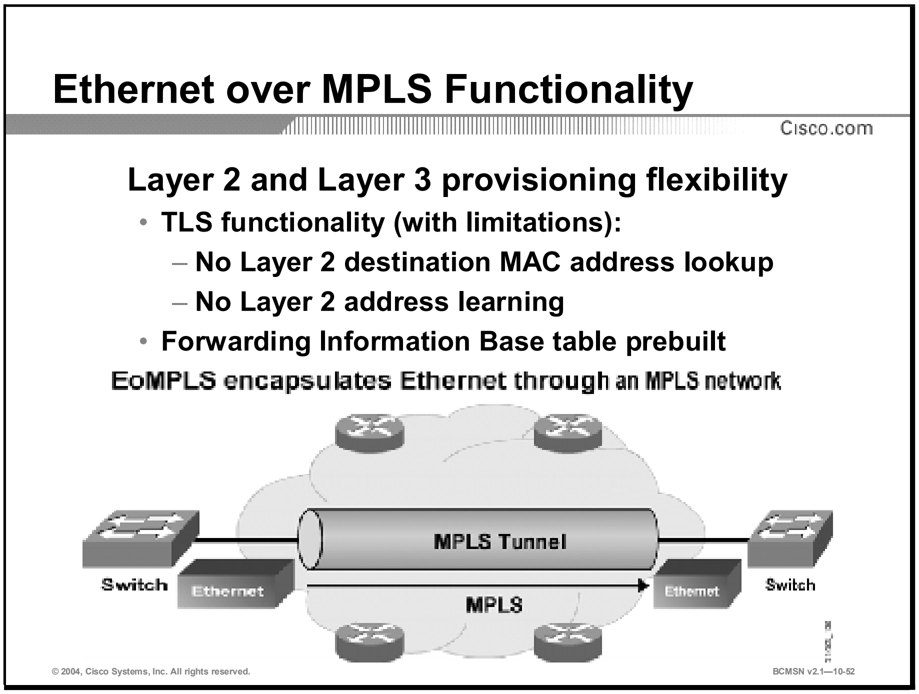

Ethernet over MPLS Functionality

Cisco.com

Layer 2 and Layer 3 provisioning flexibility

- **TLS functionality (with limitations):**
  - **No Layer 2 destination MAC address lookup**
  - **No Layer 2 address learning**
- **Forwarding Information Base table prebuilt**

EoMPLS encapsulates Ethernet through an MPLS network

BCMSN v2.1—10-52

EoMPLS technology leverages an existing MPLS backbone network to deliver TLS-based Ethernet connectivity to the customer site. The concept of a TLS is straightforward: it is the ability to connect two Ethernet networks, which are geographically separate, and have the two networks appear as a single logical Ethernet or VLAN domain. The introduction of such a VLAN transport capability allows service providers to deliver a service that allows VLAN networks in different locations within a Metro service area to be cost-effectively connected at transmission speeds equivalent to Fast Ethernet or Gigabit Ethernet.

When EoMPLS is deployed in conjunction with an MPLS VPN, a service provider can provide tremendous flexibility in the variety of both Layer 2 and Layer 3 network services provisioned for its Metro customers, and can do so over a single, simplified, integrated MPLS backbone network.

EoMPLS Characteristics

EoMPLS transmits Ethernet frames over the service provider core by applying two MPLS labels: the tunnel label and the VC label. EoMPLS requires the customer and service provider to work together to provision the service. This topic describes the characteristics of EoMPLS.

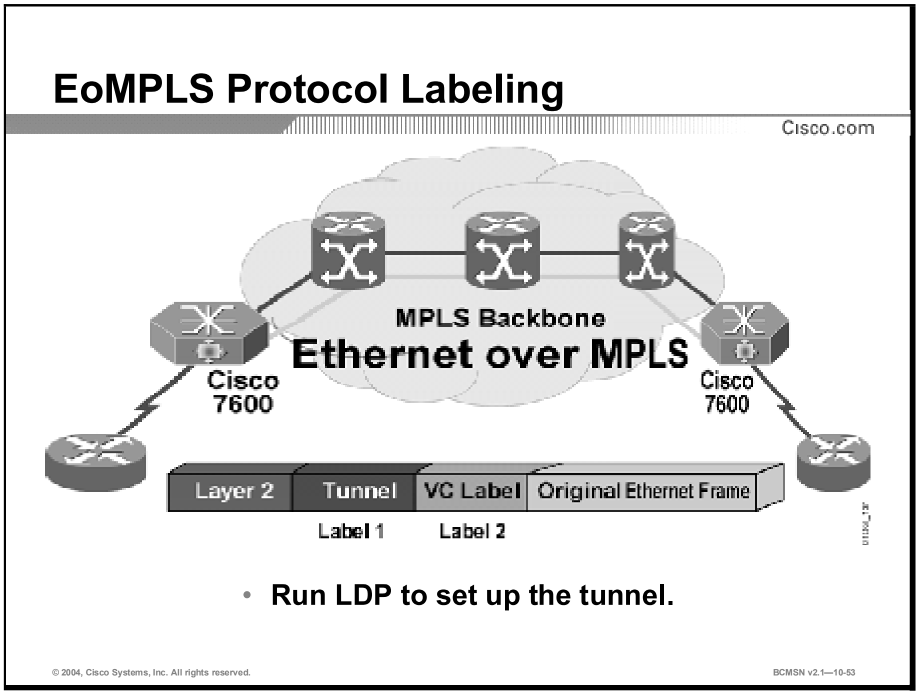

EoMPLS Protocol Labeling

Cisco.com

- Run LDP to set up the tunnel.

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—10-53

EoMPLS technology provides several important benefits to service providers that want to offer a TLS to their Metro customers. It also has several characteristics that the service provider and user should understand to make effective use of this technology:

- Establishing an EoMPLS circuit requires that the service provider customer be assigned a specific physical port on an LER device, such as a Cisco 7600. The identification of that physical port is a critical element in the binding of the MPLS label assigned to the customer EoMPLS VC.
- An enterprise may have more than one EoMPLS VC per physical port, as long as the Ethernet traffic transmitted from the enterprise site to the provider edge (PE) device has specific 802.1Q headers for each EoMPLS VC. This arrangement requires coordination between the service provider and customer in the provisioning of the EoMPLS service.
- EoMPLS VCs are point-to-point transmissions only, as explicitly specified in the IETF Martini draft specifications.
- Traffic sent between the imposition and disposition routers (between LERs) over an EoMPLS VC will take the same path across the IP and MPLS backbone. The label switch path (LSP) may change because of routing changes inside the provider network.
- Adding or removing a point-to-point Layer 2 VC requires configuration of the two VC endpoints (at the two LERs). Provisioning a VC will involve defining an endpoint of a Layer 2 VC at each of the VLAN interfaces at the PE router on the interface that connects to the customer edge (CE).

- The two LERs at the ingress and egress points of the IP and MPLS backbone (the PE routers) are the only routers with knowledge of the Layer 2 transport VCs. All other LSRs will have no table entries for the Layer 2 transport VCs. This situation means that only the PEs require software with EoMPLS functionality.

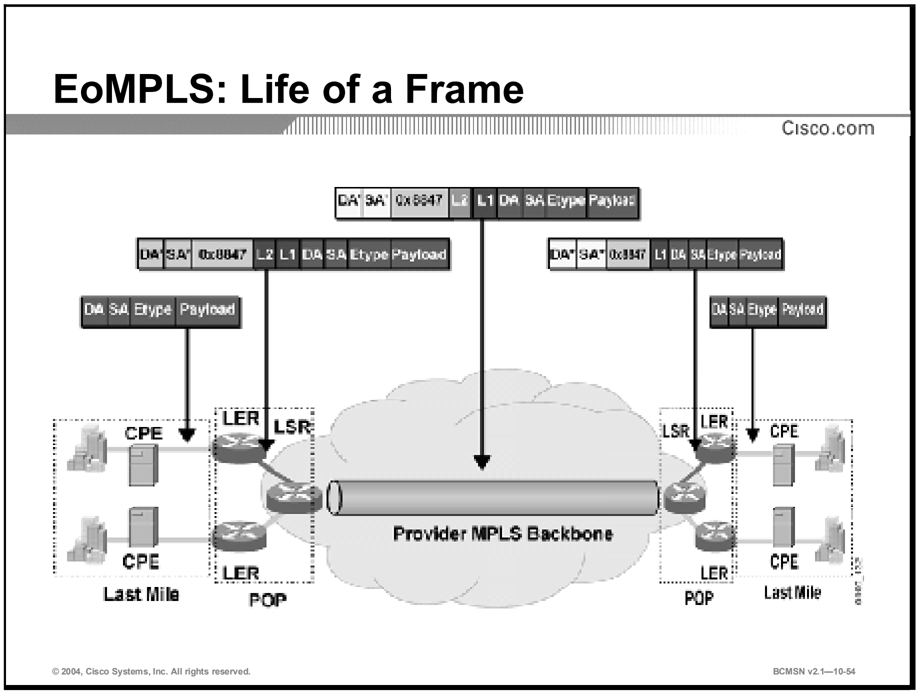

EoMPLS: Life of a Frame

Cisco.com

DA' SA' 0x8847 L2 L1 DA SA Etype Payload
DA' SA' 0x8847 L2 L1 DA SA Etype Payload
DA'' SA'' 0x8847 L1 DA SA Etype Payload
DA SA Etype Payload
DA SA Etype Payload
CPE
LER
LSR
LSR
LER
CPE
CPE
CPE
Provider MPLS Backbone
LER
LER
Last Mile
POP
POP
Last Mile

BCMSN v2.1—10-54

The Layer 2 transport service over MPLS is implemented through the use of two-level label switching between the edge routers. The label used to route the frame over the MPLS backbone to the destination PE is the tunnel label. The label used to determine the egress interface is the VC label. The egress PE allocates a VC label and binds the Layer 2 egress interface to the VC in question, and then it signals this label to the ingress PE via the targeted LDP session.

When the Layer 2 protocol data unit (PDU arrives at the ingress interface of the ingress PE, the router must perform label imposition, and switch the frame to the appropriate outgoing MPLS interface, which routes the frame to the egress LER for the VC in question.

The egress PE receives the frame with only the VC label because its neighbor (known as the “penultimate router”) pops the tunnel label prior to forwarding the frame. The egress PE uses the VC label to perform disposition and switch the frame to the appropriate egress interface.

The figure depicts the life of an EoMPLS frame. Beginning from the left, the Ethernet frame enters the MPLS network and the Layer 1 and Layer 2 labels are applied, along with the header 8847 indicating the EoMPLS encapsulation. This message is transmitted through the core network and exits the egress network, where the labels are removed. The entire Ethernet frame is tunneled through the MPLS network.

