

Building Cisco Multilayer Switched Networks

Volume 1

Version 2.2

Student Guide

CLS Production Services: 08.05.05

Copyright © 2005, Cisco Systems, Inc. All rights reserved.

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

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica
Croatia • Cyprus • 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



Copyright © 2005 Cisco Systems, Inc. All rights reserved. CCSP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0501R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS." CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

Table of Contents

Volume 1

<u>Course Introduction</u>	1
Overview	1
Learner Skills and Knowledge	1
Course Goal and Objectives	2
Course Flow	3
Additional References	4
Cisco Glossary of Terms	4
Your Training Curriculum	5
CCNP Career Certifications	5
<u>Designing a Network Using the Campus Infrastructure Module</u>	1-1
Overview	1-1
Module Objectives	1-1
<u>Describing the Campus Infrastructure Module</u>	1-3
Overview	1-3
Objectives	1-3
Devices in a Nonhierarchical Network	1-4
Layer 2 Network Issues	1-6
Routed Network Issues	1-7
What Is a Multilayer Switch?	1-8
Issues with Multilayer Switches in a Nonhierarchical Network	1-10
Enterprise Composite Network Model	1-11
Enterprise Composite Network Model Functional Areas	1-12
Enterprise Composite Network Model Benefits	1-13
Benefits of the Enterprise Composite Network Model	1-14
Modules of the Enterprise Campus	1-15
Campus Infrastructure Module	1-16
Summary	1-18
<u>Deploying Technology in the Campus Infrastructure Module</u>	1-19
Overview	1-19
Objectives	1-19
Issues in a Poorly Designed Network	1-20
Designing a Hierarchical IP Addressing Scheme	1-22
Guidelines for Applying IP Address Space in the Enterprise Network	1-23
Interconnection Technologies	1-24
Determining Equipment and Cabling Needs	1-26
References	1-27
Mapping VLANs in a Hierarchical Network	1-28
Traffic Types	1-29
Considering Traffic Source-to-Destination Path	1-31
Cisco Catalyst Configuration Interfaces	1-32
Catalyst Software Interface	1-32
Example: Using Catalyst Software Commands	1-32
Cisco IOS Interface	1-33
Example: Using IOS Commands	1-33
Configuration Interface Available on Various Catalyst Platforms	1-34
Summary	1-35
Module Summary	1-37
References	1-37
Module Self-Check	1-39
Module Self-Check Answer Key	1-40
<u>Defining VLANs</u>	2-1
Overview	2-1

Module Objectives	2-1
<i>Implementing VLANs</i>	2-3
Overview	2-3
Objectives	2-3
What Is an End-to-End VLAN?	2-4
Example: VLAN Implementation	2-5
What Is a Local VLAN?	2-6
VLAN Configuration Modes	2-8
VLAN Database Mode	2-9
Example: Creating a VLAN in VLAN Database Mode	2-9
What Are VLAN Access Ports?	2-10
Benefits of Local VLANs in the Enterprise Composite Network Model	2-12
VLAN Implementation Commands	2-14
How to Implement a VLAN	2-16
1. Create or Configure a VLAN	2-17
2. Verify VLAN Configuration	2-18
3. Associate Switch Ports with the VLAN	2-19
4. Verify Switch Port Configuration	2-19
5. Test VLAN Connectivity	2-20
6. Implement Switch and VLAN Security Measures	2-20
Summary	2-21
<i>Supporting Multiple VLANs on a Single Trunk</i>	2-23
Overview	2-23
Objectives	2-23
What Is a VLAN Trunk?	2-24
What Is a VLAN Trunking Protocol?	2-25
Comparing ISL and 802.1Q Trunking Protocols	2-26
ISL Trunking Protocol	2-27
ISL Encapsulation Process	2-28
ISL Header	2-28
ISL Trailer	2-29
References	2-30
802.1Q Trunking Protocol	2-31
802.1Q Tagging Process	2-32
What Is an 802.1Q Native VLAN?	2-33
Example: Native VLAN Implementation—Two End Devices on the Same Switch Port	2-34
Issues with 802.1Q Native VLANs	2-35
VLAN Ranges	2-36
Identifying the Modes for Dynamic Trunking Protocol	2-38
Trunking Configuration Commands	2-39
How to Configure Trunking	2-40
Configuring an ISL Trunk	2-41
Configuring a Port for ISL Trunking with No DTP	2-42
Verifying the ISL Trunk Configuration	2-43
Configuring an 802.1Q Trunk	2-44
Example: Configuring a Port for 802.1Q Trunking	2-45
Verify the 802.1Q Configuration	2-46
Example: Configure and Display Port Information for an 802.1Q Dynamic Trunk Link	2-47
Example: Displaying Trunk Information for 802.1Q Trunking	2-47
Using Trunking Protocols in the Campus Infrastructure Module	2-48
Resolving Trunk Link Problems	2-49
Summary	2-50
<i>Propagating VLAN Information with VTP</i>	2-51
Overview	2-51
Objectives	2-51
What Is a VTP Domain?	2-52
What Is the VTP Protocol?	2-53

VTP in the Campus Infrastructure Module	2-53
References	2-54
VTP Modes	2-55
Describing VTP Operation	2-57
Configuration Revision Number	2-58
VTP Advertisement Types	2-58
VTP Versions	2-58
References	2-59
VTP Configuration Commands	2-60
How to Configure a VTP Management Domain	2-62
Configuring VTP on a Switch	2-63
Verifying the VTP Configuration	2-65
VTP Counters	2-66
Common Problems with VTP Configuration	2-66
Best Practices: Configuring Switches in a VTP Domain	2-68
How to Add a New Switch to an Existing VLAN	2-69
Summary	2-71
Module Summary	2-73
References	2-73
Module Self-Check	2-75
Module Self-Check Answer Key	2-76
<i>Implementing Spanning Tree</i>	<i>3-1</i>
Overview	3-1
Module Objectives	3-1
<i>Defining the Spanning Tree Protocol</i>	<i>3-3</i>
Overview	3-3
Objectives	3-3
Transparent Bridges	3-4
Identifying Traffic Loops	3-5
Example: Flooded Unicast Frames and Bridge Loops	3-5
Preventing Loops on a Layer 2 Network	3-6
802.1D Spanning Tree Protocol	3-7
Spanning Tree Communication	3-8
What Is a Root Bridge?	3-9
BPDU Fields Associated with Root Bridge Selection	3-10
Bridge ID Field in the BPDU	3-11
Identifying the Root Selection Process	3-12
802.1D Port Roles	3-13
Forming an Association with the Root Bridge	3-14
Path Cost	3-15
Selecting the Root Port	3-16
Selecting the Designated Port	3-17
Example: Determining the Active Topology	3-18
Summary	3-19
<i>Maintaining and Configuring STP</i>	<i>3-21</i>
Overview	3-21
Objectives	3-21
Identifying Spanning Tree Port States and Timers	3-22
Spanning Tree Timers	3-23
Identifying Topology Changes	3-24
What Is a Backup Root Bridge?	3-25
Priority Commands	3-26
How to Configure a Root Bridge	3-27
Comparing CST and PVST	3-28
Example: Comparing CST and PVST	3-28
Summary	3-29
References	3-29

Configuring PortFast **3-31**

Overview	3-31
Objectives	3-31
What Is PortFast	3-32
PortFast Configuration Commands	3-33
How to Configure PortFast	3-34
Configure PortFast	3-34
Verify PortFast	3-34
Summary	3-35

Guarding Against Rogue STP Root Bridges **3-37**

Overview	3-37
Objectives	3-37
Protecting Spanning Tree	3-38
BPDU Guard	3-38
BPDU Filtering	3-38
BPDU Root Guard	3-38
BPDU Guard Configuration Commands	3-39
BPDU Filtering Applied Globally Versus Per-Port	3-39
Configuring BPDU Guard	3-39
Verifying BPDU Guard	3-40
BPDU Filtering Configuration Commands	3-41
BPDU Filtering Applied Globally Versus Per-Port	3-41
Configuring BPDU Filtering	3-42
Root Guard	3-43
Example: Using Root Guard	3-43
Root Guard Configuration Commands	3-45
How to Configure Root Guard	3-46
Configuring Root Guard	3-46
Verifying Root Guard	3-47
Summary	3-48

Configuring UplinkFast **3-49**

Overview	3-49
Objectives	3-49
What Is a Link Fault?	3-50
UplinkFast	3-51
UplinkFast Configuration Commands	3-52
How to Configure UplinkFast	3-53
Configuring UplinkFast	3-53
Verifying UplinkFast	3-54
Summary	3-56

Configuring BackboneFast **3-57**

Overview	3-57
Objectives	3-57
What Are Indirect Link Failures?	3-58
BackboneFast	3-59
Example: BackboneFast Operation	3-60
BackboneFast Configuration Commands	3-61
How to Configure BackboneFast	3-62
Configure BackboneFast	3-62
Verify BackboneFast	3-62
Summary	3-63
References	3-63

Configuring EtherChannel **3-65**

Overview	3-65
Objectives	3-65

EtherChannel	3-66
PAgP and LACP Protocols	3-67
Interface Modes	3-67
EtherChannel Configuration Commands	3-69
Configuring Port Channels Using EtherChannel	3-71
Configuring Layer 3 Etherchannel	3-72
Configure EtherChannel	3-72
Verifying EtherChannel	3-73
Example: Verifying Port-Channel Configuration	3-76
Load Balancing over EtherChannel	3-79
Configuring and Verifying EtherChannel Load Balancing	3-79
Guidelines and Best Practices for Configuring EtherChannel	3-81
Summary	3-83
Module Summary	3-85
References	3-85
Module Self-Check	3-87
Module Self-Check Answer Key	3-89

Enhancing Spanning Tree **4-1**

Overview	4-1
Module Objectives	4-1

Troubleshooting Spanning Tree **4-3**

Overview	4-3
Objectives	4-3
STP Problems	4-4
Duplex Mismatch	4-5
Unidirectional Link Failure	4-7
Frame Corruption	4-7
Resource Errors	4-7
PortFast Configuration Error	4-8
EtherChannel Issues	4-8
Spanning Tree debug Commands	4-9
How to Troubleshoot STP Problems	4-10
Refer to a Network Diagram	4-10
Identify Issues	4-10
Restore Connectivity Versus Resolve Issues	4-11
Check Ports	4-11
Look for Resource Errors	4-12
Disable Unneeded Features	4-12
STP debug Command	4-12
General Recommendations	4-12
Summary	4-13
References	4-13

Preventing STP Forwarding Loops **4-15**

Overview	4-15
Objectives	4-15
Unidirectional Link Detection	4-16
Loop Guard	4-17
Example: Before Loop Guard	4-18
Example: With Loop Guard	4-19
References	4-19
How to Prevent STP Failures Due to Unidirectional Links	4-20
Configuring UDLD and Loop Guard	4-21
Configuring UDLD	4-22
Verifying and Resetting UDLD	4-23
Example: Displaying the UDLD State	4-23
Configuring Loop Guard	4-25
Summary	4-27

References	4-27
<i>Implementing RSTP</i>	4-29
Overview	4-29
Objectives	4-29
Rapid Spanning Tree Protocol	4-30
RSTP Port States	4-31
RSTP Port Roles	4-32
What Are Edge Ports?	4-34
RSTP Link Types	4-35
Examining the RSTP BPDU	4-36
Identifying the RSTP Proposal and Agreement Process	4-37
Downstream RSTP Proposal Process	4-38
Identifying the RSTP Topology Change Notification Process	4-39
RSTP Implementation Commands	4-41
How to Implement RSTP	4-42
Explanation: Enabling PVST	4-42
Verifying the Rapid PVST Configuration	4-43
Summary	4-44
<i>Implementing MST</i>	4-45
Overview	4-45
Objectives	4-45
What Is MST?	4-46
MST Regions	4-48
Extended System ID	4-50
References	4-50
Interacting Between MST Regions and 802.1D Networks	4-51
MST Implementation Commands	4-53
How to Configure and Verify MST	4-55
Example: Displaying MST Configuration Information	4-55
Example: Displaying General MST Information	4-56
Example: Displaying MST Information for a Specific Instance	4-57
Example: Displaying MST Information for a Specific Interface	4-58
Example: Displaying MST Information for a Specific Instance and Interface	4-58
Example: Displaying Detailed MST Information	4-58
Summary	4-60
References	4-60
Module Summary	4-61
References	4-61
Module Self-Check	4-63
Module Self-Check Answer Key	4-64
<i>Implementing Multilayer Switching</i>	5-1
Overview	5-1
Module Objectives	5-1
<i>Describing Routing Between VLANs</i>	5-3
Overview	5-3
Objectives	5-3
Inter-VLAN Routing Using Multiple Interfaces on an External Router	5-4
External Router with Multiple Interface: Advantages and Disadvantages	5-5
Inter-VLAN Routing Using an External Router and a Single Trunk	5-6
External Router with Single Interface: Advantages and Disadvantages	5-8
Inter-VLAN Routing Using External Router Configuration Commands	5-9
How to Configure Inter-VLAN Routing Using an External Router	5-10
Configuring an External Router using ISL Encapsulation	5-10
Configuring an External Router using 802.1Q	5-11
Verifying the Inter-VLAN Routing Configuration using Ping	5-12

Verifying the Inter-VLAN Routing Configuration	5-13
Example: Displaying Inter-VLAN Configuration Information	5-13
Example: Displaying Routing Table Information	5-14
Summary	5-15
<i>Deploying CEF-Based Multilayer Switching</i>	5-17
Overview	5-17
Objectives	5-17
What Is Layer 2 Switching?	5-18
What Are Layer 2 Switching Tables?	5-19
Identifying the Layer 2 Switch Forwarding Process	5-20
What Is Multilayer Switching?	5-21
References	5-22
What Is a CEF-Based Multilayer Switch?	5-23
Identifying the Multilayer Switch Packet Forwarding Process	5-25
CEF-Based Tables and MLS Lookups	5-26
FIB Table Updates	5-26
Ternary Content Addressable Memory Table	5-28
ARP Throttling	5-30
CEF-Based MLS Operation	5-32
Frame Rewrite Using CEF	5-33
Configuring and Verifying CEF	5-34
Verifying CEF	5-36
Verifying Layer 3 Switching	5-37
Display CEF Statistics	5-38
Displaying Detailed Adjacency Information	5-39
Debugging CEF Operations	5-40
Common CEF Problems and Solutions	5-42
How to Troubleshoot Layer 3 Connectivity in a CEF-based Multilayer Switch	5-44
Troubleshoot Host Connectivity Using CEF	5-44
Summary	5-48
References	5-48
<i>Enabling Routing Between VLANs on a Multilayer Switch</i>	5-49
Objectives	5-49
Layer 3 Switch Virtual Interface	5-50
Routed Interfaces on a Multilayer Switch	5-51
Configuration Commands for Inter-VLAN Communication on a Multilayer Switch	5-52
How to Configure Inter-VLAN Routing on a Multilayer Switch	5-53
Summary	5-54
Module Summary	5-55
References	5-55
Module Self-Check	5-57
Module Self-Check Answer Key	5-58
<i>Implementing Redundancy in the Routing Layer</i>	6-1
Overview	6-1
Module Objectives	6-1
<i>Configuring Layer 3 Redundancy with HSRP</i>	6-3
Overview	6-3
Objectives	6-3
Identifying the Router Redundancy Process	6-4
Routing Issues	6-5
Using Default Gateways	6-5
Using Proxy ARP	6-6
Hot Standby Router Protocol	6-7
Identifying HSRP Operations	6-8
Virtual HSRP Router	6-8

Active HSRP Router	6-8
ARP Resolution with HSRP	6-9
Standby and Other HSRP Routers in the Group	6-10
HSRP Active and Standby Router Interaction	6-11
HSRP States	6-12
HSRP Initial State	6-13
HSRP Listen State	6-14
HSRP Speak State	6-15
Standby State	6-16
Active State	6-17
HSRP Configuration Commands	6-18
How to Enable HSRP	6-19
Configure HSRP Group on an Interface	6-20
Verifying HSRP Configuration	6-21
Establish HSRP Priorities	6-22
Verify the HSRP Standby Priority	6-23
Verify All HSRP Operations	6-23
Summary	6-24

Optimizing HSRP **6-25**

Overview	6-25
Objectives	6-25
Load Sharing	6-26
Addressing HSRP Groups Across Trunk Links	6-27
Supporting Multiple Subnets with Multiple HSRP Groups	6-28
HSRP Optimization Options	6-29
HSRP Standby Preempt	6-29
HSRP Hello Message Timer Adjustment	6-29
HSRP Interface Tracking	6-29
HSRP Standby Preempt	6-30
Example: Displaying HSRP Preempt	6-30
Hello Message Timers	6-32
HSRP Interface Tracking	6-33
Configuring HSRP Tracking	6-35
Tuning HSRP Operations	6-36
Subsecond Failover	6-36
Preempt Time Aligned with Router Boot Time	6-36
HSRP debug Commands	6-38
References	6-38
How to Debug HSRP Operations	6-39
Example: HSRP Debugging on Negotiation for Role of Active Router	6-39
Example: HSRP Debugging on First and Only Router on Subnet	6-40
Example: HSRP on NonPreempt Configured Router Coming Up	6-42
Example: HSRP on Preempt-Configured Router Coming Up	6-44
Summary	6-46
References	6-46

Configuring Layer 3 Redundancy with VRRP and GLBP **6-47**

Overview	6-47
Objectives	6-47
Virtual Router Redundancy Protocol	6-48
Identifying the VRRP Operations Process	6-50
Gateway Load Balancing Protocol	6-51
Identifying the GLBP Operations Process	6-52
VRRP and GLBP Configuration Commands	6-56
How to Enable VRRP and GLBP	6-58
VRRP Implementation	6-58
GLBP Implementation	6-59
Summary	6-60
References	6-60

Implementing Hardware and Software Redundancy on Modular Switches **6-61**

Overview	6-61
Objectives	6-61
What Is RPR+?	6-62
Redundant Supervisor Engine Configuration Commands	6-63
How to Implement Redundant Supervisor Engines	6-64
Cisco Catalyst 6500 Switch	6-65
What Is Stateful Switchover?	6-66
What Is Single Router Mode?	6-67
Failure with SRM and SSO	6-68
How to Configure and Verify SRM with SSO	6-69
Configure the MSFCs for SRM with SSO	6-69
Verify SRM Configuration and Operation	6-70
What Is Nonstop Forwarding?	6-72
Identifying NSF-Aware Protocols	6-74
EIGRP Operation	6-74
BGP Operation	6-75
OSPF Operation	6-75
IS-IS Operation	6-75
Failover with NFS and SSO	6-76
How to Configure NSF	6-77
Example: NSF Configuration for EIGRP	6-77
Redundant Power Supply Configuration	6-79
How to Configure Redundant Power Supplies	6-80
Turn Off or Cycle Power to Modules	6-81
Summary	6-82

Designing High Availability in a Multilayer Switch **6-83**

Overview	6-83
Objectives	6-83
What Is Redundancy in a Switched Network?	6-84
Benefits and Drawbacks of Device-Level Fault Tolerance	6-85
Benefits and Drawbacks of Redundant Network Topology	6-86
Redundancy with Stacked Switches	6-88
Layer 3 Failure with Stacked Switches	6-89
Loopback Cable to Maintain Layer 2 Path	6-90
High Availability: Access Layer Best Practices	6-91
High Availability: Distribution Layer Best Practices	6-93
Layers 2 and 3 Redundancy Alignment	6-95
Affect of Layer 3 Failure with Autostate	6-97
High Availability: Core Layer Best Practices	6-98
Summary	6-99
Module Summary	6-101
References	6-101
Module Self-Check	6-103
Module Self-Check Answer Key	6-104

Minimizing Service Loss and Data Theft in a Switched Network **7-1**

Overview	7-1
Module Objectives	7-1

Understanding Switch Security Issues **7-3**

Overview	7-3
Objectives	7-3
Switch Security Concerns	7-4
Switch Attack Categories	7-5
Describing a MAC Flooding Attack	7-7
Suggested Mitigation for MAC Flood Attacks	7-8
Describing Port Security	7-9

References	7-10
Port Security Configuration Commands	7-11
How to Configure Port Security on a Switch	7-12
Caveats to Port Security Configuration Steps	7-13
How to Verify Port Security	7-14
Verifying Network Access Security	7-14
Example: show port-security Command Output	7-15
Example: show port-security Command for a Specific Interface	7-16
Example: Displaying MAC Address Table Security Information	7-17
Port Security with Sticky MAC Addresses	7-18
Summary	7-19
References	7-19

Mitigating VLAN Attacks **7-21**

Overview	7-21
Objectives	7-21
What Is VLAN Hopping?	7-22
Switch Spoofing	7-22
Double Tagging	7-24
How to Mitigate VLAN Hopping	7-25
What Is a Private VLAN?	7-26
PVLAN Port Types	7-27
Resources	7-27
Configuring PVLANS	7-29
Example: PVLAN Configurations	7-30
Configuring VLAN Security Using Access Lists	7-31
Summary	7-34
References	7-34

Mitigating Spoof Attacks **7-35**

Overview	7-35
Objectives	7-35
Describing a DHCP Spoof Attack	7-36
Describing DHCP Snooping	7-37
DHCP Snooping Configuration Commands	7-38
References	7-38
How to Configure DHCP Snooping	7-39
Verifying the DHCP Snooping Configuration	7-40
Describing a MAC Spoof Attack	7-41
Describing ARP Spoofing	7-42
What Is Dynamic ARP Inspection?	7-44
References	7-45
How to Configure Dynamic ARP Inspection	7-46
Example: DAI Implementation	7-47
Summary	7-48
References	7-48

Implementing AAA **7-49**

Overview	7-49
Objectives	7-49
Authentication, Authorization, and Accounting	7-50
Describing the AAA Process	7-52
Authentication and Authorization Methods	7-54
Authorization Methods	7-55
Configuring AAA	7-56
Configuring Authentication	7-57
Configuring Authorization	7-59
Configuring Accounting	7-62
Comprehensive AAA Configuration Example	7-64
802.1X Port-Based Authentication	7-65

Configuring 802.1X Port-Based Authentication	7-67
Example	7-68
Reference	7-68
Summary	7-69
References	7-69

Defending Network Switches **7-71**

Overview	7-71
Objectives	7-71
CDP Security Issues	7-72
Vulnerabilities in Telnet	7-73
VTY ACLs	7-74
Commands to Configure VTY ACLs	7-75
Example: VTY Access	7-75
Secure Shell Protocol	7-76
Best Practices: Switch Security Considerations	7-77
Organizational Security Policies	7-77
Secure Switch Devices	7-78
Secure Switch Protocols	7-80
Mitigating Compromises Launched Through a Switch	7-81
Capturing Traffic in a Switched Network	7-82
Capturing Data in a Switched Network	7-82
Commands Used in Capturing Network Traffic	7-83
Configuring SPAN on a Local 3500XL	7-84
Resources	7-84
Monitoring Performance with RSPAN	7-85
RSPAN Guidelines	7-85
Configuring RSPAN	7-87
Summary	7-88
Module Summary	7-89
References	7-89
Module Self-Check	7-91
Module Self-Check Answer Key	7-92

Configuring Campus Switches to Support Voice and Video Applications **8-1**

Overview	8-1
Module Objectives	8-1

Accommodating Voice Traffic on Campus Switches **8-3**

Overview	8-3
Objectives	8-3
Voice Traffic on a Cisco Infrastructure	8-4
Benefits of IP Telephony on a Cisco Infrastructure	8-4
What Is a Voice VLAN?	8-6
Voice Considerations in Campus Submodules	8-7
Building Access Submodule	8-7
Building Distribution Submodule	8-8
Network Design Considerations for Voice	8-9
General Design Considerations	8-9
Bandwidth Provisioning	8-10
Power Considerations	8-11
Intelligent Network Services	8-12
QoS Basics	8-13
QoS and Voice Traffic in the Campus Module	8-15
Network Availability Problem Areas	8-16
QoS Trust Boundaries	8-18
QoS Traffic Classification and Marking	8-19
Layer 2 QoS Marking	8-20
Layer 3 QoS Marking	8-21
Basic Switch Commands to Support Attachment of a Cisco IP Phone	8-22

How to Configure a Switch for Attachment of a Cisco IP Phone	8-24
Example	8-25
Summary	8-26

Configuring IP Multicast **8-27**

Overview	8-27
Objectives	8-27
IP Multicast	8-28
IP Multicast Group Membership	8-30
IP Multicast Address Structure	8-31
IP Multicast to MAC Address Mapping	8-32
IP Multicast Address Ranges	8-33
Reserved Link-Local Addresses	8-33
Globally Scoped Addresses	8-34
Source-Specific Multicast Addresses	8-34
GLOP Addresses	8-34
Limited Scope Addresses	8-34
What Is RPF?	8-35
Source Distribution Trees	8-35
Shared Distribution Trees	8-36
Source Trees Versus Shared Trees	8-37
Reverse Path Forwarding Check	8-38
What Is PIM?	8-40
PIM Versions 1 and 2	8-41
References	8-41
PIM Modes	8-42
PIM Sparse Mode	8-43
What Is IGMP?	8-45
IGMP Message Format	8-46
IGMP v 3 Report Message	8-46
IGMP v 3 Query Message	8-48
Describing the IGMP Snooping Process	8-49
IP Multicast Configuration Commands	8-50
How to Enable IP Multicast	8-51
1. Enable IP Multicast Routing	8-51
2. Enable a Multicast Routing Protocol	8-51
3. Configure the RP for Sparse Mode Operation	8-52
4. Verify IP Multicast Operations	8-52
5. Verify PIM	8-53
6. Verifying Multicast Routing and Clearing the Routing Table	8-54
Summary	8-58
References	8-58
Module Summary	8-59
References	8-59
Module Self-Check	8-61
Module Self-Check Answer Key	8-62

Course Introduction

Overview

Building Cisco Multilayer Switched Networks (BCMSN) v2.2 is an instructor-led course presented by Cisco Learning Partners. This five-day course will teach how to create an efficient and expandable enterprise network by installing, configuring, monitoring, and troubleshooting network infrastructure equipment according to the Campus Infrastructure module in the Enterprise Composite Network Model (ECNM).

Learner Skills and Knowledge

This subtopic lists the skills and knowledge that learners must possess to benefit fully from the course. The subtopic also includes recommended Cisco learning offerings that learners should complete in order to benefit fully from this course.

Learner Skills and Knowledge

Cisco.com

- **Interconnecting Cisco Network Devices v2.2**
- **Complete the initial configuration of a switch**
- **Configure a switch with VLANs**
- **Create basic interswitch connections**
- **Troubleshoot a VLAN**
- **Complete the initial configuration of a router**

Course Goal and Objectives

This topic describes the course goal and objectives.

Course Goal

Cisco.com

“To create an efficient and expandable enterprise network by installing, configuring, monitoring, and troubleshooting network infrastructure equipment according to the Campus Infrastructure module in the Enterprise Composite Network Model.”

Building Cisco Multilayer Switched Networks

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—4

Upon completing this course, you will be able to meet these objectives:

- Use the Campus Infrastructure module of the ECNM to deploy an efficient and expandable enterprise network
- Define VLANs to segment network traffic and manage network utilization
- Implement the Spanning Tree Protocol to accelerate network traffic convergence in Layer 2
- Troubleshoot spanning tree and identify enhancements provided by Rapid Spanning Tree and Multiple Spanning Tree
- Implement multilayer switching to enable high-data throughput communication between isolated VLANs
- Implement redundancy in the ECNM, specifically at Layer 3, to improve and ensure end-to-end availability of network services
- Secure switches in the Campus Infrastructure module against data theft and service loss in the event of network compromise
- Configure the campus switches to optimize traffic flow when voice, video, and data applications traverse a single converged network

Course Flow

This topic presents the suggested flow of the course materials.

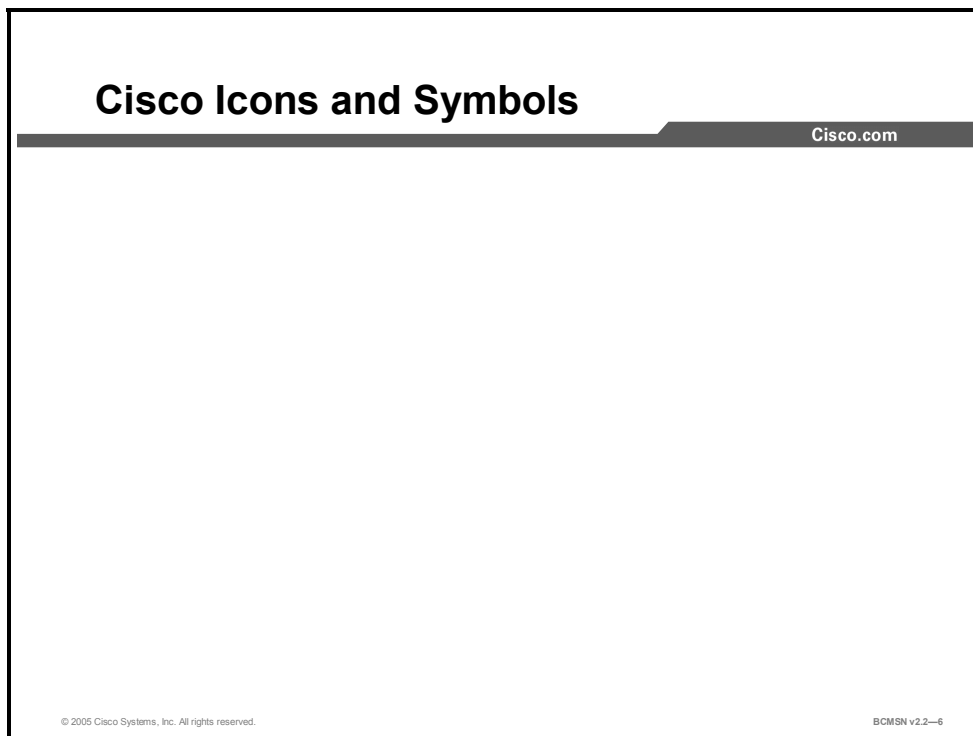
BCMSN 2.2 Course Flow Diagram					
Cisco.com					
	Day 1	Day 2	Day 3	Day 4	Day 5
A M	Course Introduction	Defining VLANS (Cont.)	Enhancing STP	Implementing MLS	Configuring Campus Switches to Support Voice and Video
	Designing a Network				
Lunch					
P M	Defining VLANS	Implementing STP	Implementing Redundancy in the Routing Layer	Minimizing Service Loss and Data Theft	

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-5

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 lab activities. The exact timing of the subject materials and labs depends on the pace of your specific class.

Additional References

This topic presents the Cisco icons and symbols used in this course as well as information on where to find additional technical references.



Cisco Glossary of Terms

For additional information on Cisco terminology, refer to the *Cisco Internetworking Terms and Acronyms* glossary of terms at <http://www.cisco.com/univercd/cc/td/doc/cisintwk/ita/index.htm>.

Your Training Curriculum

This topic presents the training curriculum for this course.



You are encouraged to join the Cisco Certification Community, a discussion forum open to anyone holding a valid Cisco Career Certification (such as Cisco CCIE[®], CCNA[®], CCDA[®], CCNP[®], CCDP[®], CCIP[®], or CCSP[™]). It provides a gathering place for Cisco-certified professionals to share questions, suggestions, and information about Cisco Career Certification programs and other certification-related topics. For more information, visit http://www.cisco.com/en/US/learning/le3/le2/le41/learning_certification_level_home.html.

CCNP Career Certifications

This subtopic describes the requirements for CCNP certification.

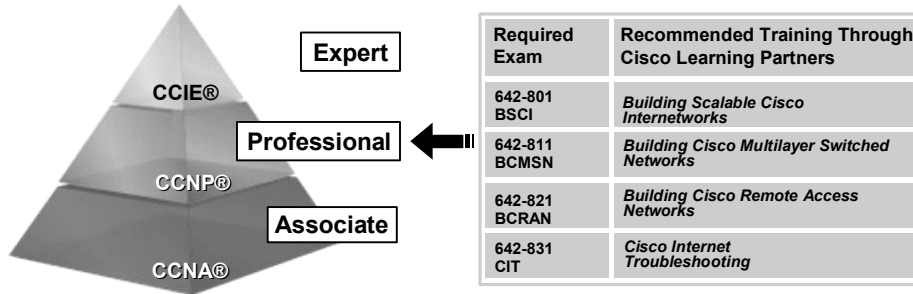
Cisco CCNP Career Certifications

Cisco.com

**Expand Your Professional Options
and Advance Your Career**

Cisco CCNP

Professional-Level Recognition in Routing and Switching



<http://www.cisco.com/go/certifications>

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—8

Designing a Network Using the Campus Infrastructure Module

Overview

Cisco Systems has developed a blueprint for designing networks around the demanding needs of the user communities of today and the vastly improved infrastructure technologies that exist to meet those needs in a modern network. This blueprint, called the Enterprise Composite Network Model (ECNM), is a modular, hierarchical approach to network design. The ECNM assists designers and engineers in developing an optimal network while reducing complexity.

This module examines the shortcomings of networks that had no clear hierarchy or design plan to accommodate organizational growth and points out that simply integrating current technologies into a poorly designed network will not solve problems. It also addresses the benefits of a modular, scalable network model and identifies how various networking technologies are deployed within the Campus Infrastructure module of the ECNM.

Module Objectives

Upon completing this module, you will be able to use the Campus Infrastructure module of the ECNM to deploy an efficient and expandable enterprise network. This ability includes being able to meet these objectives:

- Describe the Campus Infrastructure module of the ECNM and correctly identify the structure and components used to build or expand a campus network
- Identify how various network technologies are best implemented within the Campus Infrastructure module

Describing the Campus Infrastructure Module

Overview

This lesson begins by discussing operational problems found in nonhierarchical networks at Layers 2 and 3 of the Open Systems Interconnection (OSI) model. The Enterprise Composite Network Model (ECNM) is then introduced, and finally, the features and benefits are explained. Students will learn how issues that exist in traditionally designed networks can be resolved by applying this state-of-the-art design to their networks.

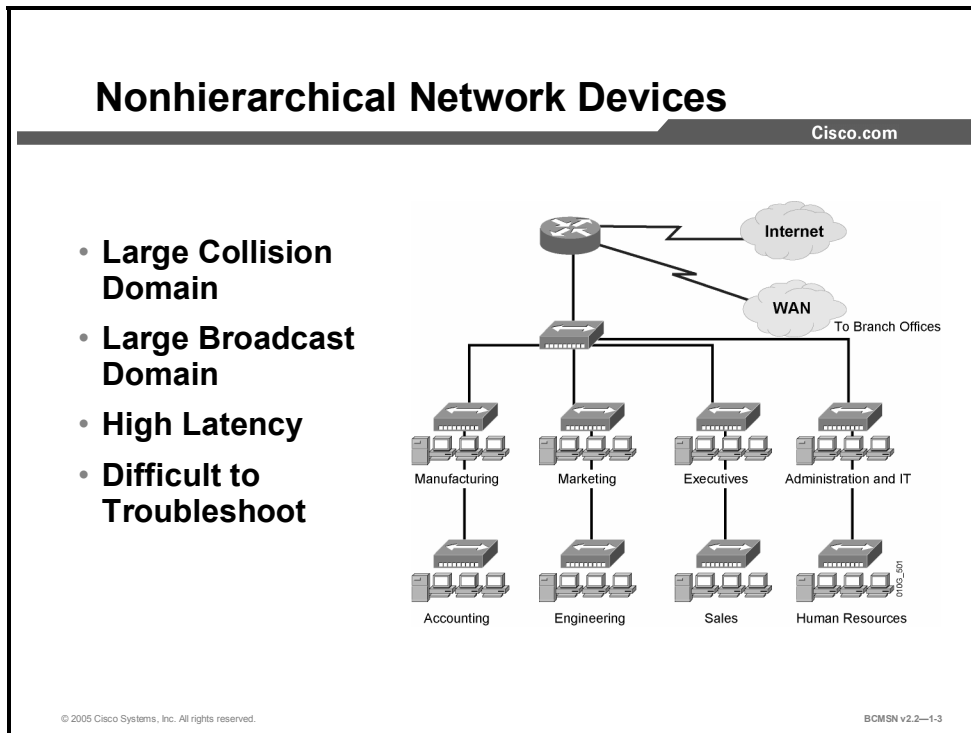
Objectives

Upon completing this lesson, you will be able to describe the Campus Infrastructure module of the ECNM. You will also be able to identify the structure and components used to build or expand a network in the Campus Infrastructure module. This ability includes being able to meet these objectives:

- Describe the devices in a nonhierarchical network
- Identify problems that can occur in a Layer 2 network
- Identify problems that can occur in a Layer 3 network
- Describe the benefits of multilayer switches in a nonhierarchical network
- List the issues that can occur with multilayer switches and VLANs in a nonhierarchical network
- Describe the ECNM used to divide the enterprise network into physical, logical, and functional boundaries
- Explain the benefits of the ECNM
- Describe the Campus Infrastructure module of the ECNM

Devices in a Nonhierarchical Network

This topic describes devices and their functions in a nonhierarchical network.



The simplest Ethernet network infrastructure is composed of a single collision and broadcast domain. This type of network is referred to as a “flat” network because any traffic that is transmitted within it is seen by all of the interconnected devices, even if they are not the intended destination of the transmission. The benefit of this type of network is that it is very simple to install and configure, so it is a good fit for home networking and small offices. The downside of a flat network infrastructure is that it does not scale well as demands on the network increase. These are some of the issues with nonhierarchical networks.

- Traffic collisions increase as devices are added, impeding traffic flow on the network.
- Broadcast traffic increases as devices are added to the network, causing over-utilization of network resources.
- Problem isolation on a large flat network can be difficult.

The following table shows the key network hardware devices in a nonhierarchical network and the function of each.

Network Devices

Device	Function
Hub	Layer 1 device used to interconnect networking components such as PCs, printers, hubs, and routers. This device creates a single broadcast and collision domain for all networking components to which it is connected. Hubs have been superseded in networks by inexpensive switches.
Switch	Layer 2 device used to interconnect networking components such as PCs, printers, hubs, and routers. In its default configuration, this device creates a single broadcast domain for devices connected to it. Each port acts as a separate collision domain.
Router	Layer 3 device used to create and interconnect network segments or broadcast domains. A router must be configured before traffic can flow through it. Each interface creates a Layer 3 segment and therefore establishes a border for the broadcast and collision domains for all devices on that segment.

Layer 2 Network Issues

This topic describes issues that can occur in a switched network.

Layer 2 Switching

Cisco.com

- **Hardware-based bridging**
- **Wire-speed performance**
- **Collision domain per port**
- **Traffic containment based on MAC address**

The diagram illustrates a network topology. At the top, a central switch is connected to two external networks: 'Internet' and 'WAN'. Below this central switch, four departmental switches are connected. Each departmental switch is associated with a group of desktop computers representing a department: 'Manufacturing and Accounting', 'Marketing and Engineering', 'Executives and Sales', and 'Administration, IT and Human Resource'. The central switch is also connected to a 'WAN' cloud.

Issues

- **No traffic between VLANs**
- **Unbounded broadcast domain**
- **Servers not centrally located**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-1.4

Layer 2 switches can significantly improve performance in a carrier sense multiple access collision detect (CSMA/CD) network when used in place of hubs. This is because each switch port represents a single collision domain, and the device connected to that port does not have to compete with other devices to access the media. Ideally, every host on a given network segment is connected to its own switch port, thus eliminating all media contention as the switch manages network traffic at Layer 2. An additional benefit of Layer 2 switching is that large broadcast domains can be broken up into smaller segments by assigning switch ports to different VLAN segments.

For all their benefits, some drawbacks still exist in nonhierarchical switched networks.

- If switches are not configured with VLANs, very large broadcast domains may be created.
- If VLANs are created, traffic cannot move between VLANs using only Layer 2 devices.
- As the Layer 2 network grows, the potential for bridge loops increases. Therefore, the use of a Spanning Tree Protocol (STP) becomes imperative.

Routed Network Issues

This topic describes problems that can occur in a Layer 3 network.

Layer 3 Routing

Cisco.com

- **Single broadcast domain per interface**
- **ACLs can be applied between segments**

The diagram illustrates a network topology. At the top, a central router is connected to the Internet and a Wide Area Network (WAN). Below the router, a switch is connected to several departments: Manufacturing and Accounting, Engineering, Marketing, Executives and Sales, and Administration, IT and Human Resource. Each department is represented by a group of computer icons.