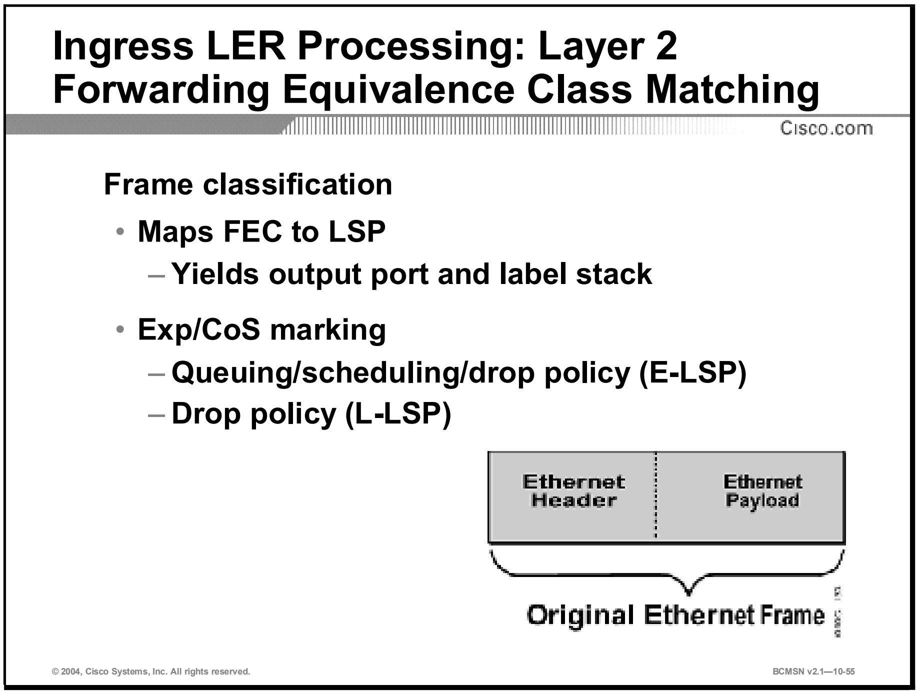

Ingress LER Processing: Layer 2 Forwarding Equivalence Class Matching

Cisco.com

Frame classification

- Maps FEC to LSP
  - Yields output port and label stack
- Exp/CoS marking
  - Queuing/scheduling/drop policy (E-LSP)
  - Drop policy (L-LSP)

Ethernet Header
Ethernet Payload
Original Ethernet Frame

"FEC" refers to the Forwarding Equivalence Class. The ingress LER receives the original Ethernet frame and maps the FEC to the LSP. Depending on the QoS selected, it will either select an E-LSP or an L-LSP for Differentiated Services (DiffServ) operation.

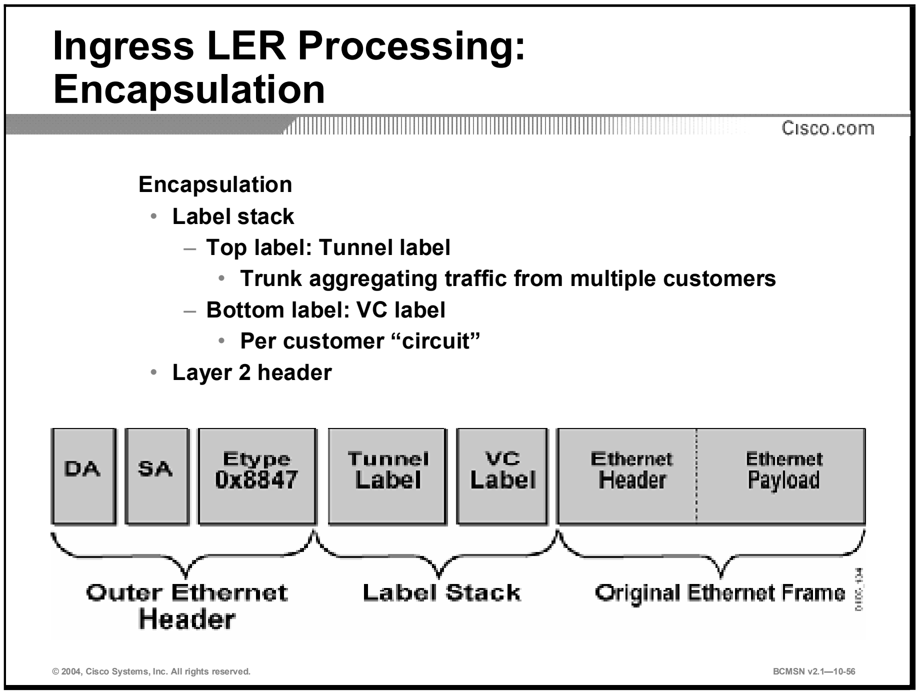

Ingress LER Processing: Encapsulation

Cisco.com

Encapsulation

- **Label stack**
  - **Top label: Tunnel label**
    - **Trunk aggregating traffic from multiple customers**
  - **Bottom label: VC label**
    - **Per customer “circuit”**
- **Layer 2 header**

DA | SA | Etype 0x8847 | Tunnel Label | VC Label | Ethernet Header | Ethernet Payload

Outer Ethernet Header | Label Stack | Original Ethernet Frame

BCMSN v2.1—10-56

The ingress LER will then encapsulate the Ethernet frame with a tunnel label and a VC label. These labels are used for the MPLS transit network. Then an outer Ethernet header is applied to transit the next link. The outer header is not required to be Ethernet; it could be a SONET header, for example.

The figure illustrates the encapsulation occurring on ingress. This encapsulation consists of the outer header, the label stack wherein the labels reside, and finally the original Ethernet frame. In the label stack, the top label is the tunnel label, which allows the aggregation of traffic from multiple customers, whereas the bottom label is the VC label, which is the enterprise circuit.

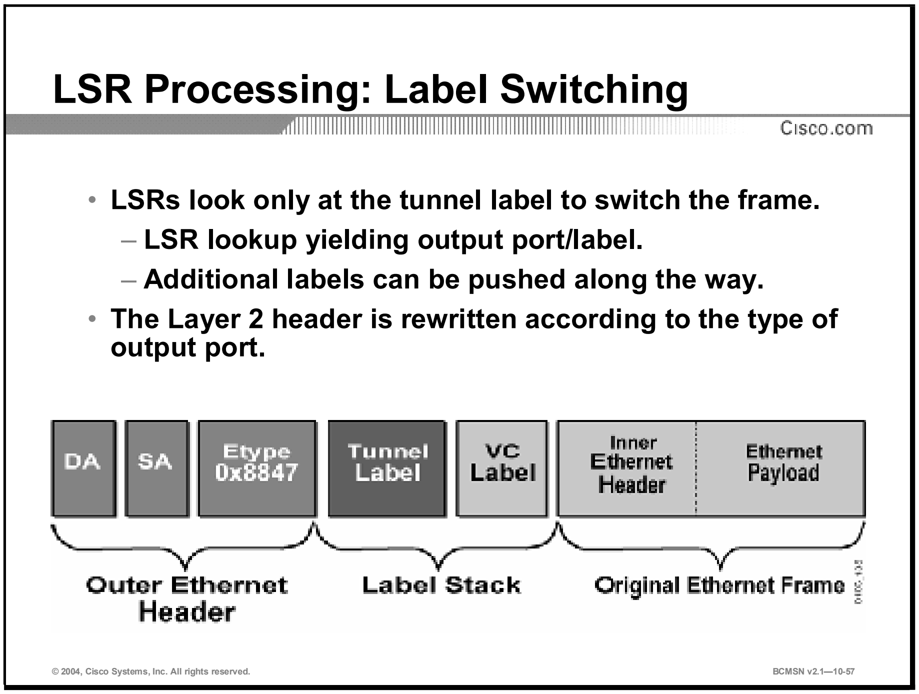

LSR Processing: Label Switching

Cisco.com

- LSRs look only at the tunnel label to switch the frame.
  - LSR lookup yielding output port/label.
  - Additional labels can be pushed along the way.
- The Layer 2 header is rewritten according to the type of output port.

BCMSN v2.1—10-57

Core LSRs read the tunnel label to move the Ethernet frame through the core network.

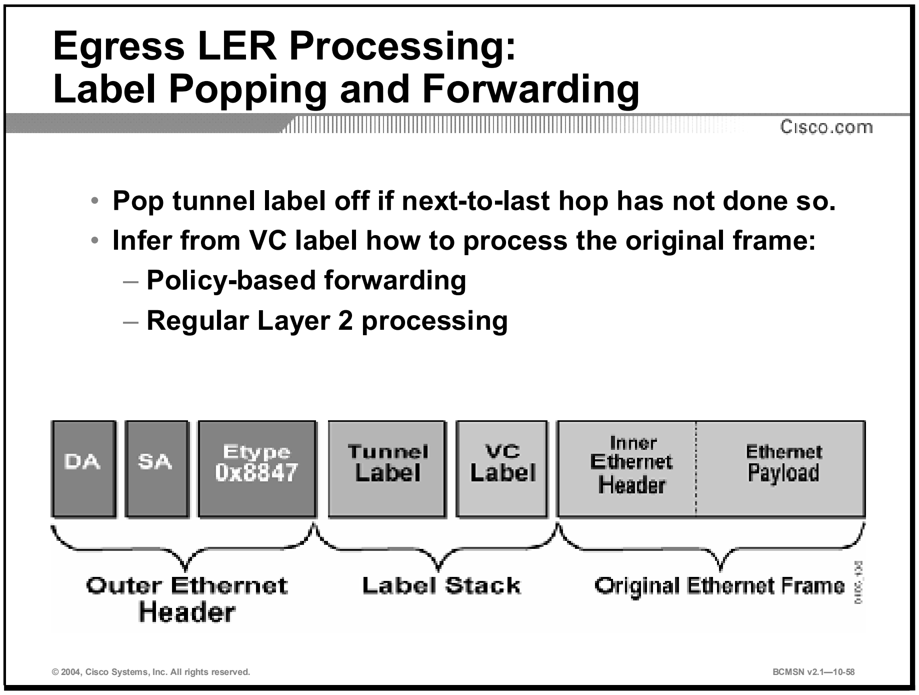

Egress LER Processing: Label Popping and Forwarding

Cisco.com

- Pop tunnel label off if next-to-last hop has not done so.
- Infer from VC label how to process the original frame:
  - Policy-based forwarding
  - Regular Layer 2 processing

BCMSN v2.1—10-58

The egress LER pops the tunnel label off if the next-to-last-hop router has not done so already. The egress LER then reads the VC label to process and forward the original frame.

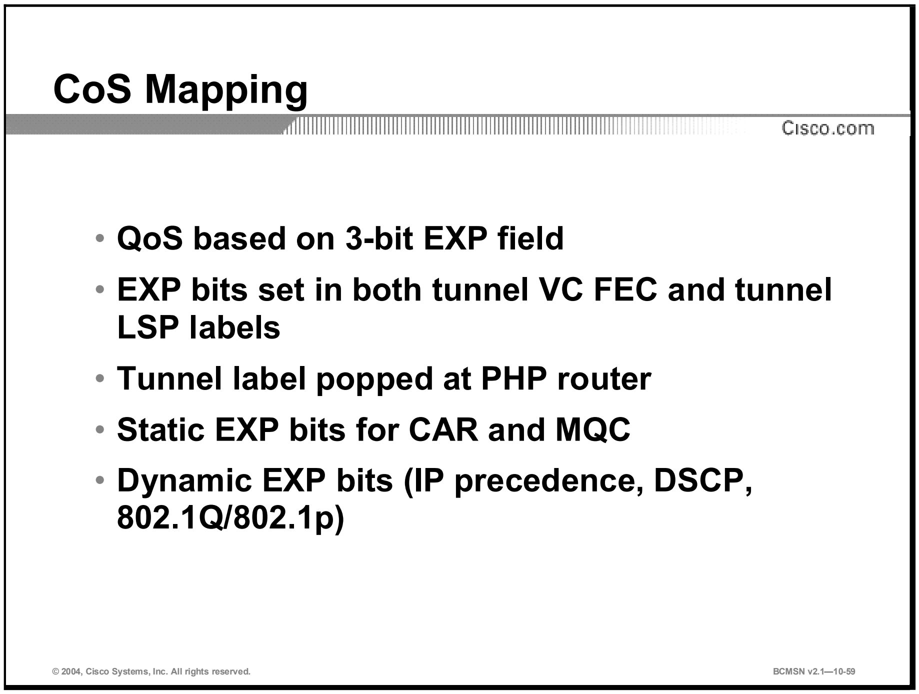

CoS Mapping

- QoS based on 3-bit EXP field
- EXP bits set in both tunnel VC FEC and tunnel LSP labels
- Tunnel label popped at PHP router
- Static EXP bits for CAR and MQC
- Dynamic EXP bits (IP precedence, DSCP, 802.1Q/802.1p)

BCMSN v2.1—10-59

MPLS provides QoS using the three experimental (EXP) bits in a label to determine the queue of packets. In order to support QoS from PE to PE, the EXP bits in both the VC and tunnel labels will have to be set. The EXP bits need to be set in the VC label because the tunnel label is popped at the penultimate router. In the case of EoMPLS, two methods of setting EXP bits are provided:

- **Static setting of EXP bits:** The service provider configures the PE to set the EXP bits in the labels to a given value based on the ingress interface.
- **Using VLAN user priority bits to determine EXP bit settings:** The three user priority bits are used to index into a table of eight values. The value for a given index is used to set the EXP bits. This method may cause out-of-order packets when packets have different user priorities.

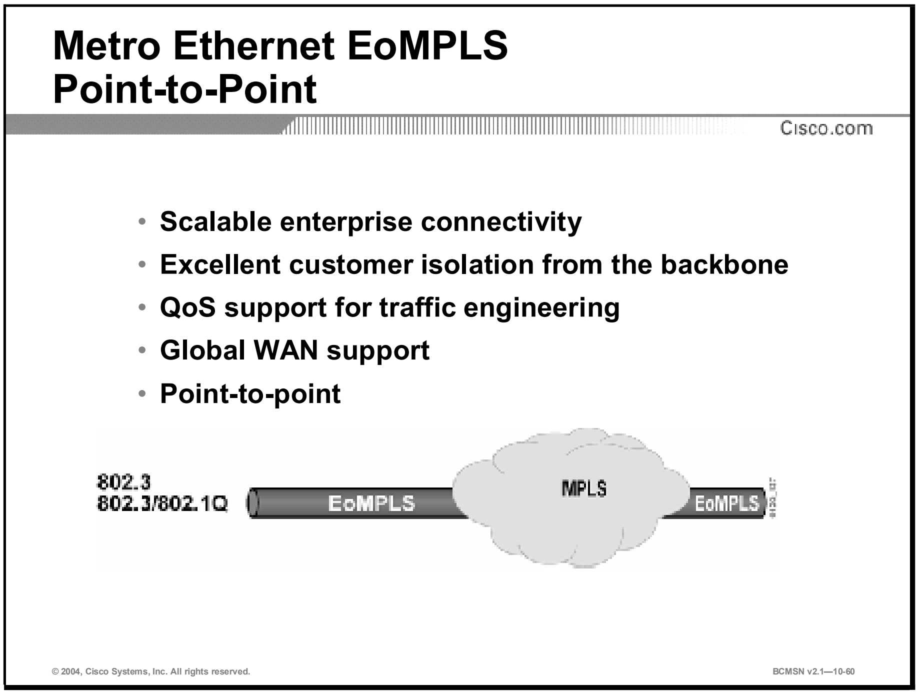

Metro Ethernet EoMPLS Point-to-Point

- Scalable enterprise connectivity
- Excellent customer isolation from the backbone
- QoS support for traffic engineering
- Global WAN support
- Point-to-point

802.3
802.3/802.1Q
EoMPLS
MPLS
EoMPLS

BCMSN v2.1—10-60

EoMPLS provides Metro Ethernet with a WAN reach. In addition, EoMPLS provides QoS support and scalability to Metro Ethernet offerings.

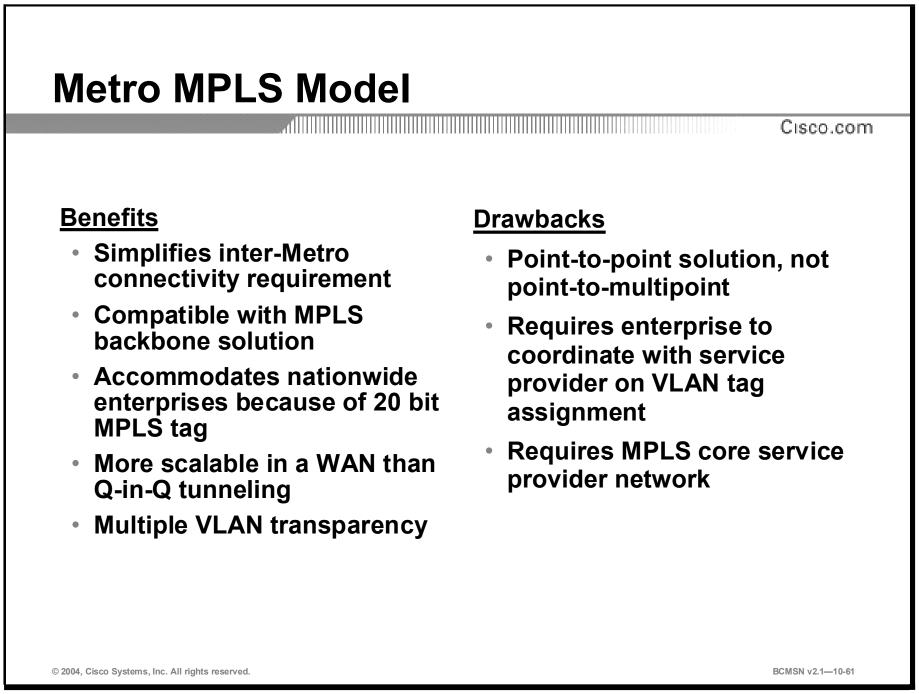

Metro MPLS Model

Benefits

- Simplifies inter-Metro connectivity requirement
- Compatible with MPLS backbone solution
- Accommodates nationwide enterprises because of 20 bit MPLS tag
- More scalable in a WAN than Q-in-Q tunneling
- Multiple VLAN transparency

Drawbacks

- Point-to-point solution, not point-to-multipoint
- Requires enterprise to coordinate with service provider on VLAN tag assignment
- Requires MPLS core service provider network

BCMSN v2.1—10-61

The advantage of EoMPLS is its highly scalable structure and interoperable design. Its drawback is that EoMPLS requires a core MPLS network. MPLS networks are not common.

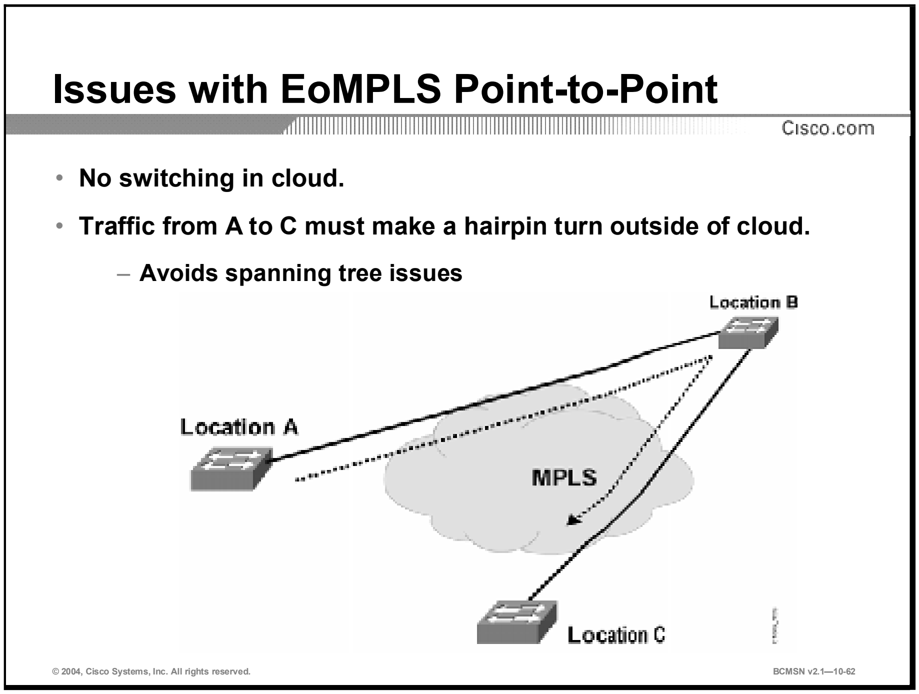

Issues with EoMPLS Point-to-Point

Cisco.com

- **No switching in cloud.**
- **Traffic from A to C must make a hairpin turn outside of cloud.**
  - **Avoids spanning tree issues**

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—10-62

EoMPLS is a point-to-point protocol. At all locations traffic is fully meshed. This traffic has to exit the network to be routed to a noncontiguous destination. This exiting from the network is typically known as a “hairpin.” A way around a hairpin is to support multipoint switching within the core.