Issues:

- **High per-port cost**
- **Layer 3 processing required**
- **High latency over Layer 2 switching**

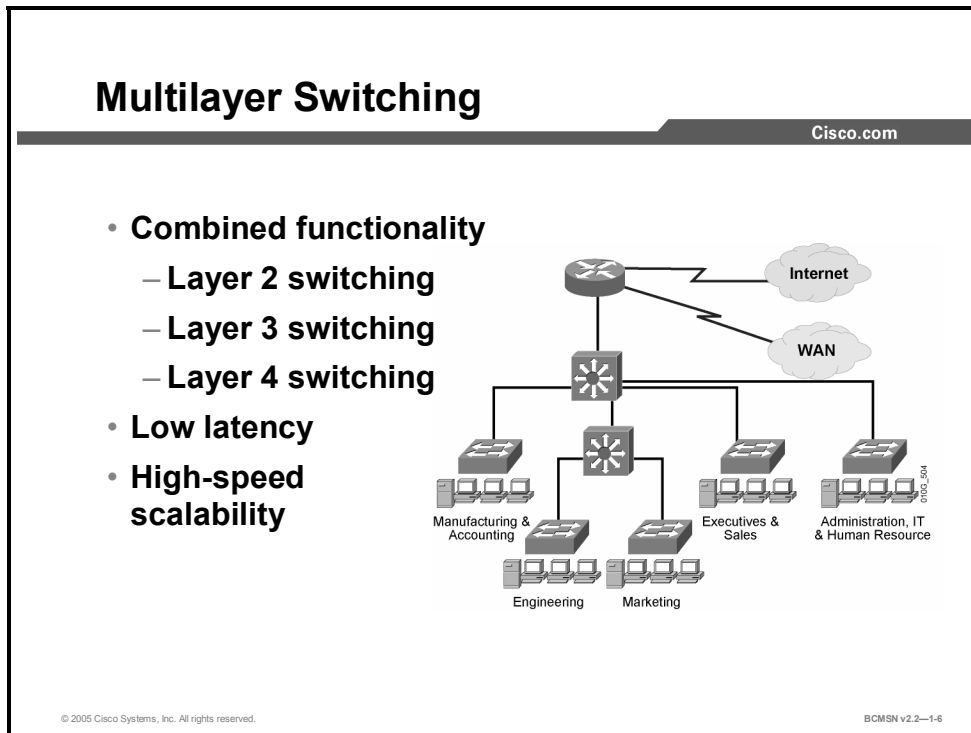
© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-1-5

A major limitation of Layer 2 switches is that they cannot switch traffic between Layer 3 network segments (IP subnets for example). Traditionally, this was done using a router. Unlike switches, a router acts as a broadcast boundary and does not forward broadcasts between its interfaces. Additionally, a router provides for an optimal path determination process. The router examines each incoming packet to determine which route the packet should take through the network. Also, the router can act as a security device, manage quality of service (QoS), and apply network policy. Although routers used in conjunction with Layer 2 switches resolve many issues, some concerns still remain.

- When security or traffic management components, such as access control lists (ACLs), are configured on router interfaces, the network may experience delays as the router processes each packet in software.
- When routers are introduced into a switched network, end-to-end VLANs are no longer supported because routers terminate the VLAN.
- Routers are more expensive per interface than Layer 2 switches, so their placement in the network should be well planned. Nonhierarchical networks by their very nature require more interconnections and, hence, more routed interfaces.
- In a nonhierarchical network, the number of router interconnections may result in peering problems between neighboring routers.
- Because traffic flows are hard to determine, it becomes difficult to predict where hardware upgrades are needed to mitigate traffic bottlenecks.

What Is a Multilayer Switch?

This topic describes multilayer switches in a nonhierarchical network.



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 ASIC and other specialized circuitry. A multilayer switch does everything to a frame and packet that a traditional switch or router does, including the following:

- 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 MIB
- Applies security and policy controls, if required
- Provides optimal path determination
- Can (if a sophisticated modular type) support a wide variety of media types and port densities
- Has the ability to support QoS
- Has the ability to support Voice over IP (VoIP) and inline power requirements

Because it is designed to handle high-performance LAN traffic, a multilayer switch can be placed anywhere within the network, cost-effectively replacing traditional switches and routers. Generally, however, a multilayer switch may be more than is required to provide end systems access to network resources.

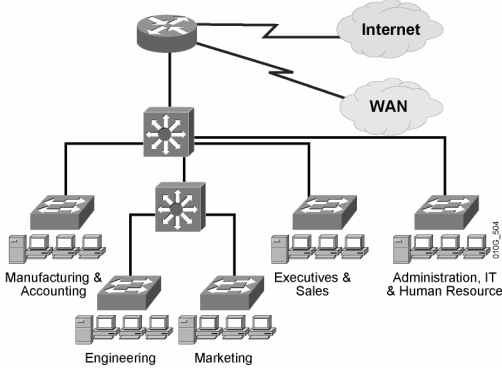
Issues with Multilayer Switches in a Nonhierarchical Network

This topic describes the issues that occur with multilayer switches and VLANs in a nonhierarchical network.

Issues with Multilayer Switches in a Nonhierarchical Network

Cisco.com

- **Single point of failure for Layers 2 and Layers 3**
- **Underutilization of Hardware**
- **Spanning tree complexity**
- **Servers not centrally located**



The diagram illustrates a nonhierarchical network topology. At the top, a central multilayer switch (represented by a square with a starburst) is connected to an Internet cloud and a WAN cloud. Below this central switch, there are four other switches (represented by squares with starbursts) connected to it. These four switches are further connected to various departments: Manufacturing & Accounting, Engineering, Marketing, Executives & Sales, and Administration, IT & Human Resource. Each department is represented by a group of computer icons. The central switch is the only point of connection to the Internet and WAN, creating a single point of failure for all traffic passing through it.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-1-7

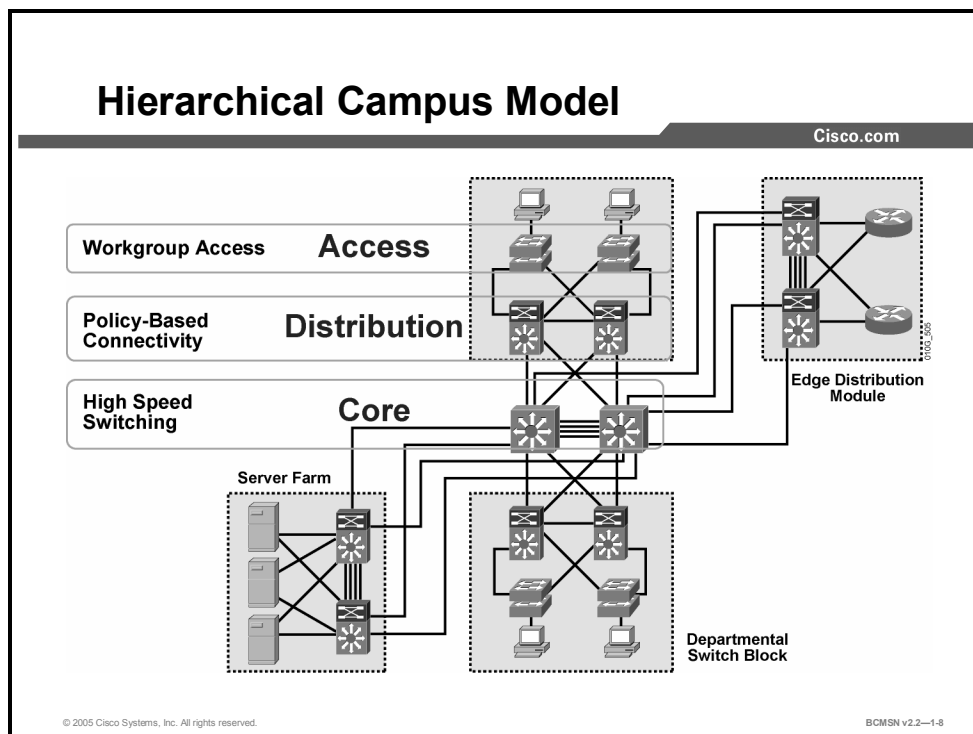
Multilayer switches combine switching and routing on a single hardware platform and can enhance overall network performance when deployed properly. Multilayer switches provide very high-speed Layer 2 and Layer 3 functionality by “caching” much of the forwarding information between sources and destinations.

Here are issues that exist when a multilayer switch is deployed in an improperly designed network.

- Multilayer switches, by condensing the functions of switching and routing in a single chassis, can create single points of failure if redundancy for these devices is not carefully planned and implemented.
- Switches in a flat network are interconnected, creating many paths between destinations. If active, these potential redundant paths will create bridging loops. To control this, the network must run an STP. Networks that use the IEEE 802.1D protocol may experience periods of disconnection and frame flooding during topology change.
- Multilayer switch functionality may be underutilized if a multilayer switch is simply a replacement for the traditional role of a router in a nonhierarchical network.

Enterprise Composite Network Model

This topic describes the ECNM, which can be used to divide the enterprise network into physical, logical, and functional areas.



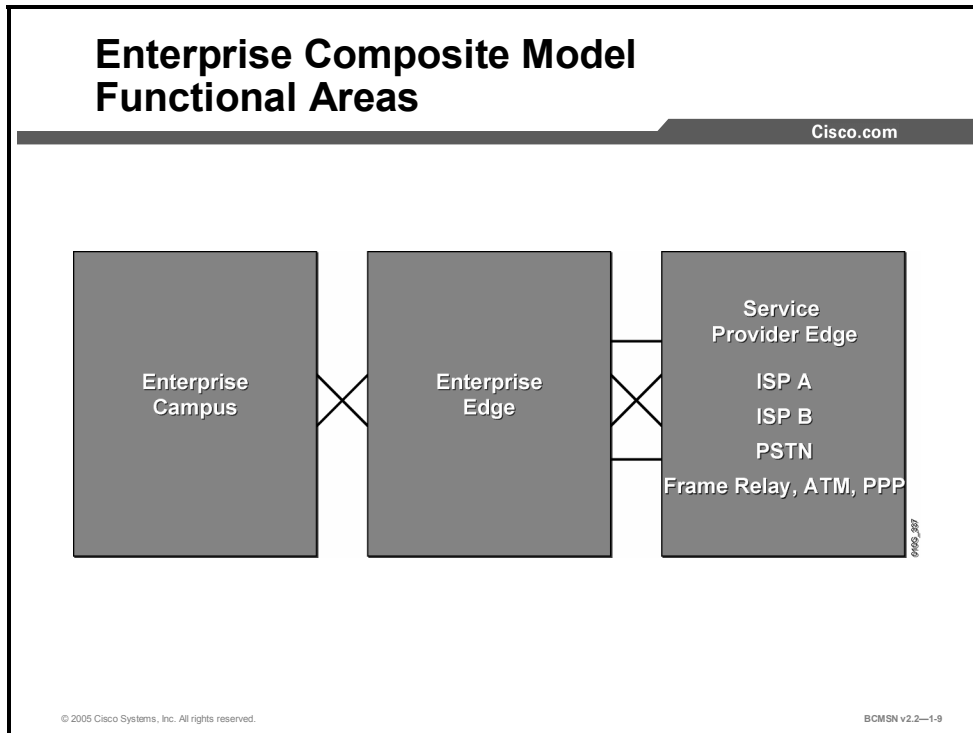
The ECNM provides a modular framework for designing networks. This modularity allows flexibility in network design and facilitates ease of implementation and troubleshooting. The hierarchical model divides networks into the Building Access, Building Distribution, and Building Core layers, as follows:

- **Building Access layer:** The Building Access layer is used to grant user access to network devices. In 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 at remote sites may provide access to the corporate network across WAN technology.
- **Building Distribution layer:** The Building Distribution layer aggregates the wiring closets and uses switches to segment workgroups and isolate network problems.
- **Building Core layer:** The Building Core layer (also known as the Campus Backbone submodule) 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 adapt to changes very quickly.

The ECNM divides the enterprise network into physical, logical, and functional areas. These areas allow network designers and engineers to associate specific network functionality on equipment based upon its placement and function in the model.

Enterprise Composite Network Model Functional Areas

This subtopic describes the functional areas of the ECNM.



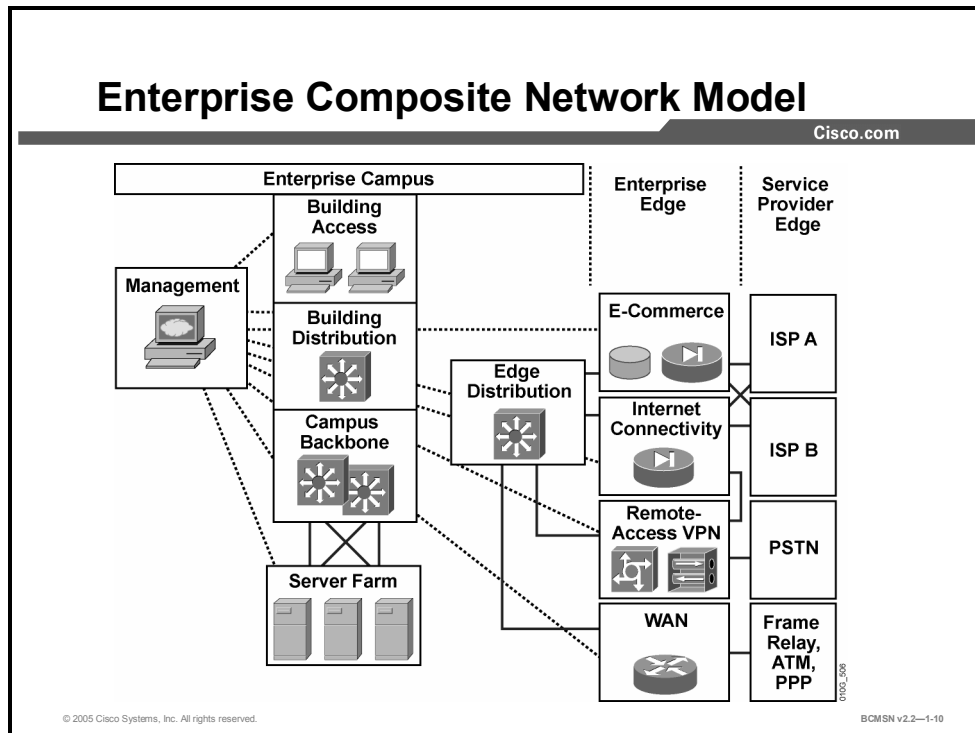
The ECNM introduces modularity by dividing the network into functional areas that ease design, implementation, and troubleshooting tasks. An enterprise campus is defined as one or more buildings, with multiple virtual and physical networks, connected across a high-performance, multilayer-switched backbone.

The ECNM 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 area contains the network elements required for independent operation within a single campus, such as access from all locations to central servers. The Enterprise Campus functional area does not offer remote connections or Internet access.
- **Enterprise Edge:** The Enterprise Edge aggregates connectivity from the various resources external to the enterprise network. As traffic comes into the campus, this area filters traffic from the external resources and routes it into the Enterprise Campus functional area. It contains all of the network elements for efficient and secure communication between the enterprise campus and remote locations, remote users, and the Internet. The Enterprise Edge would replace the “Demilitarized Zone (DMZ)” of most networks.
- **Service Provider Edge:** This functional area represents connections to resources external to the campus. This area facilitates communication to WAN and Internet service providers’ technologies.

Enterprise Composite Network Model Benefits

This topic describes the benefits of the ECNM.



To scale the hierarchical model, Cisco introduced the ECNM, which further divides the enterprise network into physical, logical, and functional areas. The ECNM contains functional areas, each of which has its own Building Access, Building Distribution, and Building Core (or Campus Backbone) layers.

The ECNM meets these criteria:

- It is 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.

Benefits of the Enterprise Composite Network Model

This subtopic describes the benefits of implementing the ECNM.

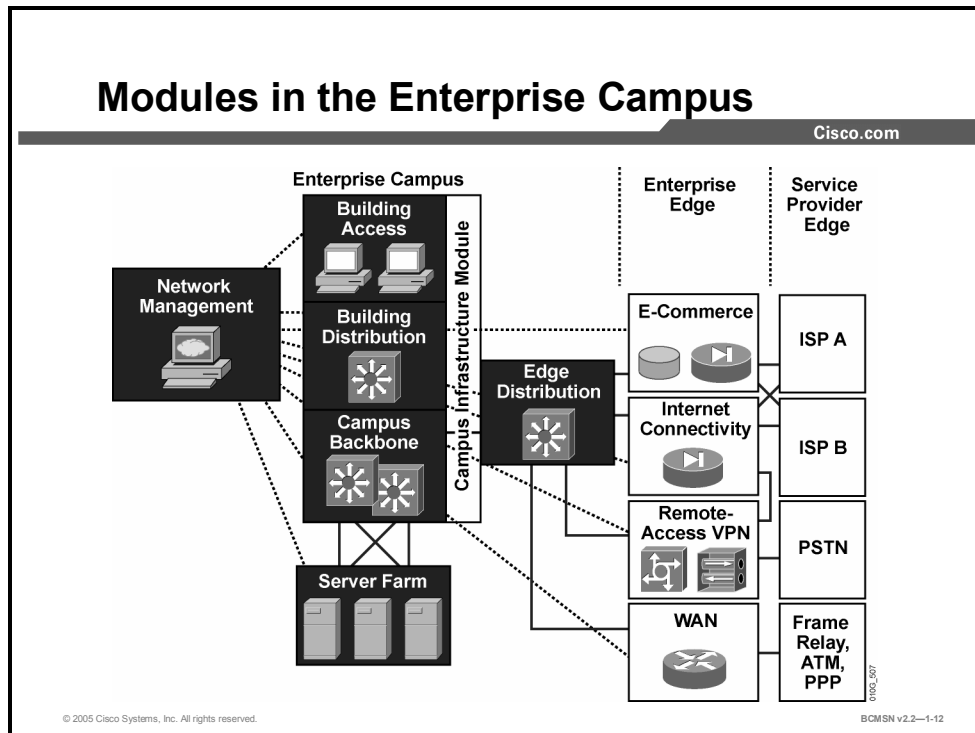
Enterprise Composite Network Model Benefits			
Cisco.com			
	Performance	Scalability	Availability
Building Access	Critical to desktop performance	Provides port density	Important to provide redundancy
Building Distribution	Critical to campus performance	Provides switch modularity	Critical to provide redundancy
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
Critical to WAN and internet performance	Critical to WAN and Internet performance	Provides switch modularity	Important to provide redundancy

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—1-11

The ECNM has a number of benefits for each of the submodules where it is implemented.

Modules of the Enterprise Campus

This topic describes the Enterprise Campus functional area.



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.

Campus Infrastructure module: Includes Building Access and Building Distribution submodules. It connects users within the campus to the Server Farm and Edge Distribution modules. The Campus Infrastructure module is composed of one or more floors or buildings connected to the Campus Backbone submodule.

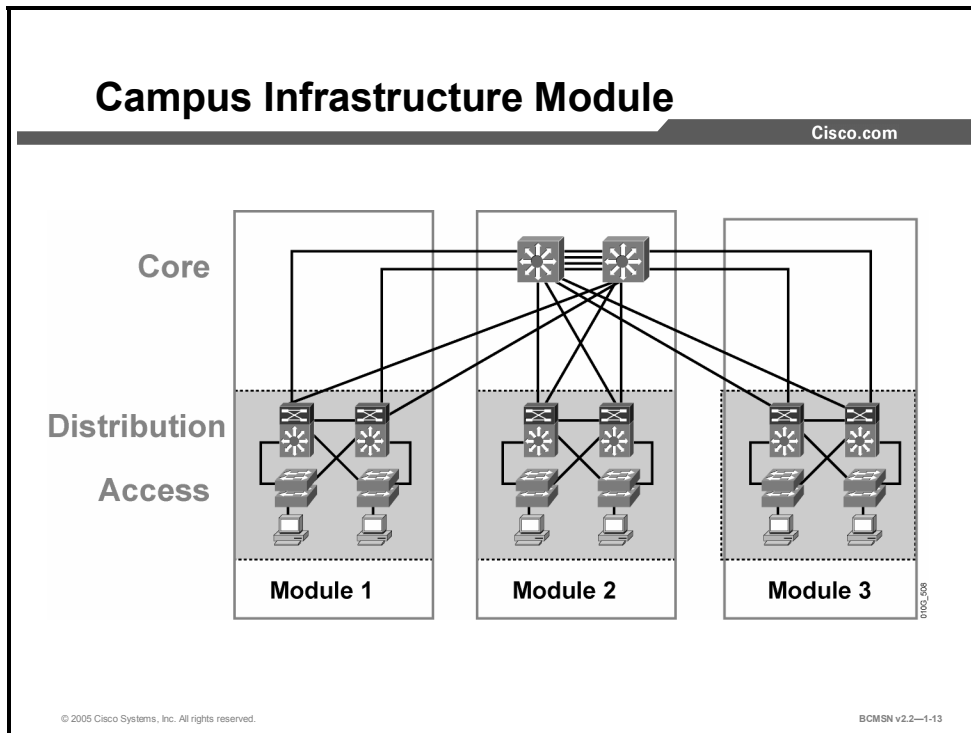
Network Management module: Performs system logging and authentication as well as network monitoring and general configuration management functions.

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

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.

Campus Infrastructure Module

This topic describes the Campus Infrastructure module of the ECNM.



The Campus Infrastructure module connects users within a campus to the Server Farm and Edge Distribution modules. The Campus Infrastructure module comprises Building Access and Building Distribution switches connected through the Campus Backbone to campus resources.

A 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 services such as support for multiple VLANs, private VLANs, and establishment of trunk links to the Building Distribution layer and IP phones. Each building access switch has connections to redundant switches in the Building Distribution submodule.
- **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. Traffic generally flows through the building distribution switches and onto the campus core or backbone. This submodule provides fast failure recovery because each building distribution switch maintains two equal-cost paths in the routing table for every Layer 3 network number. Each building distribution switch has connections to redundant switches in the core.

- **Campus Backbone submodule (also known as Building Core layer):** Provides redundant and fast-converging connectivity between buildings and the Server Farm and Edge Distribution modules. The purpose of the Campus Backbone submodule is to switch traffic as fast as possible between Campus Infrastructure submodules and destination resources. Forwarding decisions should be made at the ASIC level whenever possible. Routing, ACLs, and processor-based forwarding decisions should be avoided at the core and implemented at building distribution devices whenever possible. High-end Layer 2 or Layer 3 switches are used at the core for high throughput, with optimal routing, QoS, and security capabilities available when needed.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Flat networks extend broadcast and failure domains.**
- **Layer 2 devices reduce the size of a collision domain but cannot route between VLANs.**
- **Routing passes traffic between VLANs but has drawbacks in speed and versatility.**
- **Multilayer switches offer many improvements over routers if applied in a hierarchical manner.**
- **The Enterprise Composite Network Model provides a scalable and reliable network.**
- **The Campus Infrastructure module is a key component of the Enterprise Composite Network Model.**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—1-14

Deploying Technology in the Campus Infrastructure Module

Overview

This lesson addresses an organization's business and technology needs and addresses how those needs can be met by applying the appropriate resources to the Campus Infrastructure module.

Objectives

Upon completing this lesson, you will be able to identify how various technologies are best implemented within the Campus Infrastructure module. This ability includes being able to meet these objectives:

- List issues that can occur in a poorly designed network
- Develop a hierarchical IP addressing scheme that maps well to the Campus Infrastructure module
- Describe the different network interconnection technologies and identify their appropriate use in the Campus Infrastructure module
- Determine the equipment and cabling needs on the various links of the Campus Infrastructure module
- Map a hierarchical IP addressing scheme to the access VLANs in the Campus Infrastructure module
- Identify the most common traffic types on the network
- Identify the most common traffic sources and their destination on a campus network
- Identify the two interfaces used to configure Cisco Catalyst switches

Issues in a Poorly Designed Network

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

Issues in a Poorly Designed Network

Cisco.com

- **Unbounded failure domains**
- **Large broadcast domains**
- **Large amount of unknown MAC unicast traffic**
- **Unbounded multicast traffic**
- **Management and support challenges**
- **Possible security vulnerabilities**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-1-3

A poorly designed network has increased support costs, reduced service availability, and limited support for new applications and solutions. Less than optimal performance will affect end users directly as well as affecting access to central resources. 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 minimize the extent of a network problem when it occurs. When Layer 2 and Layer 3 boundaries are not clearly defined, failure in one network area can have a far-reaching effect.
- **Broadcast domains:** Broadcasts exist in every network. Many applications and many network operations require broadcasts to function properly; therefore, it is not possible to completely eliminate them. Just as with failure domains, in order to minimize the negative impact of broadcasts, broadcast domains should have clear boundaries and include an optimal number of devices.
- **Large amount of unknown MAC unicast traffic:** Catalyst switches limit unicast frame forwarding to ports associated with the specific unicast address. However, frames arriving for a destination MAC address not recorded in the MAC table are flooded out all switch ports and this is known as an “unknown MAC unicast flooding.” Because this causes excessive traffic on switch ports, Network Interface Cards (NICs) have to attend to a larger number of frames on the wire, and security can be compromised as data is being propagated on a wire for which it was not intended.
- **Multicast traffic on ports where not intended:** IP multicast is a technique that allows IP traffic to be propagated from one source to a multicast group identified by a single IP and MAC destination group address pair. Similar to unicast flooding and broadcasting, multicast frames will be flooded out all ports on the same VLAN where they were received.

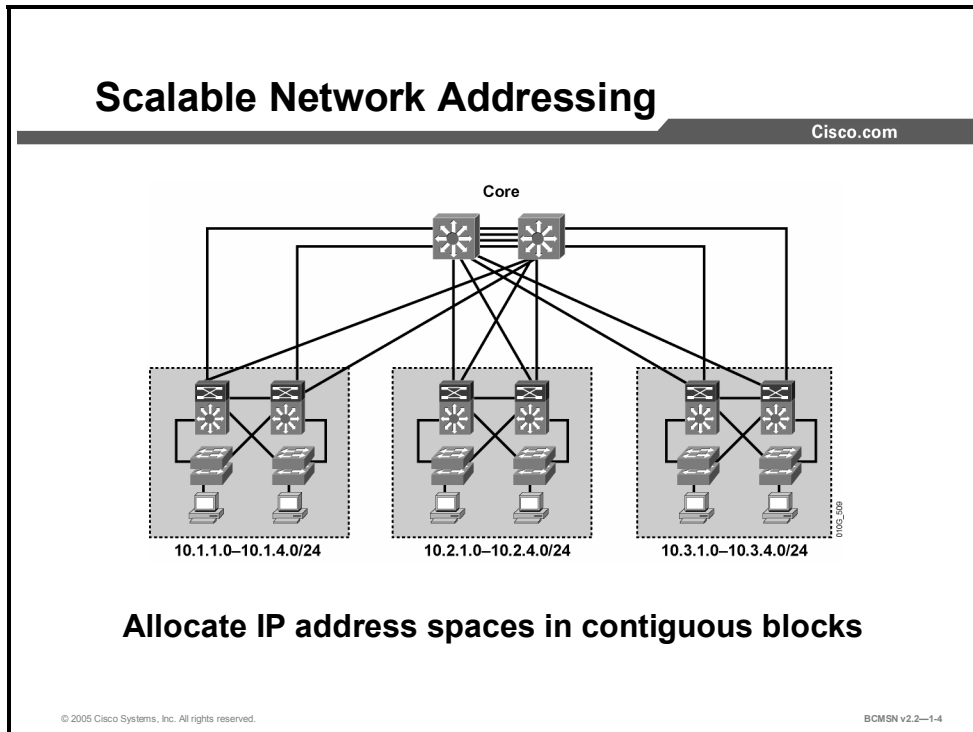
A proper design allows for containment of multicast frames while allowing them to be functional.

- **Difficulty in management and support:** Because a poorly designed network may be disorganized, poorly documented, and lacking easily identified traffic flows, support, maintenance, and problem resolution become time-consuming and arduous tasks.
- **Possible security vulnerabilities:** A poorly designed switched network with little attention to security requirements at the access layer can compromise the integrity of the entire network.

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

Designing a Hierarchical IP Addressing Scheme

This topic describes how to develop a hierarchical IP addressing scheme to migrate a network to the Campus Infrastructure module of the Enterprise Composite Network Model (ECNM).



Hierarchical network addressing means that IP network numbers are applied to the network segments or VLANs in an orderly fashion that takes into consideration the network as a whole. Blocks of contiguous network addresses are reserved for, and configured on, devices in a specific area of the network.

Here are some benefits of hierarchical addressing.

- **Ease of management and troubleshooting:** Hierarchical addressing groups network addresses contiguously. Network management and troubleshooting are more efficient, as a well-known IP addressing scheme will make problem components easier to locate.
- **Minimizing of error:** Orderly network address assignment can minimize error and duplicate address assignment.
- **Reduced number of routing table entries:** In a hierarchical addressing plan, routing protocols are able to invoke route summarization which allows a single routing table entry to represent a collection of IP network numbers. Route summarization makes routing table entries manageable and provides the following benefits:
 - Reduced number of CPU cycles when recalculating a routing table or sorting through the routing table entries to find a match
 - Reduced router memory requirements
 - Faster convergence after a change in the network
 - Easier troubleshooting

Guidelines for Applying IP Address Space in the Enterprise Network

The ECNM provides a modular framework for designing and deploying networks. It also provides the ideal structure for overlaying a hierarchical IP addressing scheme. Here are some guidelines to follow.

- Design the IP addressing scheme so that blocks of 4, 8, 16, 32, or 64 contiguous network numbers can be assigned to the subnets in a given building distribution and access switch block.
- At the Building Distribution layer, continue to assign network numbers contiguously out toward to the access layer devices.
- Have a single IP subnet correspond with a single VLAN.
- Subnet at the same binary value on all network numbers, avoiding variable length subnet masks when possible in order to minimize error and confusion when troubleshooting or configuring new devices and segments.

Interconnection Technologies

This topic describes the different network interconnection technologies and identifies their appropriate use in the ECNM.

Interconnection Technologies	
Cisco.com	
Technology	Use
Fast Ethernet	Connects end-user devices to the access layer switch
Gigabit Ethernet	Access to distribution switch, high-use servers
10-Gigabit Ethernet	High-speed switch to switch links, backbones
EtherChannel	High-speed switch to switch links, backbones with redundancy

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—1-5

A number of technologies are available to interconnect devices in the campus network. Some of the more common technologies are listed here. The interconnection technology selected will depend on the amount of traffic the link must carry. A mixture of copper and fiber-optic cabling will likely be used, based on distances, noise immunity requirements, security, and other business requirements.

- **Fast Ethernet (100-Mbps Ethernet):** This LAN specification (IEEE 802.3u) operates at 100 Mbps over twisted-pair cable. The Fast Ethernet standard raises the speed of Ethernet from 10 Mbps to 100 Mbps with only minimal changes to the existing cable structure. A switch with port functioning at both 10 and 100 Mbps can move frames between ports without Layer 2 protocol translation.
- **Gigabit Ethernet:** An extension of the IEEE 802.3 Ethernet standard, Gigabit Ethernet increases speed tenfold over that of Fast Ethernet, to 1000 Mbps, or 1 gigabit per second (Gbps). IEEE 802.3z specifies operations over fiber optics, and IEEE 802.3ab specifies operations over twisted-pair cable.
- **10-Gigabit Ethernet:** 10-Gigabit Ethernet was formally ratified as an IEEE 802.3 Ethernet standard in June 2002. This technology is the next step for scaling the performance and functionality of an enterprise. With the deployment of Gigabit Ethernet becoming more common, 10-Gigabit will become the norm for uplinks.

- **EtherChannel:** This feature provides link aggregation of bandwidth over Layer 2 links between two switches. EtherChannel bundles individual Ethernet ports into a single logical port or link, providing aggregate bandwidth of 1600 Mbps (8-100Mbps links, full duplex) or 16 Gbps (8-Gigabit links, full duplex) between two Catalyst switches. All interfaces in each EtherChannel bundle must be configured with similar speed, duplex, and VLAN membership.

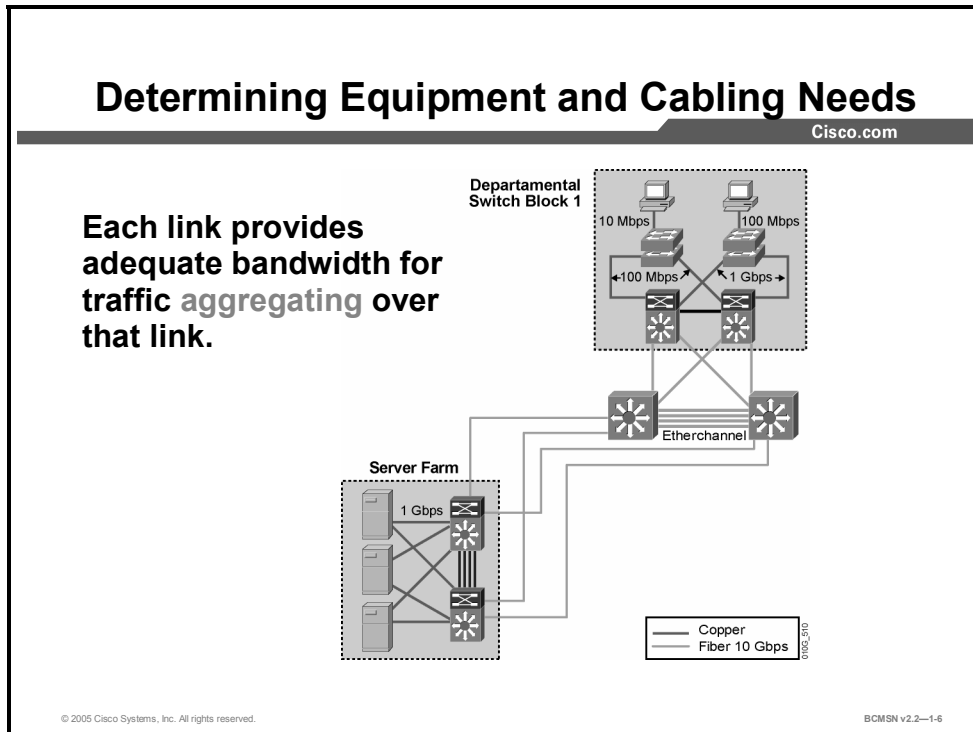
The “Interconnection Technologies” table discusses the use of each technology in the Campus Infrastructure module.

Interconnection Technologies

Technology	Use in Campus Infrastructure Module
Fast Ethernet	Often used to connect end-user devices to the access layer switch. If user connections are running at 10 mbps, Fast Ethernet links between access and distribution switches may be adequate. It is adequate for servers in small- to medium-sized networks if full duplex is invoked.
Gigabit Ethernet	High-speed LAN backbones connecting building distribution switches to campus backbone switches. Widely used internal or Internet-accessible servers might be connected via gigabit Layer 2 technology.
10-Gigabit Ethernet	Very high-speed LAN backbone and link aggregation. As gigabit links become more common, 10-Gigabit will be necessary to scale the uplinks.
EtherChannel	Any connection between switches with multiple physical links that requires high bandwidth and redundancy. Links between Building Distribution and Campus Backbone might be Gigabit EtherChannel. Links between access and distribution layer devices might be Fast Ethernet EtherChannel.

Determining Equipment and Cabling Needs

This topic describes the equipment and cabling needed to migrate a network to the Campus Infrastructure module.



The graphic highlights the changes that will take place as the nonhierarchical network is redeployed in the new ECNM strategy.

There are four objectives in the design of any high-performance network: security, availability, scalability, and manageability. The ECNM, when implemented properly, provides the framework to meet these objectives. In the migration from a current network infrastructure to the ECNM, a number of infrastructure changes may be needed, including the replacement of current equipment and existing cable plant.

This list describes the equipment and cabling decisions that should be considered when altering infrastructure.

1. Replace hubs and legacy switches with new switches at the Building Access layer. Select equipment with the appropriate port density at the access layer to support the current user base while preparing for growth. Some designers begin by planning for about 30 percent growth. If the budget allows, use modular access switches to accommodate future expansion. Consider planning for support of inline power and quality of service (QoS) if IP telephony may be implemented in the future.
2. When building the cable plant from the Building Access layer to the Building Distribution layer devices, remember that these links will carry aggregate traffic from the end nodes at the access layer to the building distribution switches. Ensure that these links have adequate bandwidth capability. EtherChannel bundles can be used here to add bandwidth as necessary.

3. At the Building Distribution layer, select switches with adequate performance to handle the load of the current Building Access layer. Also plan some port density for adding trunks later to support new access layer devices. The devices at this layer should be multilayer switches that support routing between the workgroup VLANs and network resources. Depending on the size of the network, the building distribution layer devices may be fixed chassis or modular. Plan for redundancy in the chassis and in the connections to the access and core layers, as the business objectives dictate.
4. The campus backbone equipment must support high-speed data communications between other submodules. Be sure to size the backbone for scalability and plan on redundancy.

Cisco has online tools to assist designers in making the proper selection of devices and uplink ports based on business and technology needs. Cisco suggests oversubscription ratios that can be used to plan bandwidth requirements between key devices on a network with average traffic flows.

- **Access to distribution layer links:** The oversubscription ratio should be no higher than 20:1. That is, the link can be 1/20 of the total bandwidth available cumulatively to all end devices using that link.
- **Distribution to core links:** The oversubscription ratio should be no higher than 4:1.
- **Between core devices:** There should be little to no oversubscription planning. That is, the links between core devices should be able to carry traffic at the speed represented by the aggregate number bandwidth of all the distribution uplinks into the core.

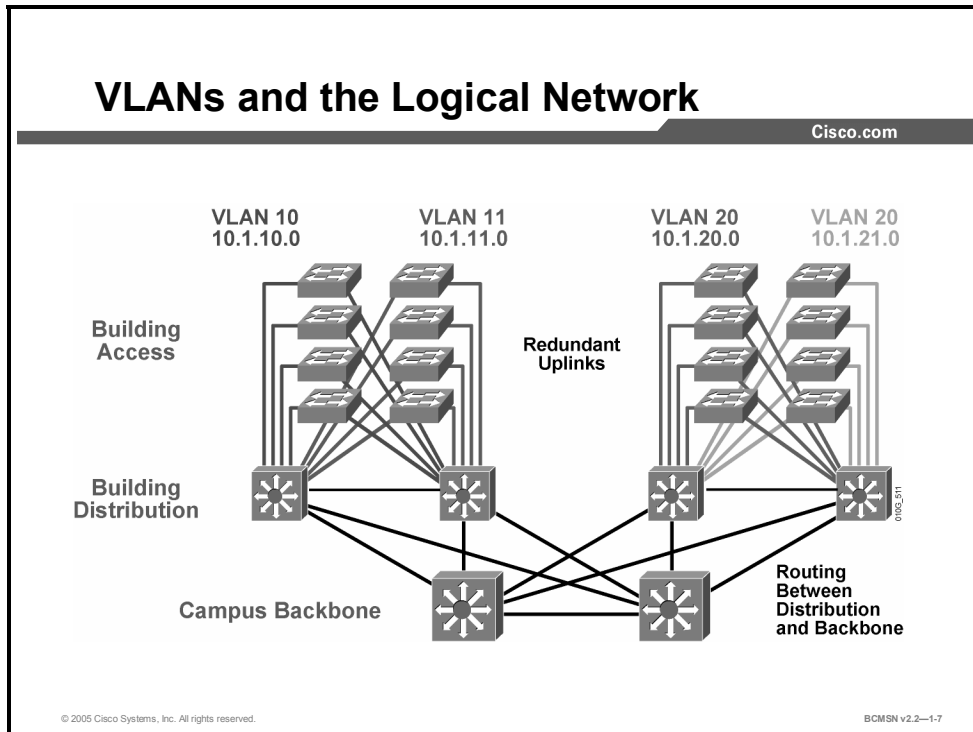
Caution These ratios are appropriate for estimating average traffic from access layer, end-user devices. They are not accurate for planning oversubscription from the Server Farm or Edge Distribution modules. They are also not accurate for planning bandwidth needed on access switches hosting atypical user applications with high bandwidth consumption (for example, non-client server databases or multimedia flows to unicast addresses. Using QoS end to end prioritizes the traffic that would need to be dropped in the event of congestion.

References

For additional information, refer to Cisco Systems, Inc., Cisco Product Advisor:
http://www.cisco.com/en/US/products/prod_tools_index.html.

Mapping VLANs in a Hierarchical Network

This topic describes a methodology to assign VLANs to different network segments in a hierarchical network.

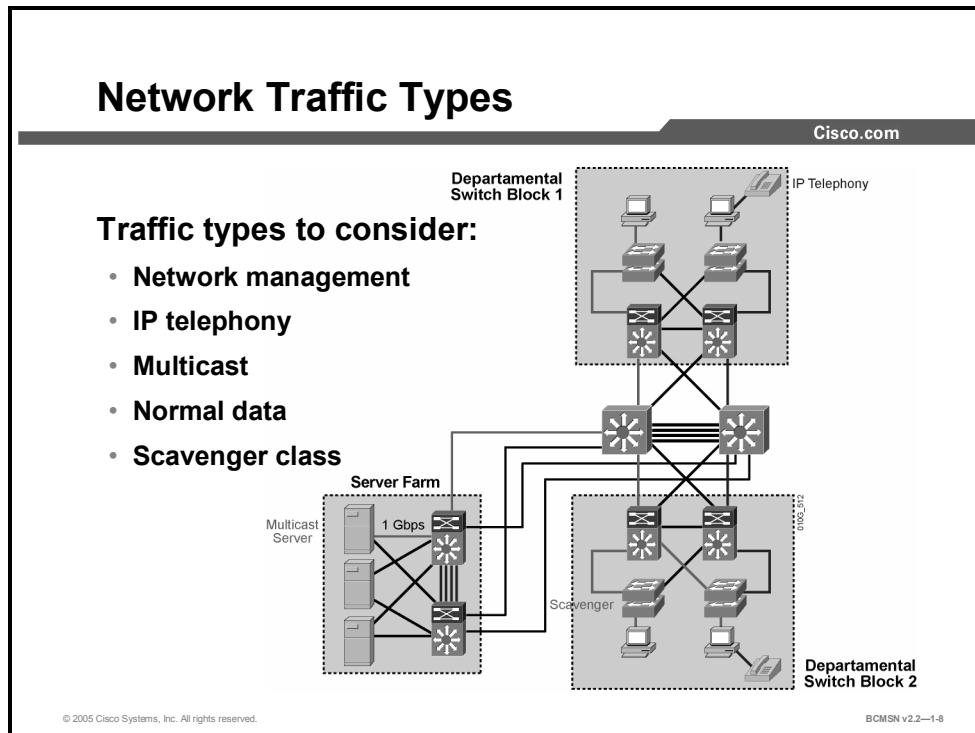


When mapping VLANs onto the new hierarchical network design, keep these parameters in mind.