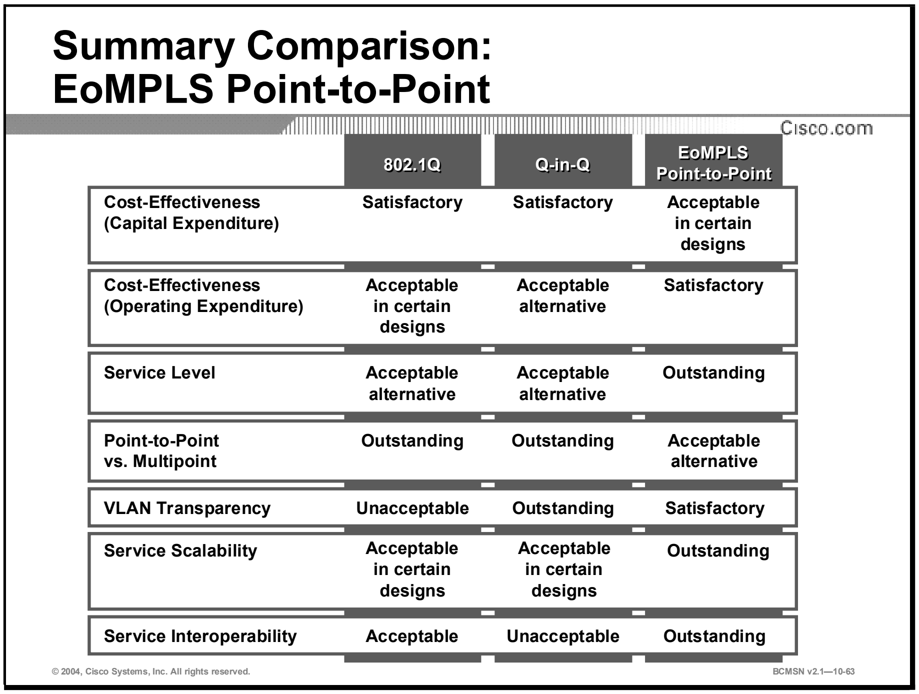

Summary Comparison: EoMPLS Point-to-Point

	802.1Q	Q-in-Q	EoMPLS Point-to-Point
Cost-Effectiveness (Capital Expenditure)	Satisfactory	Satisfactory	Acceptable in certain designs
Cost-Effectiveness (Operating Expenditure)	Acceptable in certain designs	Acceptable alternative	Satisfactory
Service Level	Acceptable alternative	Acceptable alternative	Outstanding
Point-to-Point vs. Multipoint	Outstanding	Outstanding	Acceptable alternative
VLAN Transparency	Unacceptable	Outstanding	Satisfactory
Service Scalability	Acceptable in certain designs	Acceptable in certain designs	Outstanding
Service Interoperability	Acceptable	Unacceptable	Outstanding

In a comparison of EoMPLS to 802.1Q-in-Q, EoMPLS provides outstanding service interoperability and scalability. In addition, MPLS provides QoS operation and support.

Metro Ethernet: EoMPLS Encapsulation Point-to-Multipoint

EoMPLS multipoint extends EoMPLS to include point-to-multipoint capabilities. This topic explains the operations of EoMPLS point-to-multipoint.

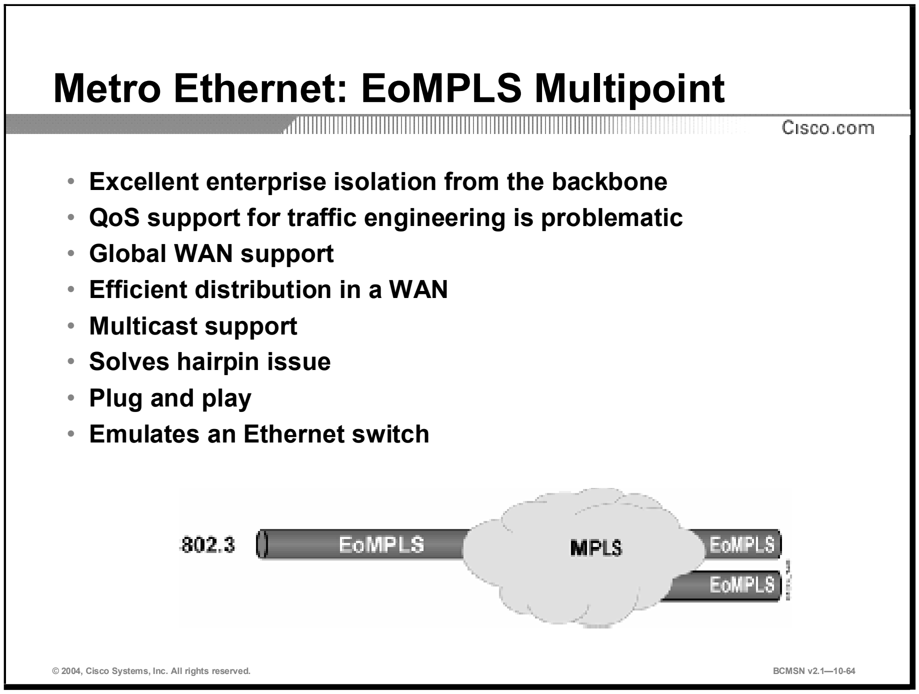

Metro Ethernet: EoMPLS Multipoint

Cisco.com

- Excellent enterprise isolation from the backbone
- QoS support for traffic engineering is problematic
- Global WAN support
- Efficient distribution in a WAN
- Multicast support
- Solves hairpin issue
- Plug and play
- Emulates an Ethernet switch

© 2004, Cisco Systems, Inc. All rights reserved.

BCMSN v2.1—10-64

EoMPLS multipoint provides all the advantages of point-to-point EoMPLS with the addition of multipoint switching in the core network. This option solves the issue associated with hairpin operation in a point-to-point network.

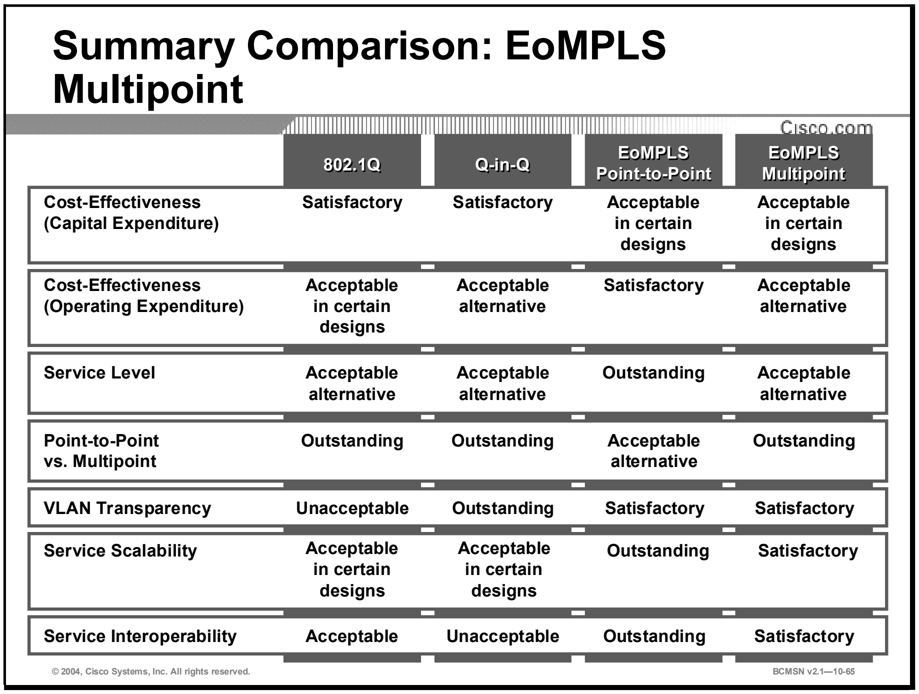

Summary Comparison: EoMPLS Multipoint

	802.1Q	Q-in-Q	EoMPLS Point-to-Point	EoMPLS Multipoint
Cost-Effectiveness (Capital Expenditure)	Satisfactory	Satisfactory	Acceptable in certain designs	Acceptable in certain designs
Cost-Effectiveness (Operating Expenditure)	Acceptable in certain designs	Acceptable alternative	Satisfactory	Acceptable alternative
Service Level	Acceptable alternative	Acceptable alternative	Outstanding	Acceptable alternative
Point-to-Point vs. Multipoint	Outstanding	Outstanding	Acceptable alternative	Outstanding
VLAN Transparency	Unacceptable	Outstanding	Satisfactory	Satisfactory
Service Scalability	Acceptable in certain designs	Acceptable in certain designs	Outstanding	Satisfactory
Service Interoperability	Acceptable	Unacceptable	Outstanding	Satisfactory

EoMPLS multipoint offers the advantages of service interoperability and scalability associated with EoMPLS point-to-point with an additional advantage in the point-to-multipoint category. Now core switching can be Metro Ethernet-aware. However, as with EoMPLS, the service provider must build an MPLS core, which can be quite costly.

Summary

This topic summarizes the key points discussed in this lesson.

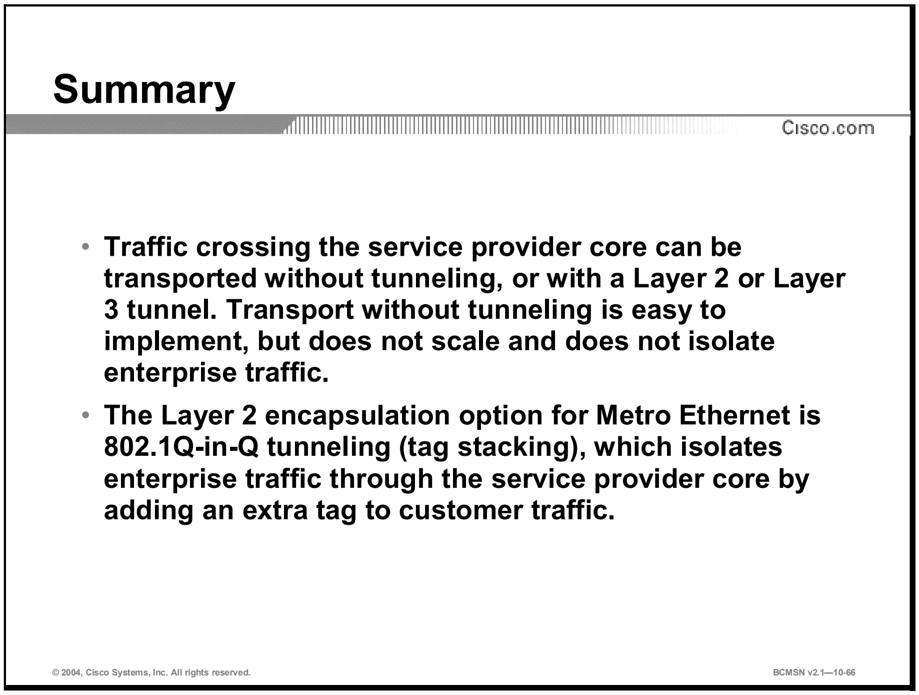

Summary

- Traffic crossing the service provider core can be transported without tunneling, or with a Layer 2 or Layer 3 tunnel. Transport without tunneling is easy to implement, but does not scale and does not isolate enterprise traffic.
- The Layer 2 encapsulation option for Metro Ethernet is 802.1Q-in-Q tunneling (tag stacking), which isolates enterprise traffic through the service provider core by adding an extra tag to customer traffic.

BCMSN v2.1—10-66

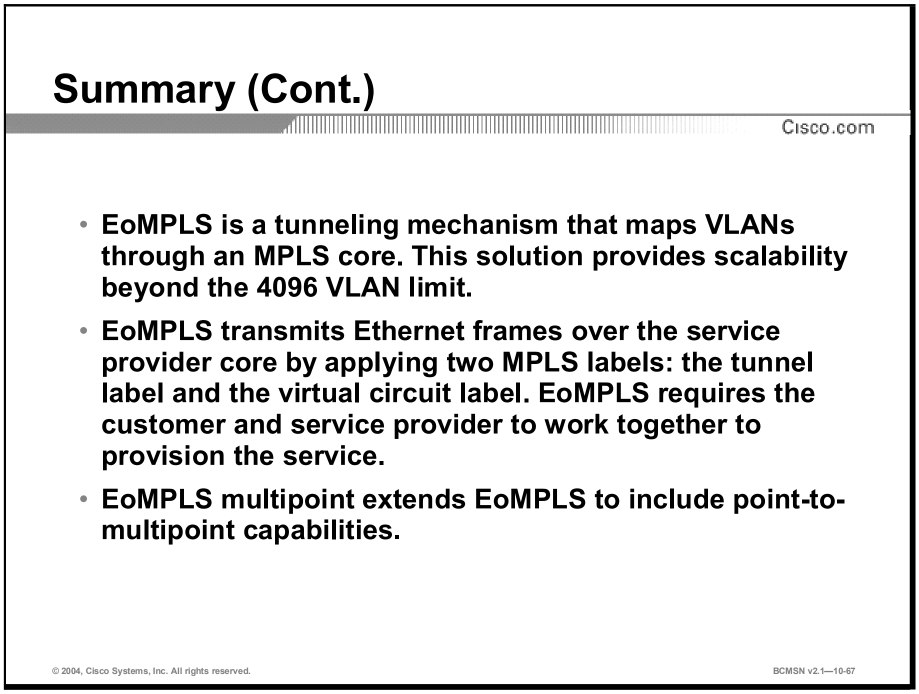

Summary (Cont.)

- EoMPLS is a tunneling mechanism that maps VLANs through an MPLS core. This solution provides scalability beyond the 4096 VLAN limit.
- EoMPLS transmits Ethernet frames over the service provider core by applying two MPLS labels: the tunnel label and the virtual circuit label. EoMPLS requires the customer and service provider to work together to provision the service.
- EoMPLS multipoint extends EoMPLS to include point-to-multipoint capabilities.

BCMSN v2.1—10-67

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) What is the main drawback of 802.1Q transport with no tunneling for Metro Ethernet?

A) complexity

B) lack of QoS

C) bandwidth reduction

D) lack of traffic isolation

Q2) Which topology is supported with 802.1Q-in-Q tunneling?

A) one customer switch with multiple links to the same service provider switch

B) multiple customer switches with multiple links to multiple customer switches

C) one customer switch with multiple links to multiple service provider switches

D) multiple customer switches with multiple links to multiple service provider switches

Q3) Which MPLS term describes the device that applies the first MPLS label to a frame?

A) label switch router

B) label switch controller

C) edge label switch router

D) label distribution protocol

Q4) What are two primary advantages of EoMPLS? (Choose two.)

A) scalability

B) interoperability

C) requirement for MPLS core

D) transparency to users

E) ease of implementation

Q5) What is the primary advantage of EoMPLS multipoint over EoMPLS?

A) scalability

B) cost-effectiveness

C) VLAN transparency

D) point-to-multipoint operation

Quiz Answer Key

Q1) D

Relates to: Practice: Metro Ethernet Tunneling Options

Q2) A

Relates to: Practice: Tag Stacking (Q-in-Q Tunneling)

Q3) C

Relates to: Practice: Metro Ethernet: EoMPLS Encapsulation

Q4) A, B

Relates to: Practice: EoMPLS Characteristics

Q5) D

Relates to: Practice: Metro Ethernet: EoMPLS Encapsulation Point-to-Multipoint

Lesson Assessments

Overview

Use the lesson assessments here to test what you learned in this module. The correct answers and solutions are found in the Lesson Assessment Answer Key.

Outline

This section includes these assessments:

- Quiz 10-1: Examining Metro Ethernet Connectivity Services and Layer 1 Transport Options
- Quiz 10-2: Examining Metro Ethernet Tunneling

Quiz 10-1: Examining Metro Ethernet Connectivity Services and Layer 1 Transport Options

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Describe Cisco Metro Ethernet networking solutions
- List the criteria used to evaluate Metro Ethernet connectivity options
- Describe Transparent LAN Service, its benefits and disadvantages
- Describe Directed VLAN Service, its benefits and disadvantages
- Describe the benefits and disadvantages of DWDM as a Layer 1 implementation of Metro Ethernet
- Describe the benefits and disadvantages of SONET as a Layer 1 implementation of Metro Ethernet
- Describe the benefits and disadvantages of CWDM as a Layer 1 implementation of Metro Ethernet

Quiz

Answer these questions:

Q1) What is the primary Metro deployment option for the Catalyst 4000 switch within a service provider network?

A) CPE

B) high-end access

C) provider edge core

D) high-density access

Q2) Which two considerations make up the scalability selection criteria for transmitting Ethernet over the WAN? (Choose two.)

A) bandwidth

B) number of end users

C) number of network devices

D) number of available VLANs

E) geographic distribution of customers

Q3) How does a TLS present a scalability problem for customer routers?

A) It is limited to 4096 VLANs.

B) QoS is difficult to implement.

C) Multicast traffic is not constrained.

D) Routers must all peer with each other.

Q4) What information is used to select destinations with a DVS?

A) IP address

B) VLAN ID

C) subnet mask

D) MAC address

Q5) What other service is similar to Metro Ethernet over DWDM as far as transparency is concerned?

A) PPP

B) ATM

C) leased-line

D) Frame Relay

Q6) Which feature of Metro Ethernet over SONET is both an advantage and a drawback?

A) general availability

B) bandwidth granularity

C) fault-tolerance mechanism

D) multiservice networking support

Q7) Which two redundancy mechanisms operate at Layer 1? (Choose two.)

A) GBIC

B) spanning tree

C) EtherChannel

D) routing protocols

E) protection switching

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Quiz 10-2: Examining Metro Ethernet Tunneling

Complete this quiz to assess what you learned in the lesson.

Objectives

This quiz tests your knowledge of how to:

- Describe the benefits and disadvantages of transporting traffic across the service provider network without tunneling
- Describe the benefits and disadvantages of 802.1Q-in-Q tunneling
- Describe Ethernet over MPLS
- List the benefits and disadvantages of EoMPLS as a Layer 3 transport option for Metro Ethernet
- Describe EoMPLS encapsulation point-to-multipoint

Quiz

Answer these questions:

Q1) How many VLAN identification values can IEEE 802.1Q support?

A) 1024

B) 2048

C) 4096

D) 8192

Q2) What is a drawback of tag stacking for Metro Ethernet?

A) complexity

B) lack of scalability

C) lack of traffic isolation

D) no spanning tree capability

Q3) What is the purpose of the second MPLS label applied to frames with EoMPLS?

A) to determine the appropriate egress VLAN

B) to determine the appropriate ingress VLAN

C) to map the customer VLAN to an MPLS circuit

D) to switch the frame to the appropriate exit point

Q4) Which MPLS device uses the tunnel label to switch frames?

A) LSR

B) LER

C) LSC

D) ELSR

Q5) Which EoMPLS problem is solved with EoMPLS multipoint?

A) hairpins

B) QoS implementation

C) lack of VLAN transparency

D) lack of service interoperability

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Lesson Assessment Answer Key

Quiz 10-1: Examining Metro Ethernet Connectivity Services and Layer 1 Transport Options

Q1)	D
Q2)	B, E
Q3)	D
Q4)	B
Q5)	C
Q6)	C
Q7)	C, E

Quiz 10-2: Examining Metro Ethernet Tunneling

Q1)	C
Q2)	B
Q3)	D
Q4)	A
Q5)	A

Course Glossary

The Course Glossary for *Building Cisco Multilayer Switched Networks* (BCMSN) v2.1 highlights and defines key terms and acronyms used throughout this course. Many of these terms are also described in the Cisco Internetworking Terms and Acronyms resource, available via http://www.cisco.com/univercd/cc/td/doc/cisintwk/ita/.

Acronym or Term	Definition
1000BASE-T	1000-Mbps baseband Gigabit Ethernet specification using twisted pair based on the IEEE 802.3ab standard.
100BASE-TX	100-Mbps baseband Fast Ethernet specification using UTP wiring.
10BASE-T	10-Mbps baseband Ethernet specification using two pairs of twisted-pair cabling (Categories 3, 4, or 5): one pair for transmitting data and the other for receiving data. 10BASE-T, which is part of the IEEE 802.3 specification, has a distance limit of approximately 328 feet (100 meters) per segment.
AAA	authentication, authorization, and accounting (pronounced “triple a”). Network security services that provide the primary framework through which you set up access control on your router or access server. AAA protocol requirements for network access are defined in RFC 2989.
ACL	access control list. Definitions kept by routers and switches to control access to or from the device for a number of services (for example, to prevent packets with a certain IP address from leaving a particular interface on the device).
address	Data structure or logical convention used to identify a unique entity, such as a particular process, network interface (IP) or a network device (DECnet).
ARP	Address Resolution Protocol. Protocol used to map an IP address to a MAC address. Defined in RFC 826.
AVVID	Cisco Architecture for Voice, Video and Integrated Data.
backbone	Part of a network that acts as the primary path for traffic that is most often sourced from, and destined for, other networks.
BackboneFast	A feature on the switch that reduces the Spanning Tree Protocol convergence time from 50 seconds to 20 to 30 seconds.
backplane	The physical connection between an interface processor or card and the data buses and the power distribution buses inside a chassis.
backup	A way of providing high availability by using redundant links. Backup connection can be established either via dial-up or by using permanent connections.
bandwidth	The difference between the highest and lowest frequencies available for network signals. The term is also used to describe the rated throughput capacity of a given network medium or protocol.
BPDU	bridge protocol data unit. Spanning Tree Protocol mechanism to exchange information among bridges in the network.
bps	bits per second.
Bps	bytes per second.
bridge	Device that connects and passes frames between two network segments that use the same communications protocol. Bridges operate at the data link layer (Layer 2) of the OSI reference model. In general, a bridge filters, forwards, or floods an incoming frame based on the MAC address of that frame.
bridge forwarding	A process that uses entries in a filtering database to determine whether frames with a given MAC destination address can be forwarded to a given port or ports. Described in the IEEE 802.1 standard.
broadcast	Data packets or frames that are sent to all nodes on a network.
broadcast address	A special address reserved for sending a message to all stations. Generally, a data-link broadcast address is a MAC destination address of all ones. An IPv4 broadcast address is one in which the host portion of the address is all ones. There is no corresponding capability in IPv6.