1. Examine the subnetting scheme that has been applied to the network and associate a VLAN to each subnet.
2. Configure routing between VLANs at the distribution layer. Routing always occurs at the distribution layer switch.
3. Make end-user VLANs and subnets local to a specific switch block.
4. Ideally, limit a VLAN to one access switch or switch stack. However, it may be necessary to extend a VLAN across multiple access switches within a switch block to support a capability such as wireless mobility.

Traffic Types

This topic describes some of the different traffic sources on the network.



This table lists different types of traffic that may exist on the network and will require consideration before device placement and VLAN configuration.

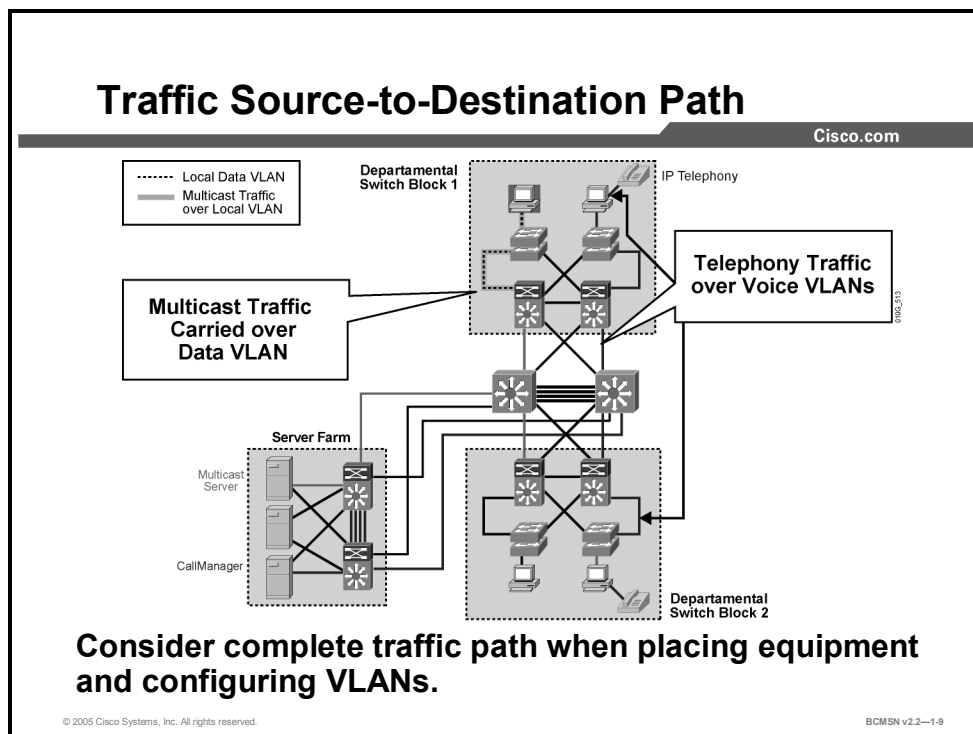
Traffic Types

Traffic Type	Description
Network Management	Many different types of network management traffic may be present on the network. Examples include bridge protocol data units (BPDUs), Cisco Discovery Protocol (CDP) updates, Simple Network Management Protocol (SNMP) and RMON (Remote Monitoring) traffic. Some designers will assign a separate VLAN to the task of carrying certain types of network management traffic in order to make network troubleshooting easier.
IP Telephony	There are two types of IP telephony traffic: signaling information between the end devices (for example, IP phones, softswitch, such as Cisco CallManager) and the data packets of the conversation itself. Often, the data to and from IP phones will be configured on a separate VLAN for voice traffic, as the designer will want to apply QoS measures to that type of data.
IP Multicast	IP multicast traffic is sent from a particular source address to non-unique MAC addresses. Examples of applications that generate this type of traffic are IPTV broadcasts and imaging software used to configure workstations and servers quickly. Multicast traffic can produce a large amount of data streaming across the network. Switches need to be configured to keep this traffic from flooding to devices that have not requested it, and routers need to ensure that multicast traffic is forwarded to the network areas where it is requested.

Traffic Type	Description
Normal Data	This is typical application traffic related to file and print services, e-mail, Internet browsing, database access, and other shared network applications. This data may have to be treated in the same or different ways in different parts of the network, based on the volume of each type. Examples of this type of traffic are small to medium business (SMB), Netware Core Protocol (NCP), Simple Mail Transfer Protocol (SMTP), and HTTP.
Scavenger Class	Scavenger class includes all traffic with protocols or patterns that exceed their normal data flows. It is used to protect the network from exceptional traffic flows that may be the result of malicious programs executing on end-system PCs.

Considering Traffic Source-to-Destination Path

This topic describes the complete source-to-destination path taken by network traffic.



The size of an enterprise network drives the design and placement of certain types of devices. If the network is designed according to the ECNM, there will be distinct devices separating the access, distribution, and backbone areas of the network. The network design and the types of applications supported will determine where certain traffic sources are located. In the case of multicast and IP telephony applications, they do share some common traffic types. Specifically, if a Cisco CallManager is providing music on hold, it may need to multicast that traffic stream. Similarly, if there is an IPTV broadcast server on the network, it will also be sending information via multicast to a specific set of devices. Consider the following points when determining where to place the servers.

- IP multicast servers may exist within a server farm or be distributed throughout the network at appropriate locations. Select distribution layer switches to act as rendezvous points and place them where they are central to the location of the largest distribution of receiving nodes.
- Cisco CallManager servers must be accessible throughout the network at all times. Ensure that there are redundant NICs in the publisher and subscriber servers and redundant connections between those NICs and the upstream switch from the server. It is recommended that voice traffic be configured on its own VLAN.
- VLAN trunks must be configured appropriately to carry IP telephony traffic throughout the network or to specific destinations.


Cisco Catalyst Configuration Interfaces

This topic compares Catalyst configuration interfaces.

Cisco Catalyst Software Interface

Cisco.com

- Use catalyst software interface for Layer 2 configuration.
- Some platforms have options for Layer 2 configuration via Cisco IOS interface.



**Cisco Catalyst 4000, 5500,
and 6500 switches**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—1-10

In the era of the early high-end Catalyst switches, the operating system and the command interface were significantly different from the Cisco IOS mode navigation interfaces available on all newer Catalyst platforms. The two interfaces have different features and a different prompt.

Catalyst Software Interface

The original Catalyst interface is sometimes referred to as the “set-based” or, more recently, “Catalyst software” command-line interface.

In the Catalyst software, commands are executed at the switch prompt, which can be either nonprivileged (where a limited subset of user-level commands is available) or at a password-protected privileged mode (where all commands are available). Configuration commands are prefaced with the keyword **set**. In the example below, the Catalyst software commands execute the following: first, show the status of a port; second, move to enable mode that requires a password; third, enable the port.

Example: Using Catalyst Software Commands

```
Console> show port 3/5
Console> enable

Enter password:
Console(enable) set port enable 3/5
```

Cisco IOS Interface


This topic describes the Cisco IOS interface that is used on most Catalyst switches.

Cisco IOS Interface

Cisco.com

On most Catalyst switches, Cisco IOS interface is standard for

- **Layer 2 configuration**
- **Layer 3 configuration on Multilayer Switch**



© 2005 Cisco Systems, Inc. All rights reserved. BCM5N v2.2-1-11

Catalyst switch platforms have had a number of different operating systems and user interfaces. Over the years, Cisco has made great strides in converting the interface on nearly every Catalyst platform to the IOS interface familiar to Cisco users on routing platforms. Unlike the Catalyst software, various modes are navigated to execute specific commands.

Here is an example of how switch port 3 might be enabled on an access layer switch using the IOS interface and how its status is verified after configuration. Compare how the IOS interface is navigated here to the previous example, showing how the same function is performed in the Catalyst software.

Example: Using IOS Commands

```
Switch# config terminal
Switch(config)# interface fastethernet 0/3
Switch(config-if)# no shut
Switch(config-if)# end
Switch# show interface fastethernet 0/3
```

Configuration Interface Available on Various Catalyst Platforms

Some widely used Catalyst switch platforms that support the IOS interface are 2950, 3500, 3700, 4500*, 6500*, and 8500.

* These platforms have an option to use either IOS or Catalyst software for Layer 2 configuration.

The Catalyst software interface exists on several modular Catalyst platforms, including 4000, 4500, 5500, 6000, and 6500.

For example, on the Catalyst 6500, you have the option of using the Catalyst software, Catalyst software plus IOS software, or IOS software functionality.

Catalyst 6500 Interfaces

Operating System	Where Installed	Purpose
Catalyst software	On Catalyst switch supervisor module.	Catalyst software interface provided to configure Layer 2 switch functions. Suitable if unit is used in a Layer 2 environment only.
Catalyst software + IOS software	If switch contains routing capability, where the supervisors run Catalyst software, and the Multilayer Switch Feature Card (MSFC) or Route Switch Module (RSM) runs IOS software.	This allows the Layer 2 switch functionality to be separate from the Layer 3 (and above) IOS functionality.
Native IOS	A single instance of IOS software is installed on the Catalyst Supervisor Engine, which also controls MSFC.	A single IOS kernel provides all Multilayer Switching functions (Layers 2 and above).

The IOS interface is used across a wide variety of Catalyst switch platforms, particularly the fixed and stackable switches, and is therefore the interface assumed through the remainder of this courseware. Labs may provide direction on the use of specific Catalyst software commands, depending on the equipment provided.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Many issues result from a poorly designed network.**
- **A hierarchical IP addressing scheme scales well in the Campus Infrastructure module.**
- **Select the best equipment, cabling, and interconnection technologies to connect devices.**
- **VLANs should map to the IP hierarchy and access links for the Campus Infrastructure module.**
- **Identify common campus traffic types.**
- **Identify common traffic sources and destinations.**
- **Identify the two interfaces used to configure Catalyst switches.**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—1-12

Module Summary

This topic summarizes the key points discussed in this module.

Module Summary

Cisco.com

- **The Campus Infrastructure module of the Enterprise Composite Network model provides a structure and set of components that can be used to build or expand a campus network.**
- **The Campus Infrastructure module provides a structure for the implementation of specific network technologies at various points in the module**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v32.2—1-1

In this module you learned best practices for network design. The ECNM provides a design pattern with integrity and resiliency for medium-to-large switched networks. The Campus Infrastructure module and building switch blocks provide an outstanding design to use on a physical campus network.

You have learned what decisions need to be made regarding technology components and other considerations when designing and implementing a campus network.

References

For additional information, refer to Cisco Systems, Inc., Cisco Product Advisor: http://cisco.com/en/US/products/prod_tools_index.html.

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which two of the following are benefits of the Enterprise Composite Network Model? (Choose two.) (Source: Describing the Campus Infrastructure Module)
- A) provides reliable communications between modules
 - B) offers more network integrity in network design
 - C) defines a deterministic network with clearly defined boundaries between modules
 - D) increases network scalability but complicates the design task
- Q2) Which module of the Enterprise Campus functional area contains e-mail and corporate servers providing applications, file, print, e-mail, and Domain Name System services to internal users? (Source: Describing the Campus Infrastructure Model)
- A) Campus Infrastructure module
 - B) Network Management module
 - C) Server Farm module
 - D) Enterprise DMZ module
- Q3) Which two of the following are considered appropriate interconnection technologies for servers in the Enterprise Composite Network Model? (Choose two.) (Source: Deploying Technology in the Campus Infrastructure Module)
- A) Fast Ethernet
 - B) Gigabit Ethernet
 - C) 10-Gigabit Ethernet
 - D) 10-Gigabit Etherchannel

Module Self-Check Answer Key

Q1) B, C

Q2) C

Q3) A, B

Defining VLANs

Overview

This module defines the purpose of Virtual Local Area Networks (VLANs) and describes how VLAN implementation can simplify network management and troubleshooting and can improve network performance. When VLANs are created, their names and descriptions are stored in a VLAN database that can be shared between switches. You will see how design considerations determine which VLANs will span all the switches in a network and which VLANs will remain local to a switch block. The configuration components of this module will describe how individual switch ports may carry traffic for one or more VLANs, depending on their configuration as access or trunk ports. This module explains both why and how VLAN implementation occurs in an enterprise network.

Module Objectives

Upon completing this module, you will be able to define VLANs, to segment network traffic, and to manage network use. This ability includes being able to meet these objectives:

- Configure access VLANs on access switches so that traffic is isolated to the individual VLANs as planned in the campus infrastructure design
- Configure trunk access and distribution switches so that a single trunk supports multiple VLANs
- Configure VTP modes and domains so VLAN configuration information is sent between switches as intended

Implementing VLANs

Overview

VLANs are used to create logical Layer 3 segments in a given network. A VLAN is considered a logical segment because the traffic it carries may traverse multiple physical network segments. This lesson will examine how switch ports can be statically configured to belong to one or more VLANs and how various ports on a single switch can belong to different VLANs. End-to-end VLANs will be differentiated from local VLANs. Local VLANs exist within the context of a single switch or switch stack, while end-to-end VLANs span multiple network segments interconnected by switches.

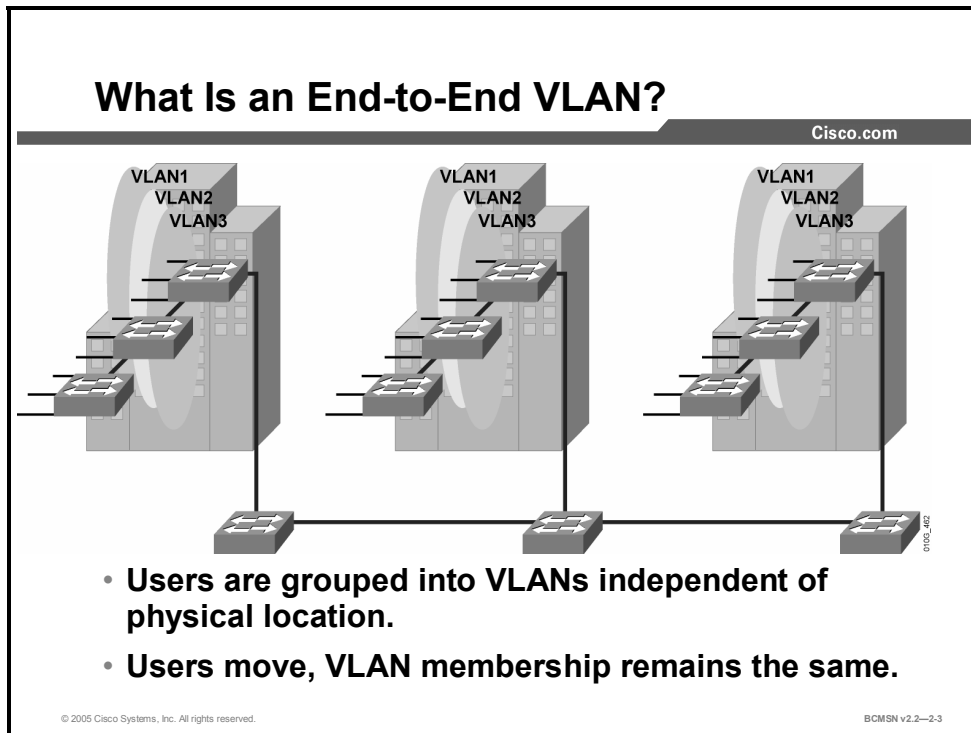
Objectives

Upon completing this lesson, you will be able to configure VLANs on access switches to confine traffic to individual VLANs in accordance with the Campus Infrastructure module design. This ability includes being able to meet these objectives:

- Define an end-to-end VLAN
- Define a local VLAN
- Describe the VLAN configuration modes and their functions
- Define a VLAN access port
- List the benefits of VLAN implementation in the Enterprise Network Composite Model
- Identify the commands to implement a VLAN
- List the steps needed to create a VLAN and associate it with an access port

What Is an End-to-End VLAN?

This topic describes VLANs that span multiple switches in a campus environment.



The term end-to-end VLAN refers to a single VLAN associated with switch ports that are widely dispersed throughout an enterprise network. Traffic for this VLAN is carried throughout the switched network. If many VLANs in a network are end-to-end, special links are required between switches to carry the traffic.

An end-to-end VLAN has these characteristics:

- The VLAN is geographically dispersed throughout the network.
- Users are grouped into the VLAN regardless of physical location.
- As a user moves throughout a campus, the VLAN membership of that user remains the same.
- Users are typically associated with a given VLAN for network management reasons.
- All devices on a given VLAN typically have addresses on the same IP subnet.

Because a VLAN represents a Layer 3 segment, end-to-end VLANs allow a single Layer 3 segment to be geographically dispersed throughout the network. Reasons for implementing this design might include the following:

- **Grouping users:** Users can be grouped on a common IP segment even though they are geographically dispersed.
- **Security:** A VLAN may contain resources that should not be accessible to all users on the network, or there may be a reason to confine certain traffic to a particular VLAN.
- **Applying quality of service (QoS):** Traffic from a given VLAN can be given higher or lower access priority to network resources.

- **Routing avoidance:** If much of the VLAN user traffic is destined for devices on that same VLAN and routing to those devices is not desirable, users can access resources on their VLAN without their traffic being routed off the VLAN, even though the traffic may traverse multiple switches.
- **Special purpose VLAN:** Sometimes a VLAN is provisioned to carry a single type of traffic that must be dispersed throughout the campus (for example, multicast, voice, or visitor VLANs).
- **Poor design:** For no clear purpose, users are placed in VLANs that span the campus or even WAN networks.

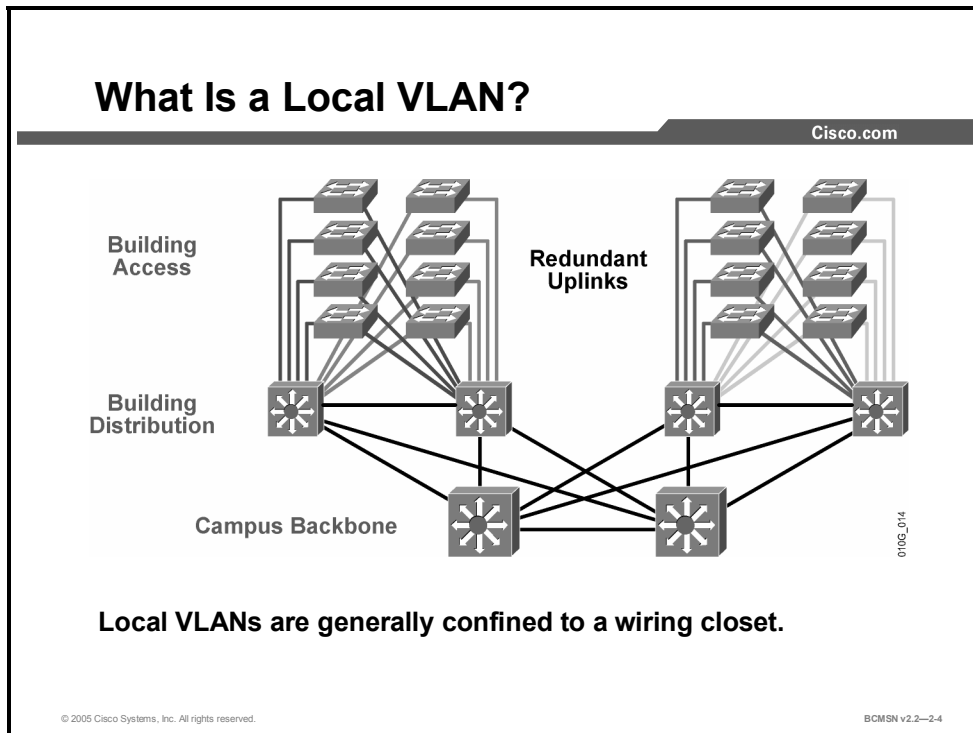
Some items should be considered when implementing end-to-end VLANs. Switch ports are provisioned for each user and associated with a given VLAN. Because users on an end-to-end VLAN may be anywhere in the network, all switches must be aware of that VLAN. This means that all switches carrying traffic for end-to-end VLANs are required to have identical VLAN databases. Also, flooded traffic for the VLAN is, by default, passed to every switch even if it does not currently have any active ports in the particular end-to-end VLAN. Finally, troubleshooting devices on a campus with end-to-end VLANs can be challenging, as the traffic for a single VLAN can traverse multiple switching in a large area of the campus.