Acronym or Term	Definition
broadcast domain	Set of all devices that receive broadcast frames originating from any device within the set. Routers typically bound data link broadcast domains because routers do not forward data link broadcast frames.
broadcast storm	An undesirable network event in which many broadcasts are repeatedly replicated and sent simultaneously across all network segments. A broadcast storm uses substantial network bandwidth and typically causes network timeouts.
buffer	A storage area used for handling data in transit. Buffers are used in internetworking to compensate for differences in processing speed between network devices. Bursts of data can be stored in buffers until they can be handled by slower processing devices. Sometimes referred to as a packet buffer.
Building Access submodule	A submodule within the Enterprise Composite Network model. Contains end-user workstations, IP Phones, and Layer 2 access switches for connecting devices to the Building Distribution component.
Building Distribution submodule	A submodule within the Enterprise Composite Network model. Provides aggregation of access networks using Layer 3 switching. Performs routing, QoS, and access control.
campus	One or more buildings with multiple virtual and physical networks, connected across a high-performance backbone.
Campus Backbone submodule	A submodule within the Enterprise Composite Network model that connects distribution modules.
Campus Infrastructure module	A module within the Enterprise Composite Network model that comprises the Building Access, Building Distribution, and Campus Backbone submodules.
CBWFQ	class-based weighted fair queuing extends WFQ functionality to provide support for user-defined traffic classes.
CDP	Cisco Discovery Protocol. Media- and protocol-independent device-discovery protocol that runs on all equipment manufactured by Cisco, including routers, access servers, bridges, and switches. Using CDP, a device can advertise its existence to other devices and receive information about other devices on the same LAN or on the remote side of a WAN. Runs on all media that support SNAP, including LANs, Frame Relay, and ATM media.
CEF	Cisco Express Forwarding. Scalable, distributed Layer 3 switching technology designed to enhance network performance within supported platforms.
Cisco IOS	Cisco system software that provides common functionality, scalability, and security for Cisco products. Cisco IOS software allows centralized, integrated, and automated installation and management of internetworks while ensuring support for a wide variety of protocols, media, services, and platforms.
CLI	command-line interface. A syntactic user interface that allows interaction with the application or operating system through commands and optional arguments entered from a keyboard.
codec	coder-decoder. Integrated circuit device that transforms analog acoustic signals into a digital bit stream (coder) and digital signals back into analog signals (decoder).
collision domain	A single CSMA/CD network in which there will be a collision if two devices attached to the system transmit at the same time. Ethernet uses CSMA/CD. Repeaters and hubs extend the collision domain.
congestion	Traffic in excess of network capacity.
congestion avoidance	Mechanism by which a network controls the traffic entering the network to minimize delays.

Acronym or Term	Definition
CoS	class of service. An indication of how an upper-layer protocol requires a lower-layer protocol to treat its messages. In Ethernet networks, CoS is signaled using three bits in the frame header. Closely related to ToS in networks implemented using Cisco routers and switches.
CPE	customer premises equipment. Terminating equipment, such as terminals, telephones, and modems installed at customer sites, and connected to the telephone company network. Can also refer to any telephone equipment supplied by the telephone company residing on the customer site.
CPU	Central Processing Unit. Computing part of a computer or networking device.
crossbar	A type of high-performance switching fabric.
CSMA/CD	carrier sense multiple access collision detect.
dark fiber	Unused fiber-optic cable. When it is carrying a signal, it is called lit fiber.
data link layer	Layer 2 of the OSI reference model. This layer responds to service requests from the network layer and issues service requests to the physical layer. It provides reliable transit of data across a physical link. The data link layer is concerned with physical addressing, network topology, line discipline, error notification, ordered delivery of frames, and flow control. The IEEE divided this layer into two sublayers: the MAC sublayer and the LLC sublayer. Sometimes simply called link layer.
designated bridge	Bridge that incurs the lowest path cost when forwarding a frame from a segment to the root bridge.
destination address	Address of a network device that is receiving data.
deterministic load distribution	Technique for distributing traffic between two bridges across a circuit group. Guarantees packet ordering between source-destination pairs and always forwards traffic for a source-destination pair on the same segment in a circuit group for a given circuit-group configuration.
differentiated service	A paradigm for providing QoS on the Internet by employing a small, well-defined set of building blocks from which a variety of services can be built.
E-Commerce module	A module within the Enterprise Composite Network model. The E-commerce module enables enterprises to successfully deploy e-commerce applications.
Edge Distribution module	A module within the Enterprise Composite Network model that aggregates the connectivity from the various elements at the Enterprise Edge module and routes the traffic into the Campus Backbone submodule.
Edge LSR	Edge Label Switch Router. The role of an Edge LSR is to turn unlabeled packets into labeled packets, and vice versa.
EDI	Electronic Data Interchange. Electronic communication of operational data, such as orders and invoices, between organizations.
EGP	Exterior Gateway Protocol. Internet protocol for exchanging routing information between autonomous systems. Documented in RFC 904. Not to be confused with the general term exterior gateway protocol. EGP is an obsolete protocol that was replaced by BGP.
EIGRP	Enhanced Interior Gateway Routing Protocol. Advanced version of IGRP developed by Cisco. Provides superior convergence properties and operating efficiency. A hybrid, it combines the advantages of link state protocols with those of distance vector protocols.
e-mail	Electronic Mail. Widely used application in which text messages are transmitted electronically between end users over various types of networks using various network protocols. Underlying network application protocols include SMTP and POP.

Acronym or Term	Definition
encapsulation	Wrapping of data in a particular protocol format. For example, Ethernet data is wrapped in a specific Ethernet frame before network transit.
encryption	Application of a specific algorithm to data so as to alter the representation of the data making it incomprehensible to those who do not have access to the algorithm and key required to reverse the process.
Enterprise Campus	A functional area within Enterprise Composite Network model. Comprises the modules required to build a highly robust campus network in terms of performance, scalability, and availability.
Enterprise Composite Network model	A model of enterprise campus networks which logically and physically segregates the campus along functional boundaries.
Enterprise Edge	A functional area within the Enterprise Composite Network model comprises four modules: WAN module, E-Commerce module, Internet Connectivity module, and Remote Access and VPN module.
enterprise network	The comprehensive network that connects an organization. It includes all LAN, campus, metropolitan, and WAN links and equipment.
Ethernet	Baseband LAN specification invented by Xerox Corporation and developed jointly by Xerox, Intel, and Digital Equipment Corporation. Ethernet networks use CSMA/CD and run over a variety of cable types at 10 Mbps. Current Ethernet implementations are defined in the IEEE 802.3 series of standards.
EXEC	Interactive command processor of Cisco IOS software.
failover	A backup operational mode in which the functions of a system component (such as a processor, server, network, or database, for example) are assumed by secondary system components when the primary component becomes unavailable through either failure or scheduled down time.
failure domain	A group of Layer 2 switches connected together is called a Layer 2 switched domain. The Layer 2 switched domain can be considered as a failure domain because a misconfigured or malfunctioning workstation can introduce errors that will impact or disable the entire domain.
Fast EtherChannel	Bundled Fast Ethernet links that appear as one logical interface.
Fast Ethernet	Any of a number of 100-Mbps Ethernet specifications. Fast Ethernet offers a speed increase 10 times that of the 10BASE-T Ethernet specifications while preserving such qualities as frame format, MAC mechanisms, and MTU. These similarities allow the use of existing Ethernet applications and network management tools on Fast Ethernet networks. Based on an extension to the IEEE 802.3 specification.
flat addressing	Scheme of network addressing that does not use a logical hierarchy to determine association. For example, MAC addresses are flat. Bridging protocols must flood packets throughout a flat network to deliver the packet to the appropriate location.
frame	Logical grouping of information sent as a data link layer unit over a transmission medium. Often refers to the header and the trailer, used for synchronization and error control that surround the user data contained in the unit.
full duplex	A link segment capable of transferring signals in both directions simultaneously.
Gb	gigabit. Approximately 1,000,000,000 bits.
GbE	Gigabit Ethernet. Standard for a high-speed Ethernet at 1 Gbps, approved by the Institute of Electrical and Electronics Engineers (IEEE) 802.3z standards committee in 1996.
Gbps	gigabits per second.
GBps	gigabytes per second.

Acronym or Term	Definition
Gigabit EtherChannel	Bundled multiple Gigabit Ethernet links, which appear as one logical interface.
half duplex	A link segment capable of transferring signals in either direction along the link, but not in both directions simultaneously.
header	Control information placed before data when encapsulating data for network transmission.
hello packet	Mechanism that many protocols use to indicate that a device is operational.
high availability	An intelligent network service that, when carefully implemented, ensures adequate connectivity for mission-critical applications through fault tolerance, device redundancy, redundant physical connections, and route redundancy.
host	In internetworking, a device connected to a network that provides data and services to other computers. Services may include data storage, file transfer, data processing, e-mail, bulletin board services, World Wide Web, DHCP, etc.
HSRP	Hot Standby Router Protocol. Provides high network availability and transparent network topology changes. HSRP creates a Hot Standby router group with a lead router that services all packets sent to the Hot Standby address. Other routers in the group monitor the lead router, and if it fails, one of the standby routers inherits the lead position and the Hot Standby group address. HSRP is documented in RFC 2281.
ICMP	Internet Control Message Protocol. Network layer Internet protocol that reports errors and provides other information relevant to IP packet processing. Documented in RFC 792.
IEEE	Institute of Electrical and Electronics Engineers. Professional organization whose activities include the development of communications and network standards. IEEE LAN standards are the predominant LAN standards today.
IEEE 802.2	IEEE LAN protocol that specifies an implementation of the LLC sublayer of the data-link layer. IEEE 802.2 handles errors, framing, flow control, and the network layer (Layer 3) service interface. Used in IEEE 802.3 and IEEE 802.5 LANs.
IEEE 802.3	IEEE LAN protocol that specifies an implementation of the physical layer and the MAC sublayer of the data link layer. IEEE 802.3 uses CSMA/CD access at a variety of speeds over a variety of physical media. Extensions to the IEEE 802.3 standard specify implementations for Fast Ethernet. Physical variations of the original IEEE 802.3 specification include 10BASE-2, 10BASE-5, 10BASE-FL, 10BASE-T, and 10Broad36. Physical variations for Fast Ethernet include 100BASE-TX, 100BASE-T2, 100BASE-T4, and 100BASE-FX.
IETF	Internet Engineering Task Force. IETF is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual. The IETF operates under the auspices of ISOC.
IGMP	Internet Group Management Protocol. Protocol used by IP hosts to report their multicast group membership requests to an adjacent multicast router. IGMP is defined in RFC 3376.
intelligent network services	Services that enable application awareness within the network. Intelligent network services add intelligence to the network infrastructure beyond that required to just move a datagram between two points. Examples of intelligent network services include network management, security, high availability, quality of service (QoS), and IP multicasting.

Acronym or Term	Definition
Internet	The largest global internetwork, connecting tens of thousands of networks worldwide and having a "culture" that focuses on research and standardization based on real-life use. Many leading-edge network technologies come from the Internet community. The Internet evolved in part from Advanced Research Projects Agency Network (ARPANET), at one time, called the Defense Advanced Research Projects Agency (DARPA) Internet. Not to be confused with the general term internet.
Internet Connectivity module	A module within the Enterprise Edge functional area of the Enterprise Composite Network model. This module provides internal enterprise users with connectivity to Internet services.
internetwork	Collection of networks interconnected by routers and other devices that functions (generally) as a single network. Sometimes called an internet, which is not to be confused with the Internet.
intranet	Intranet is a closed, organization-wide network that includes LANs and WANs. It frequently uses open standards such as TCP/IP instead of proprietary protocols traditionally used for LANs and WANs.
IP	Internet Protocol. IP allows for transmitting blocks of data called datagrams from sources to destinations, where sources and destinations are hosts identified by fixed-length addresses. Defined in RFC 791.
IP address	An address that indicates where a device is. IPv4 addresses are a fixed length of four octets (32 bits). The most significant bits of an address are the network number, followed by the local address.
IP datagram	Fundamental unit of information passed across the Internet. Contains source and destination addresses along with data and a number of fields that define such things as the length of the datagram and the header checksum. The IPv6 header is similar in purpose, though very different in structure.
IP multicast	Internet Protocol multicast. A packet routing technique that allows IP traffic to be propagated efficiently from one source to a number of destinations, or from many sources to many destinations. Rather than sending duplicate packets, one to each destination, only one packet is sent out each interface on which a multicast group identified by a single IP destination group address is registered. This can greatly reduce the required bandwidth.
IP precedence	Use of three bits from the ToS octet in the IP header to provide limited prioritization for IP packets in a routed network.
IP telephony	Internet Protocol telephony. The transmission of voice calls over data networks that use the Internet Protocol (IP).
ISL	Inter-Switch Link. Cisco proprietary protocol that maintains VLAN information while traffic from multiple VLANs flows between switches and routers on a single physical link.
ISO	International Organization for Standardization. International organization that is responsible for a wide range of standards, including those relevant to networking. ISO developed the OSI reference model, a popular networking reference model.
ITU-T	International Telecommunication Union Telecommunication Standardization Sector. International body that develops worldwide standards for telecommunications technologies. A United Nations agency, the ITU-T carries out the functions of the former Consultative Committee for International Telegraph and Telephone (CCITT).
jitter	The interpacket delay variance; that is, the difference, in time, between interpacket arrival and departure intervals. Reducing jitter is important for real-time applications such as voice and video. Analog communication line distortion caused by the variation of a signal from its reference timing positions. Jitter can cause data loss, particularly at high speeds.

Acronym or Term	Definition
jitter buffer	Dejitter buffers are used at the receiving end to smooth delay variability and allow time for decoding and decompression. They help on the first talk spurt to provide smooth playback of voice traffic.
kbps	kilobits per second.
LAN	local-area network. High-speed, low-error-rate data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. The most widely used LAN implementation technology is Ethernet at various rates.
latency	Delay in time.
Layer 2 switching	Switching based on Layer 2 (data link layer) information. The current generation of Ethernet Layer 2 switches are functionally equivalent to bridges. The exposures in a large bridged network include broadcast storms, spanning-tree loops, and address limitations.
Layer 3 switching	Integrates routing with switching to yield very high routing throughput rates typical of Layer 2 switches while offering network layer (Layer 3) routing services and data link layer (Layer 2) termination.
LLQ	low latency queuing. Feature that brings strict priority queuing to CBWFQ. Strict priority queuing allows delay-sensitive data, such as voice, to be dequeued and sent first (before packets in other queues are dequeued), giving delay-sensitive data preferential treatment over other traffic.
load balancing, load-sharing	In routing, the capability of a router to distribute traffic over all its network ports that are the same distance from the destination address. Good load-balancing algorithms use both line speed and reliability information. Load balancing increases the use of network segments, thus increasing effective network bandwidth.
LRE	Long-Reach Ethernet. Ethernet standard frames over single-pair wiring at distances of up to 5000 feet.
LSR	label switch router. The role of an LSR is to forward packets in an MPLS network by looking only at the fixed-length label.
MAC	Media Access Control. Lower of the two sublayers of the data link layer defined by the IEEE. The MAC sublayer handles access to shared media, and makes such decisions as whether token passing or contention will be used.
MAC address	Standardized data-link layer address that is required for every port or device that connects to an Ethernet-based LAN. Other devices in the network use these addresses to locate specific ports in the network and to create and update routing tables and data structures. MAC addresses are six bytes long and are controlled by the IEEE. Also known as a hardware address, MAC layer address, and physical address.
MAC address learning	Service that characterizes a learning bridge, in which the source MAC address of each received packet is stored so that future packets destined for that address can be forwarded only to the bridge interface on which that address is located. Packets destined for unrecognized addresses are forwarded out every bridge interface. This scheme helps minimize traffic on the attached LANs. MAC address learning is defined in the IEEE 802.1D standard.
MAN	metropolitan-area network. Network that spans a metropolitan area. Generally, a MAN spans a larger geographic area than a LAN, but a smaller geographic area than a WAN.
Mb	megabit. Approximately 1,000,000 bits.
MB	megabyte. Depending on the context, it can mean either 1,000,000 or 1,048,576 (2^20) bytes.

Acronym or Term	Definition
Mbps	megabits per second. A bit rate expressed in millions of binary bits per second.
MIB	Management Information Base. Database of network management information that is used and maintained by a network management protocol, such as SNMP or Common Management Information Protocol (CMIP). The value of a MIB object can be changed or retrieved using SNMP or CMIP commands, usually through a GUI network management system. MIB objects are organized in a tree structure that includes public (standard) and private (proprietary) branches defined in ASN.1.
MLS	Multilayer Switching.
MM fiber	Multimode Fiber. A less costly fiber-optic medium in which light travels in multiple modes.
MPLS	Multiprotocol Label Switching. Switching method that forwards Network Layer traffic using a label. This label instructs the routers and the switches in a network where to forward the packets based on preestablished routing information determined as the packet entered the network. MPLS is defined in RFC 3031.
MTU	maximum transmission unit. Maximum packet size, in bytes, that a particular interface can transmit without fragmentation.
multicast	The transmission of packets from a single source to multiple destinations in a way which conserves network bandwidth by reducing the duplication of packets sent.
multicast router	Router used to send IGMP query messages on their attached local networks. Host members of a multicast group respond to a query by sending IGMP reports noting the multicast groups to which they belong. The multicast router takes responsibility for forwarding multicast datagrams from one multicast group to all other networks that have members in the group.
multilayer switch	Switch that filters and forwards packets based on MAC addresses and network addresses.
NetFlow	NetFlow technology efficiently provides the metering base for a key set of applications including network traffic accounting, usage-based network billing, network planning, as well as DoS (denial of services) monitoring capabilities, network monitoring, outbound marketing, and data mining capabilities.
network address	Network layer address referring to a logical, rather than a physical, network device. Also called a protocol address.
network layer	Layer 3 of the OSI reference model. This layer provides connectivity and path selection between two end systems. The network layer is the layer at which routing occurs.
Network Management module	A module within the Enterprise Composite Network model. This module performs intrusion detection logging, system logging, and Terminal Access Controller Access Control System Plus (TACACS+)/RADIUS and One-Time Password (OTP) authentication, as well as network monitoring and general configuration management functions.
NIC	Network Interface Card. Board that provides network communication capabilities to and from a computer system. Also called an adapter.
node	Endpoint of a network connection or a junction common to two or more lines in a network. Nodes can be processors, controllers, or workstations. Nodes, which vary in routing and other functional capabilities, can be interconnected by links, and serve as control points in the network. Node sometimes is used generically to refer to any entity that can access a network, and frequently is used interchangeably with device.
OSI	Open System Interconnection. International standardization program created by ISO and ITU-T to develop standards for data networking that facilitate multivendor equipment interoperability.

Acronym or Term	Definition
OSI reference model	Open System Interconnection reference model. Network architectural model developed by ISO and ITU-T. The model consists of seven layers, each of which specifies particular network functions, such as addressing, flow control, error control, encapsulation, and reliable message transfer. The lowest layer (the physical layer) is closest to the media technology. The highest layer (the application layer) is closest to the user application. The OSI reference model is used universally as a method for teaching and understanding network functionality.
packet	Logical grouping of information that includes a header containing control information and (usually) user data. Packets most often are used to refer to network layer units of data. The terms datagram, frame, message, and segment also are used to describe logical information groupings at various layers of the OSI reference model and in various technology circles.
packet sniffer	Device that monitors traffic on a network and reports on problems on the network.
PIM	Protocol Independent Multicast. Multicast routing protocol that allows the addition of IP multicast routing on existing IP networks. PIM is unicast routing protocol-independent and can be operated in two modes: dense and sparse.
PIM dense mode	One of the two PIM operational modes. PIM dense mode is data-driven and resembles typical multicast routing protocols. Packets are forwarded on all outgoing interfaces until pruning and truncation occur. In dense mode, receivers are densely populated, and it is assumed that the downstream networks want to receive and will probably use the datagrams that are forwarded to them. The cost of using dense mode is its default flooding behavior. Sometimes called dense mode PIM or PIM dense mode.
PIM sparse mode	One of the two PIM operational modes. PIM sparse mode tries to constrain data distribution so that a minimal number of routers in the network receive it. Packets are sent only if they are explicitly requested at the rendezvous point (RP). In sparse mode, receivers are widely distributed, and the assumption is that downstream networks will not necessarily use the datagrams that are sent to them. The cost of using sparse mode is its reliance on the periodic refreshing of explicit join messages and its need for RPs. Sometimes called sparse mode PIM or PIM SM.
ping	An application that uses an ICMP echo message and its reply. Often used in IP networks to test the reachability of a network device.
policy domain	A collection of networks under single management.
POP	point of presence. A physical location within a service provider network where users dial in. Post Office Protocol. Internet application protocol providing e-mail services. An Internet Standard, POP is defined by RFC 1939.
PortFast	Feature used on switched ports where only end-user stations are directly connected. There is no delay in passing traffic, because the switch immediately puts the port to the forward state.
pps	packets per second.
PQ	priority queuing. Queue management and service discipline that prioritizes traffic at a network interface. Four traffic priorities can be configured. A series of filters based on packet characteristics (source IP address and port) is defined to cause the router to place critical traffic in the highest queue and other traffic in the lower three queues. The queue with the highest priority is serviced first until empty; the lower queues are then serviced in sequence. It is possible for higher-priority traffic to starve lower-priority traffic by consuming all the bandwidth.
PQ/CBWFQ	priority queuing/class-based weighted fair queuing. Feature that joins strict priority queuing and CBWFQ. Strict priority queuing allows delay-sensitive data, such as voice, to be dequeued from a single priority queue and sent first (before packets in other queues are dequeued), giving delay-sensitive data preferential treatment over other traffic. It is also called low latency queuing (LLQ).