Example: VLAN Implementation

In a military setting, one VLAN is designated to carry top-secret data. Users with access to that data are widely dispersed throughout the network. Because all devices on that VLAN have similar security requirements, security is handled by access lists at the Layer 3 devices that route traffic onto the segment (VLAN). Security can be applied VLAN-wide without addressing security at each switch in the network, which might have only a single user on the top-secret VLAN.

What Is a Local VLAN?

Here is an explanation of how a local VLAN differs from an end-to-end VLAN.



In the past, network designers attempted to implement the 80/20 rule when designing networks. The rule was based on the observation that, in general, 80 percent of the traffic on a network segment went between local devices, and only 20 percent of the traffic was destined for remote network segments. Designers now consolidate servers in central locations on the network and provide access to external resources such as the Internet through one or two paths on the network, as the bulk of traffic now traverses a number of segments. Therefore, the paradigm now is closer a 20/80 proportion, in which the greater flow of traffic leaves the local segment.

Additionally, the concept of end-to-end VLANs was very attractive when IP address configuration was a manually administered and burdensome process; therefore, anything that reduced this burden as users moved between networks was a good thing. But, given the ubiquity of DHCP, the process of configuring IP at each desktop is no longer a significant issue. As a result, there are few benefits to extending a VLAN throughout an enterprise. It is often more efficient to group all users of a set of geographically common switches into a single VLAN regardless of the organizational function of those users, especially from a troubleshooting perspective. VLANs that have boundaries based upon campus geography rather than organizational function are called “local VLANs.” Local VLANs are generally confined to a wiring closet.

Here are some local VLAN characteristics and use guidelines:

- Local VLANs should be created with physical boundaries in mind, rather than the job functions of the users on the end devices.
- Traffic from a local VLAN is routed to reach destinations on other networks.
- A single VLAN does not extend beyond the Building Distribution submodule.
- VLANs on a given access switch should not be advertised to all other switches in the network.

VLAN Configuration Modes

This topic identifies the steps and the commands for creating a VLAN in both global and database modes.

VLAN Configuration Modes

Cisco.com

Global Mode

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

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-2-5

VLANs are created in either global configuration or VLAN database mode on most Cisco IOS software-based switches. Global configuration mode is the preferred way of creating and managing VLANs because the user interface is familiar. When a VLAN is created or deleted, the change occurs as soon as the user hits the Enter key on the VLAN configuration line. The commands in this courseware will delineate VLAN creation and management using global configuration mode as shown in the figure.

Note Global configuration mode can be used to configure VLANs in the range 1-1005 and must be used to configure extended-range VLANs (VLAN IDs 1006 to 4094). The VLAN Trunk Protocol (VTP) configuration revision number is incremented each time a VLAN is created or changed.

VLAN Database Mode

This subtopic describes VLAN database mode.

VLAN Configuration Modes

Cisco.com

Database Mode

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

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

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—2-6

Alternatively, VLANs can be created and managed using VLAN database mode.

VLAN database mode is session oriented. When you add, delete, or modify VLAN parameters, the changes are not applied until you enter the **apply** or **exit** command. You can also exit VLAN database mode and not apply the changes by entering the **abort** command.

To access this mode, the **vlan database** command is executed from privileged EXEC mode. From this mode, you can add, delete, and modify VLAN configurations for VLANs in the range 1 to 1005.

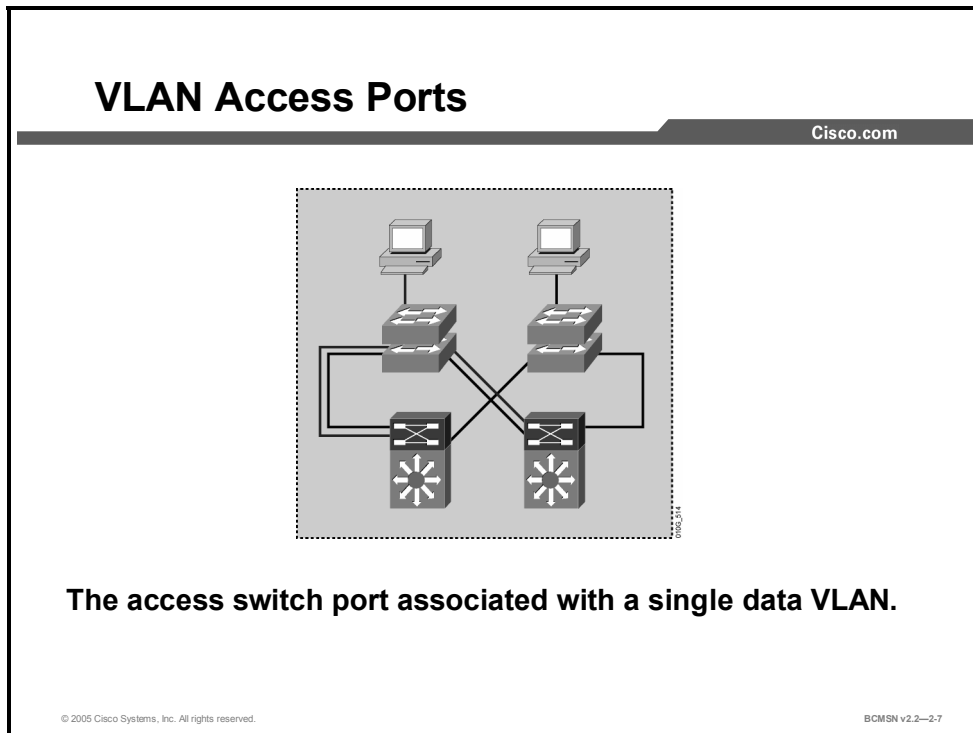
Example: Creating a VLAN in VLAN Database Mode

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

Note This mode has been deprecated and will be removed in some future release.

What Are VLAN Access Ports?

This topic describes VLAN access ports.



When an end system is connected to a switch port, it needs to be associated with a VLAN, in accordance with the network design. To associate a device with a VLAN, the switch port to which the device connects will be assigned to a single data VLAN and thus becomes an access port. A switch port can become an access port through static or dynamic configuration.

On most switches, VLAN membership results from execution of a specific **switchport** configuration command. In a local VLAN strategy, the switch port is associated with the VLAN of other devices on that same switch or switch cluster.

Attributes and characteristics of access ports:

- An access port is associated with a single VLAN.
- The VLAN to which the access port is assigned must exist in the VLAN database of the switch, or the port will be associated with an inactive VLAN that does not forward frames.
- Because an access switch port is part of a VLAN or Layer 2 domain, that port will receive broadcasts, multicasts, unicast floods, and so forth that are sent to all ports in the VLAN.
- The end device will typically have an IP address that is common to all other devices on the access VLAN.

Dynamic Access Port Association

Switch ports can be dynamically associated with a given VLAN based upon the MAC address of the device connecting on that port. This requires that the switch query a VLAN Membership Policy Server (VMPS) to determine what VLAN to associate with a switch port, when a specific source MAC address is seen on the switch port.

This might be beneficial with a set of workstations that rove throughout the enterprise. Regardless of what switch or switch port the workstation connected to, that switch port would become an access port on a single, specific VLAN. Some security situations may require dynamic VLAN associations. However, dynamic VLANs are not consistent with the Enterprise Composite Network Model and will not be discussed further in this course.

Benefits of Local VLANs in the Enterprise Composite Network Model

This topic outlines some benefits of implementing VLANs in an enterprise network.

Benefits of Local VLANs in the Enterprise Composite Network Model

Cisco.com

- **Deterministic traffic flow**
- **Finite failure domain**
- **High availability**
- **Ease of management**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-2-8

Local VLANs are part of the Enterprise Composite Network Model design. VLANs used at the access layer should extend no further than their associated distribution switch. Traffic is routed from the local VLAN as it is passed from the distribution layer into the core. This design can mitigate Layer 2 troubleshooting issues that occur when a single VLAN traverses the switches throughout an enterprise campus network. Implementing the Enterprise Composite Network Model using local VLANs provides the following benefits:

- **Deterministic traffic flow:** The simple layout provides a predictable Layer 2 and Layer 3 traffic path. In the event of a failure that was not mitigated by the redundancy features, the simplicity of the model facilitates expedient problem isolation and resolution within the switch block.
- **Finite failure domain:** If VLANs are local to a switch block and the number of devices on each VLAN is kept small, failures at Layer 2 are confined to a small subset of users.
- **High availability:** Redundant paths exist at all infrastructure levels. Local VLAN traffic on access switches can be passed to the building distribution switches across an alternative Layer 2 path in the event of primary path failure. Redundant Layer 3 protocols can provide failover should the default gateway for the access VLAN fail. When both the Spanning Tree Protocol (STP) instance and VLAN are confined to a specific access and distribution block, then Layer 2 and Layer 3 redundancy measures and protocols can be configured to failover in a coordinated manner.
- **Ease of management:** Local VLANs, typically confined to the Building Access submodule, are easier to plan and manage than VLANs spanning various switches and network areas. Also, local VLANs, when used in combination with dynamically assigned

IP addresses, allow workstations to move from one VLAN to another with limited administrative overhead.

VLAN Implementation Commands

This topic describes commands used to implement VLANs.

VLAN Implementation Commands

Cisco.com

Configuring VLANs

- **vlan 101**
- **switchport mode access**
- **switchport access vlan 101**

Verifying VLANs

- **show interfaces**
- **show vlan**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-2.9

The “VLAN Implementation Commands” table describes the primary commands used to implement VLANs and to verify their configuration in the Catalyst switch IOS interface.

VLAN Implementation Commands

Command	Description
<code>switch(config)# vlan <i>vlan-id</i></code>	Creates VLAN <i>vlan-id</i> .
<code>switch(config)# no vlan <i>vlan-id</i></code>	Deletes VLAN <i>vlan-id</i> .
<code>switch(config-vlan)# name <i>vlan_name</i></code>	Assigns a specific name to the VLAN.
<code>switch(config-if)# switchport mode access</code>	Specifies this port is to function at Layer 2 and places the port in VLAN access mode. Typically followed by the switchport access vlan <i>vlan-id</i> command.
<code>switch(config-if)# switchport access vlan <i>vlan-id</i></code>	Associates a single switch port as an access port to a single VLAN <i>vlan-id</i> .
<code>switch(config-if)# no shutdown</code>	Enables a switch port or interface.
<code>switch(config)# interface vlan <i>vlan-id</i></code>	Creates or moves to interface configuration mode for a switch virtual interface on VLAN <i>vlan-id</i> . IP addresses can be assigned and other commands can be executed as they would be on any Layer 3 router interface.

Command	Description
<pre>switch(config)# shutdown vlan <i>vlan-id</i></pre>	Suspends local traffic on the specified VLAN. Does not change the VLAN information in the VTP database, and the switch still advertises VTP information.
<pre>switch# show vlan</pre>	Displays the parameters for all configured VLANs or for one VLAN on the switch.

How to Implement a VLAN

This topic describes how to implement a VLAN.

How to Implement a VLAN

Cisco.com

- **Create or configure a VLAN.**
- **Verify VLAN configuration**
- **Associate switch ports with the VLAN.**
- **Verify switch port configuration.**
- **Test VLAN connectivity.**
- **Implement VLAN and switch security.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—2-10

To create or configure a VLAN and associate switch ports, follow these steps:

- Step 1** Create the VLAN.
- Step 2** Verify the VLAN configuration.
- Step 3** Associate switch ports with the VLAN.
- Step 4** Verify the switch port configuration.
- Step 5** Test VLAN connectivity.
- Step 6** Implement switch and VLAN security measures.

These steps are explained in greater detail in the remainder of this topic.

1. Create or Configure a VLAN

Configuring an Access VLAN

Cisco.com

```
Switch(config)# vlan vlan_id
```

Create a VLAN.

```
Switch(config-vlan)# description vlan_description
```

Provide a VLAN description.

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

Place the switch port into access mode.

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

Associate the access switch port with a VLAN.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2--2-11

Before assigning a switch port to a specific VLAN, the VLAN may need to be created. The example below shows the syntax for creating a VLAN using the Cisco IOS interface.

To create a VLAN or enter VLAN configuration mode, use the **vlan** command:

```
Switch(config)# vlan vlan_id
```

VLAN Creation Arguments

Argument	Description
<i>vlan_id</i>	Required: Any valid VLAN number from 1-4094 if accepted by the switch platform, 1-1024 if not. If VLAN does not presently exist, a VLAN with this <i>vlan_id</i> will be created and prompt will change to VLAN config mode. If the VLAN already exists, prompt will change and VLAN parameters can be altered for this <i>vlan_id</i> .

To enter a description of the VLAN from the VLAN configuration mode:

```
Switch(config-vlan)# description vlan_description
```

2. Verify VLAN Configuration

Verifying the Access VLAN Configuration

Cisco.com

```

Switch#show vlan

VLAN Name                Status   Ports
-----
1    default                 active   Fa0/1, Fa0/2, Fa0/3, Fa0/4
                    Fa0/5, Fa0/7, Fa0/9

11   asw11_data               active
12   asw12_data               active
95   VLAN0095                 active   Fa0/8
99   Trunk Native             active
100  Internal Access          active
111  voice-for-group-11       active
112  voice-for-group-12       active
1002 fddi-default              act/unsup
1003 token-ring-default      act/unsup
1004 fddinet-default          act/unsup
1005 trnet-default            act/unsup

VLAN Type  SAID      MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1
-----
1    enet  100001   1500  -     -     -     -     -         0
11   enet  100011   1500  -     -     -     -     -         0
. . .
. . .
. . .
    
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—2-12

Show VLAN

Execute the **show vlan** command from privileged EXEC mode. It displays information about a particular VLAN. The fields in the **show vlan** command output are described in the table.

Field	Description
VLAN	VLAN number
Name	Name of the VLAN, if configured
Status	Status of the VLAN (active or suspended)
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
AREHops	Maximum number of hops for all-routes explorer frames
STEHops	Maximum number of hops for spanning tree explorer frames

3. Associate Switch Ports with the VLAN

Switch ports that are to function at Layer 2 and carry traffic for a single VLAN are configured as access switch ports and are assigned an access VLAN.

To configure a Layer 2 switch port as an access port:

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

Switch Port Parameters

Parameter	Description
switchport	Required: Configures the interface to function as a Layer 2 port only. On many switches, this is the default. No switchport would reverse this process and, on some switch platforms, convert the port to a Layer 3 port.
mode access	Required: Switch port must be configured in access mode if you will next assign a specific access VLAN. Alternative mode options are available for nonaccess port functionality.

To assign the access port to a specific VLAN:

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

Switch Port Access Parameters

Parameter	Description
switchport	Required: Indicates further configuration of Layer 2 functionality of switch port.
access	Required: Indicates further configuration of access features of the switch port.
vlan vlan_id	Required: Indicates what single VLAN number is to be associated with this access port. On some switch platforms, this command will create a VLAN, not just associate an ID.

4. Verify Switch Port Configuration

The following commands are useful for verifying that a switch port is configured as intended:

```
show interface type slot/port switchport  
show running-config interface type slot/port  
show vlan
```

Show Running-Config 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
```

5. Test VLAN Connectivity

After placing a device on the configured switch port, these steps will help verify if the device is connecting to the VLAN as intended:

- Step 1** Ensure that the connected device has a correctly configured IP address and a subnet mask that places it on the same network as the default gateway.
- Step 2** Ping the default gateway.
- Step 3** If the ping to default gateway is successful, the VLAN configuration and the IP address configuration have been verified.

6. Implement Switch and VLAN Security Measures

When implementing VLANs, you should consider a few measures to secure the VLAN and the switch itself. The security policy of the organization will likely have more detailed recommendations, but these can provide a foundation. Security will be covered in more detail in the “Securing Your Multilayer Network to Minimize Service Loss and Data Theft” module.

- Create a “parking-lot” VLAN with a VLAN ID other than VLAN1 and place all unused switch ports in this VLAN. This VLAN may provide the user with some minimal network connectivity. (Check on the security policy of your organization before implementing.)
- Disable unused switch ports, depending on the security policy of the organization.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **An end-to-end VLAN may span an entire network. A local VLAN is isolated to a single switch.**
- **VLANs solve issues that arise in a Layer 2 switched network.**
- **VLANs can be configured globally or in VLAN database mode.**
- **An access switch port is associated with one VLAN.**
- **The design of the Enterprise Composite Model offers several benefits for VLAN implementation.**
- **Cisco provides a series of commands to configure VLAN and verify configuration on an access switch.**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—2-13

Supporting Multiple VLANs on a Single Trunk

Overview

Switch ports carrying traffic for multiple VLANs are called trunk ports. As frames from multiple VLANs traverse trunk ports, the switch must identify each frame to associate it with a given VLAN. This lesson will examine the differences between Inter-Switch Link (ISL) and 802.1Q, two protocols used to mark frames on a trunk link.

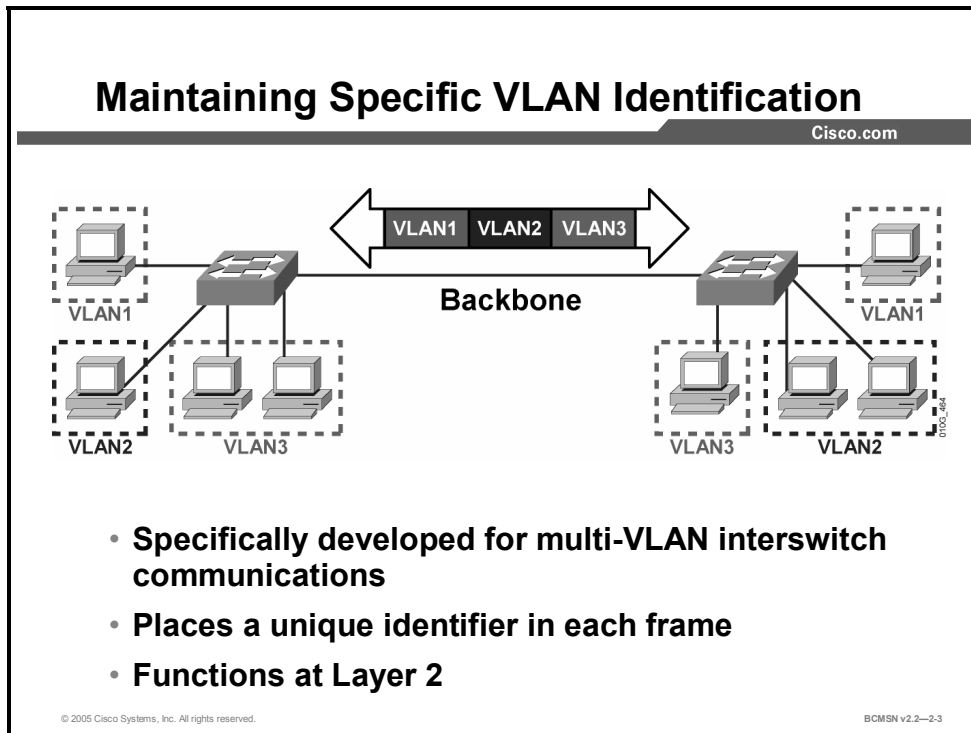
Objectives

Upon completing this lesson, you will be able to use appropriate commands to configure trunk access and distribution switches to carry multiple VLANs across a single link. This ability includes being able to meet these objectives:

- Define the purpose of a VLAN trunk link
- Define the purpose of a VLAN trunking protocol
- Identify key differences between ISL and 802.1Q trunking protocols
- Describe the ISL trunking protocol
- Describe the 802.1Q trunking protocol
- Define an 802.1Q native VLAN
- Describe VLAN ranges and their use
- Identify the purpose of DTP and its primary modes
- Identify the commands used to configure and verify ISL and 802.1Q trunking
- Configure ISL and 802.1Q trunking
- Identify key issues when implementing trunk links in the Campus Infrastructure module
- List key problems that result from trunk link configuration

What Is a VLAN Trunk?