Acronym or Term	Definition
PQ-WFQ	priority queuing-weighted fair queuing. Also called IP Real-Time Transport Protocol (RTP) Priority. Queuing mechanism that provides a strict priority queuing scheme for delay-sensitive data such as voice.
PVST+	Per VLAN Spanning Tree+. A mechanism developed by Cisco to allow running several STP instances (even over an 802.1q network). Each SPT instance has one or more VLAN(s) mapped to it.
QoS	quality of service. The intent for a transmission system to deliver guaranteed, differentiated services by giving network resource and usage control to the network operator.
queue	A data structure used to store frames or packets waiting for service. Typically the service is forwarding.
queuing delay	Amount of time that a data packet must wait in a queue before it can be transmitted onto a statistically multiplexed physical circuit.
RADIUS	Remote Authentication Dial-In User Service. Responsible for receiving user connection requests, authenticating the user, and then returning all configuration information necessary for the client to deliver service to the user. It includes a database for authenticating connections and for tracking connection time. RADIUS is defined in RFC 2865.
RED	random early detection. Congestion avoidance algorithm in which some percentage of packets are dropped when congestion is detected and before the queue in question overflows completely.
redundancy	In internetworking, the duplication of devices, services, or connections so that, in the event of a failure, the redundant devices, services, or connections can perform the work of those that failed.
Remote Access and VPN module	A module within the Enterprise Edge functional area of the Enterprise Composite Network model. This module terminates VPN traffic, forwarded by the Internet Connectivity module, from remote users and remote sites.
rendezvous point	Router specified in PIM sparse mode implementations to track membership in multicast groups and to forward messages to known multicast group addresses.
RFC	Request For Comments. Document series used as the primary means for communicating information about Internet protocols and related technical details. Some RFCs are designated by the IAB as Internet standards. Most RFCs document protocol specifications, such as Telnet and FTP, but some are humorous or historical. RFCs are available online from numerous sources.
root bridge	The root, or start, of the spanning tree in a switched network. It exchanges topology information with designated bridges in a spanning-tree instance and notifies all other bridges in the network when topology changes are required. This exchange prevents loops and provides a measure of defense against link failure.
RPF	Reverse Path Forwarding. Multicasting technique in which a multicast datagram is forwarded out of all but the receiving interface if the receiving interface is the one used to forward unicast datagrams to the source of the multicast datagram.
RSVP	Resource Reservation Protocol. Protocol that supports the reservation of resources across an IP network. Applications running on IP end systems can use RSVP to indicate to other nodes the nature (bandwidth, jitter, maximum burst, and so on) of the packet streams they want to receive. Also known as Resource Reservation Setup Protocol.
RTCP	RTP Control Protocol. RTCP carries control information about an RTP session, such as the amount of data transmitted and receipt acknowledgements. RTP uses this information to request the session source adapt accordingly.

Acronym or Term	Definition
RTP	Real-Time Transport Protocol. Protocol designed to provide end-to-end network transport functions for applications transmitting real-time data, such as audio, video, or simulation data, over multicast or unicast network services. RTP provides such services as payload type identification, sequence numbering, time stamping, and delivery monitoring to real-time applications.
Server Farm module	A module within the Enterprise Composite Network model. It contains servers providing application, file, print, e-mail, and Domain Name System (DNS) services to internal users.
Service Provider Edge	A functional area described within the Enterprise Composite Network model. The modules in this area are not implemented by the enterprise itself, but are necessary to enable communication with other networks. It most often uses different WAN technologies provided by SPs.
SNMP	Simple Network Management Protocol. Network management protocol used almost exclusively in TCP/IP networks. SNMP provides a means to monitor and control network devices, and to manage configurations, statistics collection, performance, and security.
SONET	Synchronous Optical Network. A standard format for transporting a wide range of digital telecommunications services over optical fiber. SONET is characterized by standard line rates, optical interfaces, and signal formats.
SP	service provider.
SPAN	Switched Port Analyzer. Feature of a Catalyst switch that extends the monitoring capabilities of existing network analyzers into a switched Ethernet environment. SPAN mirrors the traffic at one switched segment onto a predefined SPAN port. A network analyzer attached to the SPAN port can monitor traffic from any of the other Catalyst switched ports.
SPF	Shortest Path First algorithm or Dijkstra algorithm. Routing algorithm that iterates on length of path to determine a shortest-path spanning tree. Commonly used in link-state routing algorithms, it runs on every routing device in the network.
STP	Spanning Tree Protocol. Bridge protocol that uses the spanning-tree algorithm, enabling a learning bridge to dynamically work around loops in a network topology by creating a spanning tree. Bridges exchange BPDU messages with other bridges to detect loops, and then remove the loops by shutting down selected bridge interfaces. Refers to both the IEEE 802.1 Spanning Tree Protocol standard and the earlier Digital Equipment Corporation Spanning Tree Protocol upon which it is based. The IEEE version supports bridge domains and allows the bridge to construct a loop-free topology across an extended LAN. The IEEE version generally is preferred over the Digital Equipment version. Sometimes abbreviated as STP.
switch	Network device that filters, forwards, and floods frames based on the destination address of each frame. The switch operates at the data link layer of the OSI model. General term applied to an electronic or mechanical device that allows a connection to be established as necessary and terminated when there is no longer a session to support. In telephony, a general term for any device, such as a PBX, that connects individual phones to phone lines. See also PBX and PSTN.
switching	Process of taking an incoming frame from one interface and delivering to another interface for transmission. Routers use Layer 3 switching to route a packet, and traditional LAN switches use Layer 2 switching to forward frames. See also Layer 2 switching and Layer 3 switching.
TACACS+	Terminal Access Controller Access Control System Plus. Authentication protocol extended by Cisco that provides remote access authentication and related services, such as event logging. User passwords are administered in a central database rather than in individual routers, providing a scalable network security solution.

Acronym or Term	Definition
TCP	Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission. TCP is part of the TCP/IP protocol stack.
TDM	Time-Division Multiplexing. Technique in which information from multiple channels can be allocated bandwidth on a single wire based on preassigned time slots.
Telnet	Standard terminal emulation protocol in the TCP/IP protocol stack. Telnet is used for remote terminal connection, enabling users to log in to remote systems and use resources as if they were connected to a local system. Telnet is defined in RFC 854.
terminal server	Communications processor that connects asynchronous devices, such as terminals, printers, hosts, and modems, to any LAN or WAN that uses TCP/IP. Terminal servers provide the internetwork intelligence that is not available in the connected devices.
TFTP	Trivial File Transfer Protocol. Simplified version of FTP that allows files to be transferred from one computer to another over a network, usually without the use of client authentication (for example, username and password). TFTP is defined in RFC 1350.
topology	Physical arrangement of network nodes and media within an enterprise networking structure.
ToS	type of service. The type of service octet in the Internet Protocol (IP) header is defined in RFC 1349, which has been made obsolete by RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers. Now called the DS (for differentiated services) field for both IPv4 and IPv6. In IPv4, it defines the layout of the ToS octet. In IPv6, it defines the traffic-class octet. A base set of packet forwarding treatments, or per-hop behaviors, is defined based on the information in the DS field.
traffic management	Techniques for avoiding congestion and shaping and policing traffic. Allows links to operate at high levels of utilization by scaling back lower-priority, delay-tolerant traffic at the edge of the network when congestion begins to occur.
traffic policing	Process used to measure the actual traffic flow across a given connection and compare it to the total admissible traffic flow for that connection. Traffic outside of the agreed upon flow can be discarded immediately or tagged (where some field is changed) and discarded en route if congestion develops. Traffic policing is used in ATM, Frame Relay, and other types of networks. Also known as admission control, permit processing, and rate enforcement.
traffic shaping	Use of queues to smooth surges that can congest a network. Data is buffered and then sent into the network in regulated amounts to ensure that the traffic fits within the promised traffic envelope for the particular connection. Traffic shaping is used in ATM, Frame Relay, and other types of networks. Also known as metering, shaping, and smoothing. Most often configured on egress ports to insure compliance with agreed connection traffic rates to avoid traffic policing, it is frequently implemented using a token bucket algorithm.
tunneling	A dual encapsulation mechanism by which a protocol at some layer in the protocol stack is transported by another protocol operating at the same layer.
twisted pair	Twisted pair describes copper media in which the wires are twisted around each other in a spiral to reduce crosstalk or electromagnetic induction between the pairs of wires. The ordinary copper wire that connects homes and many business computers to the PSTN uses a single pair for each analog telephone line.
UDP	User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768.
unicast	Traffic from a single source sent to a single network destination.

Acronym or Term	Definition
UplinkFast	A spanning-tree maintenance mechanism that enables the switch to put a redundant path (port) into active state within a second.
UTP	unshielded twisted pair. A four-pair wire medium used in a variety of networks. UTP does not require the fixed spacing between connections that is necessary with coaxial-type connections.
VACL	VLAN access control list. A VACL contains an ordered list of access control entries (ACEs).
VLAN	virtual local-area network. A group of devices on one or more LANs that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.
VTP	VLAN Trunk Protocol. VTP reduces administration in a switched network by distributing VLAN information through all switches in the VTP domain. VTP is a Cisco proprietary protocol that is available on most of the Cisco Catalyst family products.
vty	virtual type terminal. Also commonly used as virtual terminal lines.
WAN module	A module within the Enterprise Edge functional area of the Enterprise Composite Network model. The WAN module includes all WAN technologies that provide circuits between geographically separated locations. FR, ATM, and PPP are frequently encountered data-link technologies.
web	World Wide Web (also called WWW). A client/server system based on HTML and HTTP.
WFQ	weighted fair queuing. Queuing algorithm that identifies conversations (in the form of traffic streams), separates packets that belong to each conversation, and ensures that capacity is shared fairly between these individual conversations. WFQ is an automatic way of stabilizing network behavior during congestion and results in increased performance and reduced retransmission. It is the default on serial interfaces at and below 2.048 Mbps.
wiring closet	Specially designed room used for wiring a data or voice network. Wiring closets serve as a central junction point for the wiring and the wiring equipment that is used for interconnecting devices. They are sometimes called distribution facilities.
workgroup	Collection of workstations and servers on a LAN that are designed to communicate and exchange data with one another.
WRED	weighted random early detection. Queuing method that ensures that high-precedence traffic has lower loss rates than other traffic during times of congestion, by dropping some percentage of packets when congestion is detected and before the queue in question overflows. The drop probability can be configured differently for each of multiple traffic classes.