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



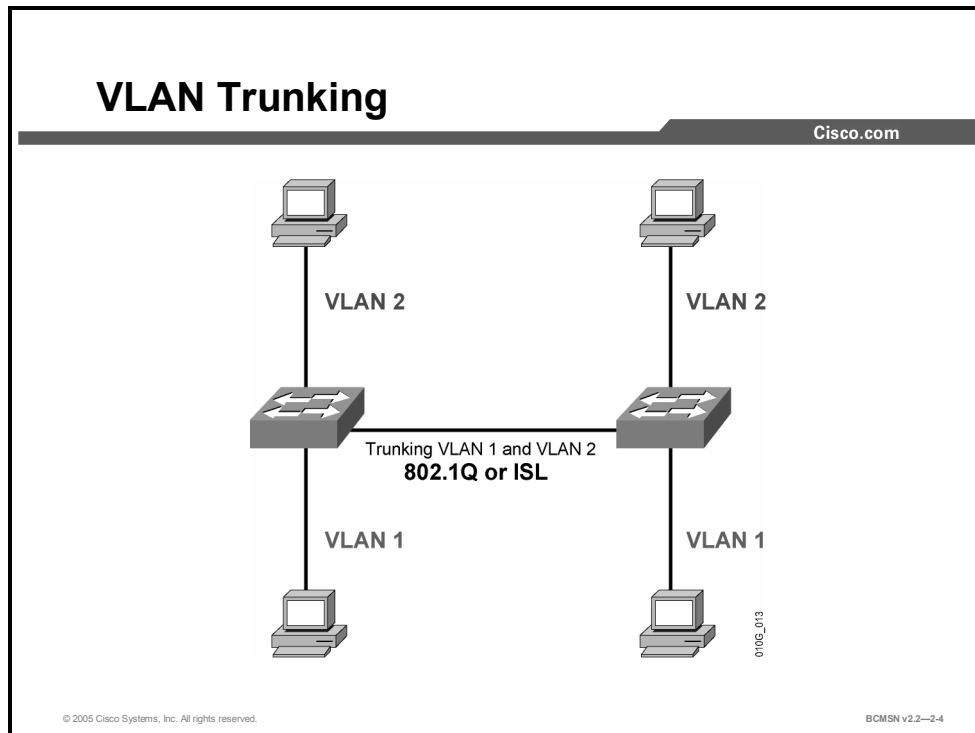
Multiple VLANs are supported between switches through the use of VLAN trunks. A trunk is a Layer 2 link between switches which are running a specialized trunking protocol. Trunks carry the traffic of multiple VLANs over physical links (multiplexing) and enable the extension of a single Layer 2 VLAN between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

If frames from a single VLAN traverse a trunk link, a trunking protocol must mark the frame to identify its associated VLAN as the frame is placed onto the trunk link. The receiving switch then knows the frame's VLAN of origin and can process the frame accordingly.

On the receiving switch, the VLAN ID (VID) is removed when the frame is forwarded onto an access link associated with its VLAN.

What Is a VLAN Trunking Protocol?

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



A special protocol is required to establish a trunk link between two devices. A trunk link may exist between these devices:

- Two switches
- A switch and a router
- A switch and a trunk-capable Network Interface Card (NIC) in a node such as a server

If a single physical link carries traffic for multiple VLANs, each frame must be “marked” with a VID so it is differentiated from frames coming from other VLANs. This marking or frame identification is accomplished through the implementation of a trunking protocol. Frame identification uniquely assigns an ID, referred to as a VID, to each frame. Each receiving switch examines this VID to determine the destination VLAN of the frame.

VLAN IDs are only associated with frames traversing a trunk link. When a frame enters or exits the switch on an access link, no VID is present. The ASIC on the switch port assigns the VID to a frame as it is placed on a trunk link and also strips off the VID if the frame exits an access switch port.

Trunk links should be managed so that they carry only traffic for intended VLANs. This practice keeps unwanted VLAN data traffic from traversing links unnecessarily. Trunk links are used between the access and distribution layers of the campus switch block. These are the trunk protocols used to carry multiple VLANs over a single link:

- ISL: Cisco Inter-Switch Link
- 802.1Q: IEEE standard trunking protocol

Comparing ISL and 802.1Q Trunking Protocols

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

Comparing ISL and 802.1Q	
ISL	802.1Q
Proprietary	Nonproprietary
Encapsulated	Tagged
Protocol independent	Protocol dependant
Encapsulates the old frame in a new frame	Adds a field to the frame header

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-2-5

Depending on the trunking protocol, data frames sent across a trunk link are either encapsulated or tagged. The purpose of encapsulating or tagging frames is to provide the receiving switch with a VID to identify the VLAN from which the frame originated. The trunking protocol ISL, a Cisco proprietary protocol, encapsulates frames, while IEEE 802.1Q inserts a tag into the original Layer 2 data frame.

802.1Q is not proprietary and can be deployed in any Ethernet standards-based Layer 2 device. It is specific to a single Layer 2 protocol (Ethernet) because it modifies the Layer 2 frame by inserting a tag between two specific fields of the frame and therefore must be aware of the frame header details.

ISL is Layer 2–protocol independent. Because the original Layer 2 frame is fully encapsulated and not altered, ISL can transport data frames from various Layer 2 media types.

ISL Trunking Protocol

This topic describes the ISL protocol.

Trunking with Inter-Switch Link

Cisco.com

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

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—2-6

ISL is a Cisco proprietary protocol option for configuring Layer 2 trunk links. It is the original standard for trunking between switches and predates IEEE trunking standards. ISL takes original Layer 2 frames and encapsulates them with a new header and trailer before placing them on the trunk link.

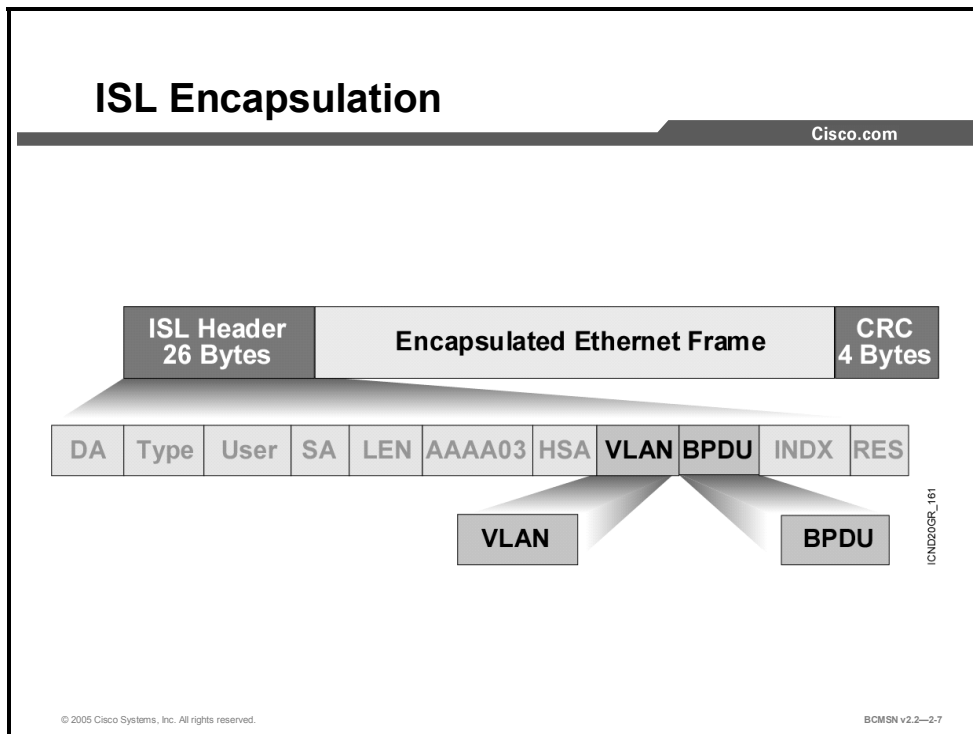
Because an entirely new header is appended to the original frame, the header offers some features not found in 802.1Q, an alternative trunking protocol.

The following are some benefits of the ISL protocol:

- It supports multiple Layer 2 protocols (Ethernet, Token Ring, FDDI, and ATM).
- It supports per-VLAN Spanning Tree Protocol
- The encapsulation process leaves original frames unmodified, less prone to error, and more secure.
- It has a large installation base.

ISL Encapsulation Process

This subtopic describes how the ISL trunking protocol encapsulates frames.



When a switch port is configured as an ISL trunk port, the entire original Layer 2 frame, including header and frame check sequence (FCS) trailer, will be encapsulated before it traverses the trunk link. Encapsulation is the process of placing an additional header in the front and a trailer at the end of the original Layer 2 frame. The ISL header will contain the VID of the VLAN where the frame originated. At the receiving end, the VID is read, the header and trailer are removed, and the original frame is forwarded like any regular Layer 2 frame on that VLAN.

Only ISL trunk ports can properly receive ISL encapsulated frames. A non-ISL port receiving an ISL frame may consider the frame size to be invalid or may not recognize the fields in the header. The frame will likely be dropped and counted as a transmission error when received by a non-ISL port.

ISL Header

The ISL header contains various fields with values that define attributes of the original Layer 2 data within the encapsulated frame. This information is used for forwarding, media identification, and VLAN identification. The population of the fields within the ISL header varies, based on the type of VLAN and the media of the link. The ASIC on an Ethernet port encapsulates the frames with a 26-byte ISL header and a 4-byte FCS. This 30-byte ISL encapsulation overhead is consistent among the Layer 2 protocols supported on Catalyst switches, but the overall size of the frame will vary and be limited by the MTU of the original Layer 2 protocol.

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 to the receiver that the packet is in ISL format.
- **Type:** 4-bit descriptor of the encapsulated frame types: Ethernet (0000), Token Ring (0001), FDDI (0010), and ATM (0011).
- **User:** 4-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 cyclic redundancy check (CRC).
- **AAAA03:** Standard Subnetwork Access Protocol (SNAP) 802.2 logical link control (LLC) header.
- **HSA (high bits of source address):** First 3 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):** 1-bit descriptor identifying whether the frame is a spanning tree BPDU. It also identifies if the encapsulated frame is a Cisco Discovery Protocol (CDP) or VLAN Trunk Protocol (VTP) frame and indicates if the frame should be sent to the control plane of the switch.
- **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).

ISL Trailer

The trailer portion of the ISL encapsulation is an FCS that carries a CRC value calculated on the original frame plus the ISL header as the ISL frame was placed onto the trunk link. The receiving ISL port recalculates this value. If the CRC values do not match, the frame is discarded. If the values match, the switch discards the FCS as a part of removing the ISL encapsulation so that the original frame can be processed. The ISL trailer consists of these fields:

- **FCS:** Consists of 4 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 for 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.

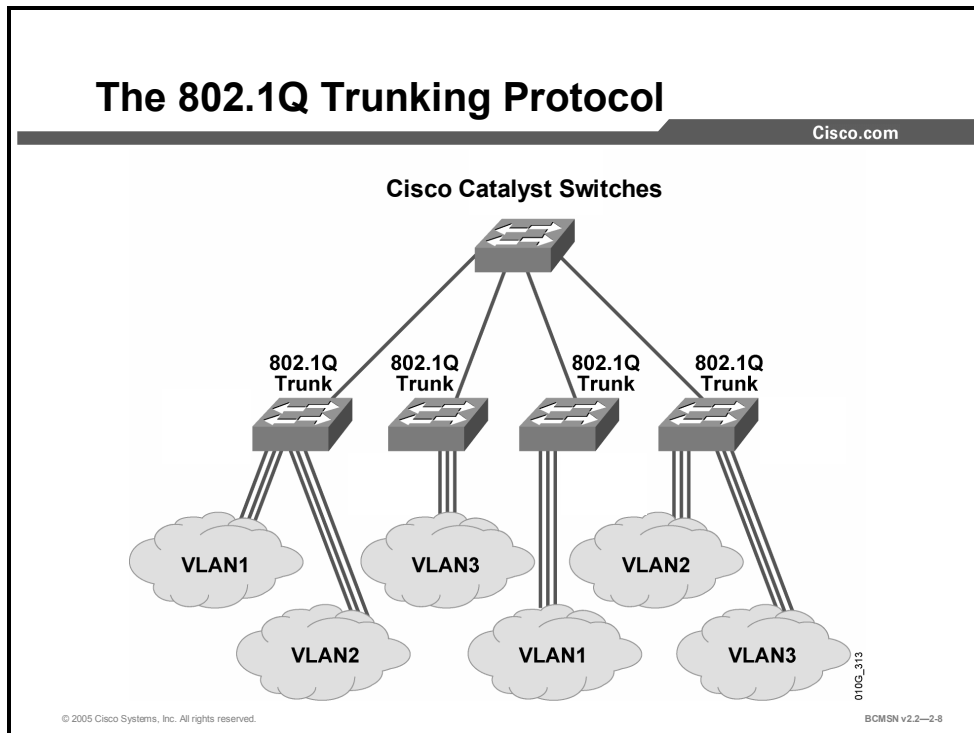
References

For more information, refer to this resource:

http://www.cisco.com/en/US/tech/tk389/tk390/technologies_tech_note09186a0080094665.shtml#field1

802.1Q Trunking Protocol

This topic describes the 802.1Q trunking protocol.



Like ISL, 802.1Q is a protocol that allows a single physical link to carry traffic for multiple VLANs. It is the IEEE standard VLAN trunking protocol. Rather than encapsulating the original Layer 2 frame in its entirety, 802.1Q inserts a tag into the original Ethernet header, then recalculates and updates the FCS in the original frame and transmits the frame over the trunk link.

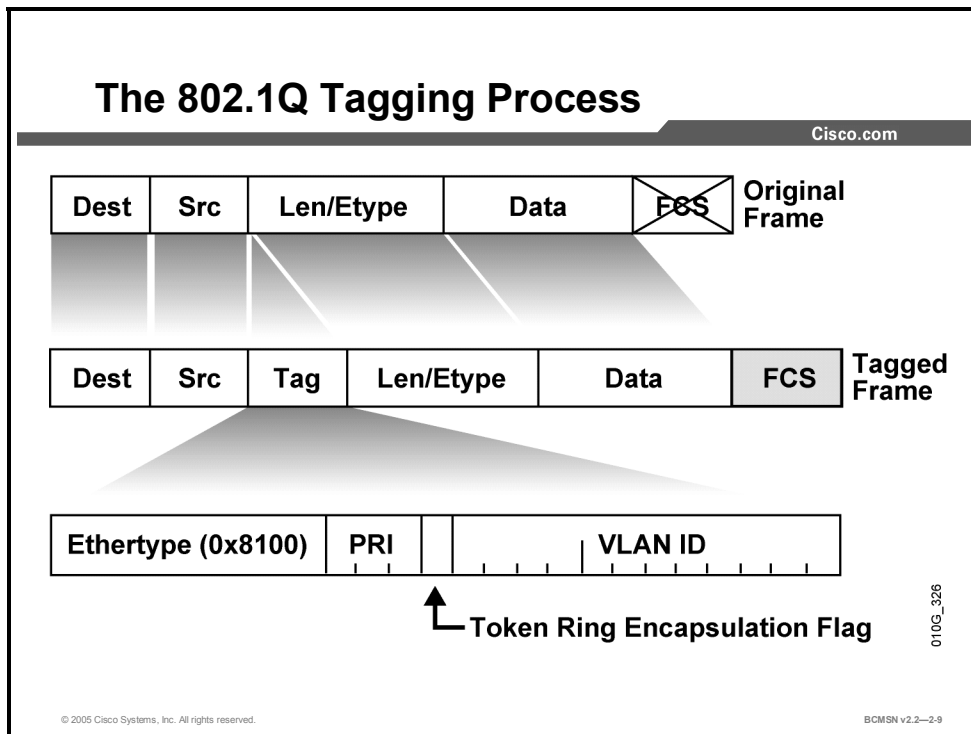
The 802.1Q protocol, often referred to as “dot-1Q,” offers the clear benefit of being the first IEEE standards-based trunking protocol for Ethernet. It allows multiple VLANs to traverse infrastructure equipment where cross vendor links exist.

Following are some additional benefits of the 802.1Q protocol:

- Support for Ethernet and Token Ring
- Support for 4096 VLANs
- Support for Common Spanning Tree (CST), Multiple Spanning Tree (MST), and Rapid Spanning Tree (RST)
- Point-to-multipoint topology support
- Support for untagged traffic over the trunk link via native VLAN
- Extended quality of service (QoS) support
- Growing standard for IP telephony links

802.1Q Tagging Process

This subtopic describes the 802.1Q process for tagging frames traversing a trunk link.



To identify a frame with a given VLAN, the 802.1Q protocol adds a tag, or a field, to the standard Layer 2 Ethernet data frame. The components of this tag are shown in the figure. Because inserting the tag alters the original frame, the switch must recalculate and alter the CRC value for the original frame before sending it out the 802.1Q trunk port. In contrast, ISL does not modify the original frame at all.

The new 802.1Q Tag field has the following components:

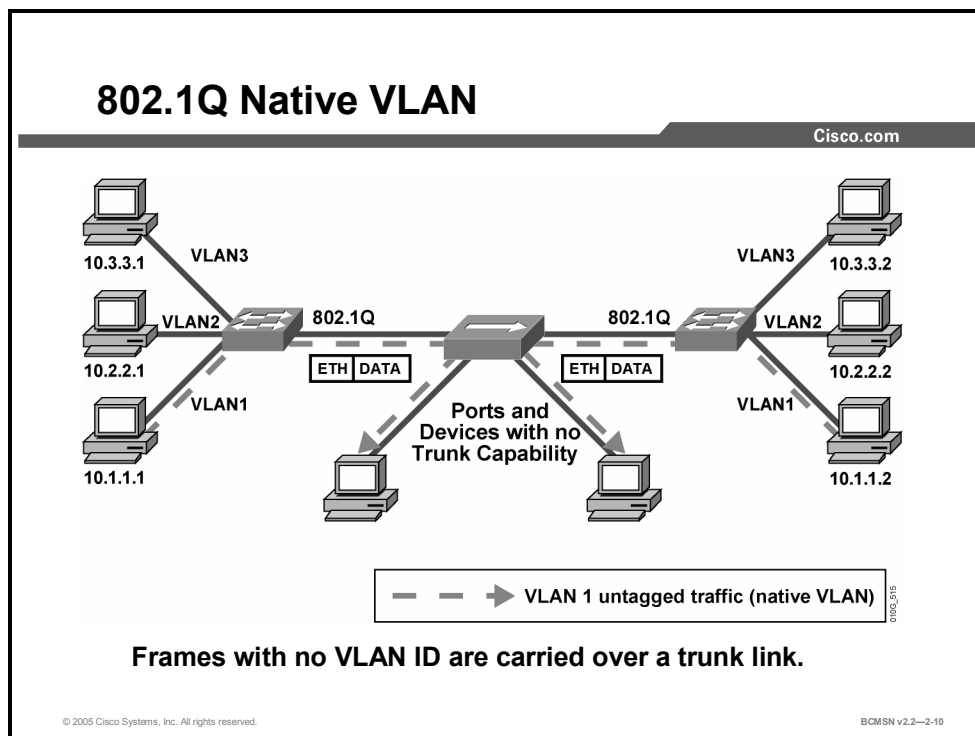
- **PRI:** 3 bits; carries priority information for the frame.
- **Token Ring Encapsulation Flag:** Indicates the canonical interpretation of the frame if it is passed from Ethernet to Token Ring.
- **VLAN ID:** VLAN association of the frame. By default, all normal and extended-range VLANs are supported.

If a non-802.1Q-enabled device or an access port receives an 802.1Q frame, the tag data is ignored, and the packet is switched at Layer 2 as a standard Ethernet frame. This allows for the placement of Layer 2 intermediate devices, such as other switches or bridges, along the 802.1Q trunk path. To process an 802.1Q tagged frame, a device must allow an MTU of 1522 or higher.

Note An Ethernet frame that has a larger MTU than expected (1518 by default for Ethernet) but no larger than 1600 bytes will register as a Layer 2 error frame called a “baby giant.” For ISL, the original frame plus ISL encapsulation can generate a frame as large as 1548 bytes and 1522 bytes for an 802.1Q tagged frame.

What Is an 802.1Q Native VLAN?

This topic describes the purpose of the native VLAN carried on an 802.1Q trunk link.



When configuring an 802.1Q trunk, a matching native VLAN must be defined on each end of the trunk link. A trunk link is inherently associated with tagging each frame with a VID. The purpose of the native VLAN is to allow frames not tagged with a VID to traverse the trunk link. An 802.1Q native VLAN is defined as one of the following:

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

Compare this to ISL, in which no frame may be transported on the trunk link without encapsulation, and any unencapsulated frames received on a trunk port are immediately dropped.

Each physical port has a parameter called a port VLAN ID (PVID). Every 802.1Q port is assigned a PVID value equal to the native VID. When a port receives a tagged frame that is to traverse the trunk link, the tag is respected. For all untagged frames, the PVID is considered the tag. This allows the frames to traverse devices that may be unable to read VLAN tag information.

Native VLANs have the following attributes:

- A trunk port will support only one native active VLAN per operational mode. The modes are access and trunk.
- By default, on Catalyst switches, all switch ports and native VLANs for 802.1Q are 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 to operate correctly.
- If the native VLAN is misconfigured for trunk ports on the same trunk link, Layer 2 loops can occur due to diverting Spanning Tree Protocol (STP) BPDUs from their correct VLAN.

Example: Native VLAN Implementation—Two End Devices on the Same Switch Port

A standard situation in which the native VLAN of 802.1Q might be used is when a single switch port supports traffic to an IP phone that then provides a connection to a PC. The port must be configured as 802.1Q so that the Layer 2 header allows the QoS marking to populate the priority (PRI) bits for the telephony traffic. A standard Ethernet packet provides no field for this marking.

The traffic arriving on the switch port from the IP phone will be tagged with VLAN information. The PC traffic arriving on the same switch port will not be tagged. The VID for the telephony traffic arriving on the 802.1Q trunk port will be respected. The PC traffic arriving with no tag will traverse the native VLAN.

Issues with 802.1Q Native VLANs

This subtopic addresses considerations that must be made when configuring the native VLAN on an 802.1Q trunk link.

802.1Q Native VLAN Considerations

Cisco.com

- **Native VLAN must match at ends of trunk.**
- **Consider default native VLAN.**
- **Consider traffic to be carried on native VLAN.**
- **CDP may have issues over native VLAN.**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2--2-11

The following issues need to be considered when configuring the native VLAN on an 802.1Q trunk link:

- The native VLAN interface configurations must match at both ends of the link or the trunk may not form.
- By default, the native VLAN will be VLAN1. For the purpose of security, the native VLAN on a trunk should be set to a specific VID that is not used for normal operations elsewhere on the network.
- If there is a native VLAN mismatch on an 802.1Q link, CDP, if used and functioning, will issue a “VLAN mismatch” error.
- On select versions of Cisco IOS software, CDP may not be transmitted or will be automatically turned off if VLAN1 is disabled on the trunk.
- If there is a native VLAN mismatch on either side of an 802.1Q link, Layer 2 loops may occur.
- When troubleshooting VLANs, note that a link can have one native VLAN association when in access mode, and another native VLAN association when in trunk mode.

VLAN Ranges

This topic describes VLAN ranges and their use.

VLAN Ranges	
Cisco.com	
VLAN Range	Use
0, 4095	Reserved for system use only
1	Cisco default
2–1001	For Ethernet VLANs
1002–1005	Cisco defaults for FDDI and Token Ring
1006–4094	Ethernet VLANs only, unusable on specific legacy platforms

Each VLAN on the network must have a unique VID. The valid range of user-configurable ISL VLANs is 1 to 1024. The valid range of VLANs specified in the IEEE 802.1Q standard is 0 to 4094. This table describes VLAN ranges and their use.

VLAN Ranges

VLAN Ranges	Range	Use	VTP Propagated
0, 4095	Reserved	For system use only. VLANs cannot be seen or used.	—
1	Normal	Cisco default VLAN. This VLAN can be used but not modified or deleted.	Yes
2–1001	Normal	These VLANs can be created, used, and deleted.	Yes
1002–1005	Normal	Cisco defaults for FDDI and Token Ring. These cannot be deleted.	Yes

VLAN Ranges	Range	Use	VTP Propagated
1006–4094	Extended	<p>For Ethernet VLANs only.</p> <p>Layer 3 ports and some software features require internal VLANs. Internal VLANs are allocated from 1006 and up. You cannot use a VLAN that has been allocated for such use. To display the VLANs used internally, enter the show vlan internal usage command.</p> <p>Switches running Catalyst product family software do not support configuration of VLANs 1006-1024. If you configure VLANs 1006-1024, ensure that the VLANs do not extend to any switches running Catalyst product family software.</p> <p>You must enable the extended system ID to use extended-range VLANs.</p>	No

In a network environment with non-Cisco devices connected to Cisco switches through 802.1Q trunks, 802.1Q VLAN numbers greater than 1000 must be mapped to ISL VLAN numbers on the Cisco switches. 802.1Q VLANs in the range 1 to 1001 are automatically mapped by VTP to a corresponding ISL VLAN. 802.1Q VLAN numbers greater than 1006 must be mapped to an ISL VLAN to be recognized and forwarded by VTP. Alternatively, configure VTP in transparent mode, to allow the use of extended system IDs and the manual management of VLAN configuration information between switches.

As a best practice, assign extended VLANs beginning with 4094 and work downward, as some switches use extended-range VLAN IDs for internal use starting at the low end of the extended range. Refer to "Configuring Extended-Range VLANs" in the software configuration guide associated with your switch platform and software release.

Identifying the Modes for Dynamic Trunking Protocol

This topic identifies the functions of the various modes for Dynamic Trunking Protocol (DTP).

DTP and 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.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—2-13

Trunk links should be configured statically whenever possible; however, Catalyst switch ports run DTP, which can automatically negotiate a trunk link. This Cisco proprietary protocol can determine an operational trunking mode and protocol on a switch port when connected to another device that is also capable of dynamic trunk negotiation.

A DTP mode can be configured to turn the protocol off or to instruct it to negotiate a trunk link only under certain conditions, as described in the figure.

Note General best practice is to set trunking to “on” and “nonegotiate” when a trunk link is required. DTP should be turned off on links where trunking is not intended.

Trunking Configuration Commands

This topic describes the commands used to configure trunking.

Trunking Configuration Commands

Cisco.com

Configuring a Trunk

- **switchport trunk**
- **switchport mode**
- **switchport nonegotiate**

Verifying a Trunk

- **show running-config**
- **show interfaces switchport**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2–2-14

Commands for configuring a trunk will vary depending on the operating system version of your switch. The commands shown here are for a Cisco IOS software-based switch.

Trunking Commands

Command	Description
Switch(config)# interface <i>number</i>	Selects the interface to configure.
Switch(config-if)# switchport trunk [<i>allowed vlan range or list</i>]	Sets the trunk characteristics when the interface is in trunking mode. Use the no form of this command to reset all trunking characteristics to the defaults. A range of VLANs to be carried on the trunk can optionally be specified.
Switch(config-if)# switchport trunk native vlan	Sets the native VLAN for sending and receiving untagged traffic when the interface is in 802.1Q trunking mode. Valid IDs are from 1 to 4094 when the Enhanced Image (EI) software is installed and 1 to 1001 when the Standard Image (SI) software is installed. Do not enter leading zeros.
Switch (configure-if)# switchport nonegotiate	Enables trunking on this port regardless of DTP messages received by a neighboring switch on this link.
Switch(config-if)# switchport mode	Configures the mode of a port to trunk or access.
Switch# show interfaces switchport	Displays the administrative and operational status of a switching (nonrouting) port.

How to Configure Trunking

This topic describes the steps for configuring an ISL and an 802.1Q trunk.

How to Configure Trunking

Cisco.com

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.
7. Use the **no shutdown** command on the interface to activate the trunking process.
8. Verify the trunk configuration.

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—2-15

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

- Step 1** Enter interface configuration mode.
- Step 2** Shut down the interface to prevent the possibility of premature autoconfiguration.
- Step 3** Select the trunking encapsulation. Note that 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 for 802.1Q links. This number *must* match at both ends of an 802.1Q trunk.
- Step 6** Configure the allowable VLANs for this trunk. This is necessary if VLANs are restricted to certain trunk links. This is best practice with the Enterprise Composite Network Model and leads to the correct operation of VLAN interfaces.
- Step 7** Use the **no shutdown** command on the interface to activate the trunking process.
- Step 8** Verify the trunk configuration using **show** commands.

Configuring an ISL Trunk

This subtopic discusses the configuration of an ISL trunk.

ISL Trunk Configuration

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 trunk native vlan 99
Switch(config-if)#switchport mode trunk
Switch(config-if)#switchport nonegotiate
Switch(config-if)#no shutdown
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2--2-16

In the example, interface Fast Ethernet 2/1 has been configured as a trunk link for ISL that is permanently on. DTP negotiation is not allowed. The trunk link will carry VLAN traffic for VLANs 1-5 and 1002-1005. VLANs 2-5 are configured on various access ports on the switch, and the trunk links need to carry the frames for these VLANs as well as the frames for the system VLANs 1, 1002-1005.

Configure Switch Port As ISL Trunking Port

Step	Action	Notes
1.	Enter interface configuration mode. Switch(config)# interface { fastethernet gigabitethernet } <i>slot/port</i>	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 <i>isl</i> or <i>dot1q</i> argument. Negotiate is the default. This command is only supported if the switch hardware supports both ISL and dot-1Q encapsulation.
3.	Configure the allowable VLANs for this trunk. Switch(config-if)# switchport trunk allowed vlan { add except all remove } <i>vlan_num1[,vlan_num[,vlan_num[,...]]]</i>	If not specified, all VLANs are allowed on the trunk. VLANs can be specifically allowed or disallowed.

Step	Action	Notes
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 It is best practice to shut down an interface while configuring trunking attributes so that premature autonegotiation cannot occur.

Configuring a Port for ISL Trunking with No DTP

When configuring the Layer 2 trunk to not use DTP, the following syntax is used so that the trunk mode is set to “on” and no DTP messages are sent on the interface.

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

This example shows how to configure a port for ISL trunking with no DTP:

```
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 trunk native vlan 99
Switch(config-if)#switchport nonegotiate
Switch(config-if)#no shutdown
Switch(config-if)#end
Switch#exit
```

Verifying the ISL Trunk Configuration

This subtopic describes how to verify an ISL trunk link.

Verifying ISL Trunking

Cisco.com

```
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	trunk	isl	trunking	99

```
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
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2--2-17

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

Configuring an 802.1Q Trunk

This subtopic describes how to configure an 802.1Q trunk link.

802.1Q Trunk Configuration

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 trunk native vlan 99
Switch(config-if)#switchport nonegotiate
Switch(config-if)#no shutdown
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-2-18

The example shows how to configure interface Fast Ethernet 5/8 as an 802.1Q trunk. Frames from VLANS 1, 5, 11, and 1002-1005 will be allowed to traverse the trunk link. The switchport mode for the interface is trunk (on), and there are to be no DTP messages sent on the interface.

Configure Switch Port as 802.1Q Trunk Link

Step	Action	Notes
1.	Enter interface configuration mode. <code>Switch(config)#interface {fastethernet gigabitethernet} slot/port</code>	Select the interface to configure.
2.	Select the encapsulation. <code>Switch(config-if)#switchport trunk encapsulation {isl dot1q negotiate}</code>	(Optional) If switchport trunk mode is configured, this command must be used with either the <i>isl</i> or the <i>dot1q</i> argument. Negotiate is the default.
3.	Configure the interface as a Layer 2 trunk. <code>Switch(config-if)#switchport mode {dynamic {auto desirable} trunk}</code>	The switchport mode of the interface determines if the link will perform trunking.
4.	Specify the native VLAN. <code>Switch(config-if)#switchport trunk native vlan vlan_number</code>	The default is VLAN1.
5.	Configure the allowable VLANs for this trunk. <code>Switch(config-if)#switchport trunk allowed vlan {add except all remove} vlan_num1 [,vlan_num[,vlan_num[,...]]</code>	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 trunk
Switch(config-if)#switchport nonegotiate
Switch(config-if)#switchport trunk encapsulation dot1q
Switch(config-if)#switchport trunk native vlan 99
Switch(config-if)#no shutdown
Switch(config-if)#end
Switch#exit
```

Caution Ensure that the native VLAN for an 802.1Q trunk is the same on both ends of the trunk link. If there is native VLAN mismatch, traffic cannot be transmitted correctly on the trunk.

Verify the 802.1Q Configuration

This subtopic describes how to verify an 802.1Q link.

Verifying the 802.1Q Configuration

Cisco.com

```
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: Off  
Access Mode VLAN: 1 (default)  
Trunking Native Mode VLAN: 99 (trunk_only)  
Trunking VLANs Enabled: 1,5,11,1002-1005  
Pruning VLANs Enabled: 2-1001  
  
. . .
```

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—2-19

Example: Configure and Display Port Information for an 802.1Q Dynamic Trunk Link

This subtopic describes how to verify an 802.1Q dynamic trunk link.

Verifying a 802.1Q Dynamic Trunk Link

Cisco.com

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

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: 99 (trunk_only)
Trunking VLANs Enabled: 1,5,11,1002-1005
Pruning VLANs Enabled: 2-1001
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2--2-20

Example: Displaying Trunk Information for 802.1Q Trunking

This example shows how to display trunk information for 802.1Q trunk links. 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	802.1q	trunking	99

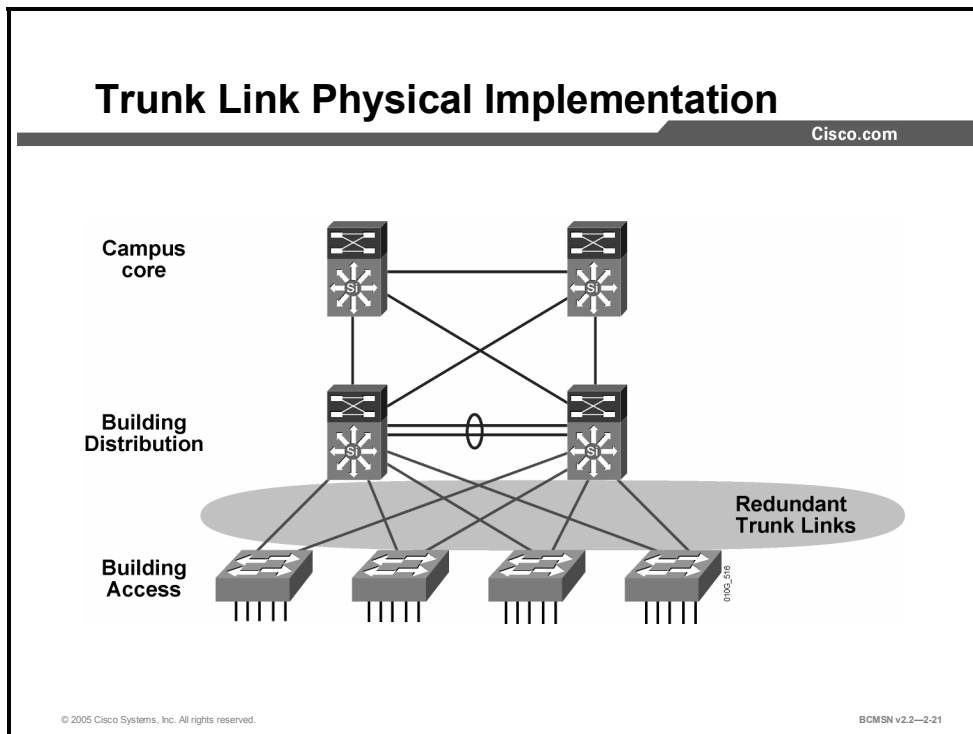
Port	Vlans allowed on trunk
Fa5/8	1,5,11,1002-1005

Port	Vlans allowed and active in management domain
Fa5/8	1,5,1002-1005

Port	Vlans in spanning tree forwarding state and not pruned
Fa5/8	1,5,1002-1005

Using Trunking Protocols in the Campus Infrastructure Module

This topic identifies how trunks are used in the Campus Infrastructure module.



Careful design and consideration should be taken when implementing VLAN trunks because they can add to overall network congestion and can also present security challenges. These are general best practices for trunk implementation in the Campus Infrastructure module.

- VLAN1 should be removed from the trunks to ensure that no user data propagates among the switches on VLAN1. Although each Catalyst switch requires VLAN1 on the actual switch and it is not possible to remove, it is possible to remove VLAN1 from trunk links.
- Limit the trunk link to only the intended VLANs required for Layer 2 access and connectivity. This improves bandwidth use by restricting unwanted VLAN traffic from the link. Explicitly permitting or denying VLANs to a specific trunk link creates a simple, deterministic Layer 2 switched domain with fewer variables to complicate troubleshooting. This also facilitates correct operation of VLAN interfaces.
- DTP should not be required. Trunk links, encapsulation types, and access ports should be statically configured across specific links according to the network design and requirements.
- Cisco is now migrating to 802.1Q as the recommended trunking protocol because of the interoperability and compatibility between the Layer 2 and Layer 3 prioritization methods. The IEEE 802.1Q/p standard provides architectural advantages over ISL; these include widely accepted QoS classification and marking standards, and the ability to carry frames that are not tagged with a VID.

Resolving Trunk Link Problems

This topic describes basic trunk link problems.

Resolving Trunk Link Issues

Cisco.com

Check the following:

- **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).**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-2.22

If a problem exists with a trunk link, or if a trunk link cannot be established, check the following:

- Verify that the interface mode configured on both ends of the link is identical or valid for negotiated links. The interface mode should be trunk, dynamic, or nonegotiate.
- Verify that the trunk encapsulation type configured on both ends of the link is valid and compatible.
- For 802.1Q links, verify that the native VLAN is the same on both ends of the trunk.

Note If the trunk appears to be configured correctly on both ends or if the configuration has changed and no trunk forms, shut down the interfaces on both ends of the trunk, check for a match in the configuration, and then bring the interfaces up to see if the trunk forms correctly.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- Trunk links carry traffic from multiple VLANs.
- Trunking protocols provide a means for frames to carry VLAN ID.
- ISL and 802.1Q protocols have specific benefits and features.
- ISL is Cisco proprietary and encapsulates the Layer 2 frames.
- 802.1Q is an IEEE standard for trunking.
- The 802.1Q native VLAN can forward frames with no VLAN ID.
- VLAN numbers have specific ranges and purposes.
- DTP can automatically negotiate a trunk link.
- Various commands are used to configure and verify ISL and 802.1Q trunk links.
- Specific guidelines should be followed when implementing trunks in the Campus Infrastructure module.
- Avoid known problems with VLAN trunk misconfiguration.

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—2-23

Propagating VLAN Information with VTP

Overview

When VLANs span multiple switches, a protocol is needed to accurately manage VLAN information at each switch. This protocol is referred to as VLAN Trunk Protocol (VTP) and is used to ensure that all switches in a given group, or VTP domain, have the same information about the VLANs present in that domain. This lesson will examine VTP and how it allows each switch to participate in the domain. The VTP mode determines if and when updates are sent by a switch.

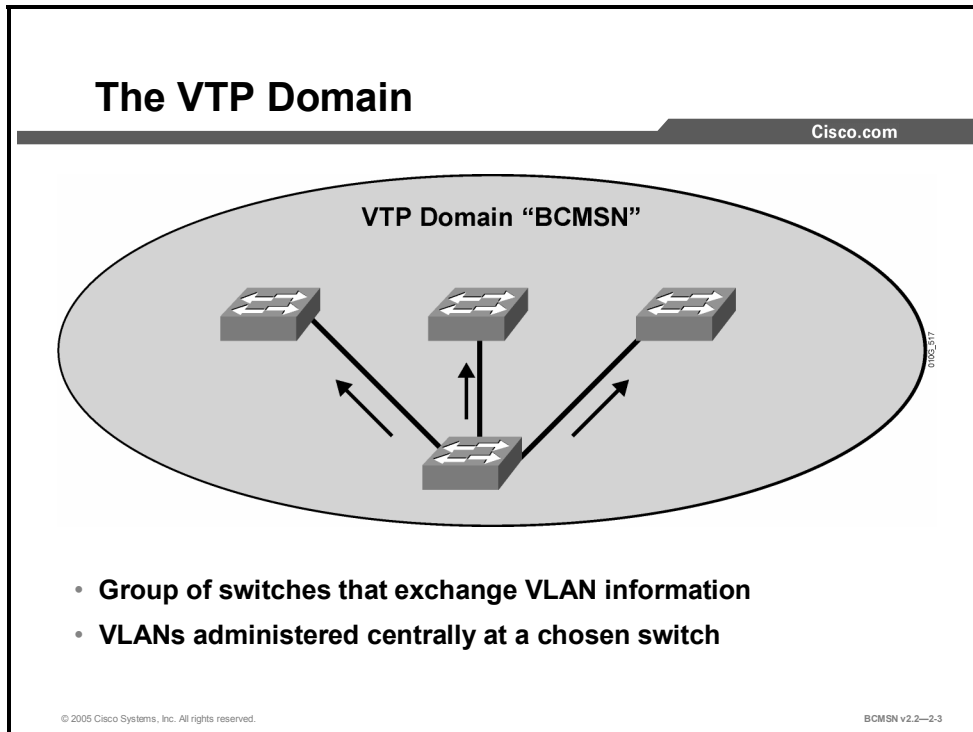
Objectives

Upon completing this lesson, you will be able to configure VTP modes and domains so that VLAN configuration information is sent between switches as intended. This ability includes being able to meet these objectives:

- Define a VTP domain
- Identify key features of the VTP protocol
- List the three VTP modes and their associated attributes
- Describe how the VTP protocol distributes and synchronizes VLAN information
- Configure and verify operations in a VTP management domain
- List the best practices when configuring VTP in the Enterprise Composite Network Model
- List the recommended steps for adding a new switch to an existing VTP domain

What Is a VTP Domain?

This topic identifies how switches can be logically grouped within a given network in order to maintain consistent VLAN information among them.



In an enterprise network with many interconnected switches, maintaining a consistent list of VLANs across those switches can be administratively cumbersome and potentially error prone. The VTP is designed to automate this administrative task.

Switches that share common VLAN information are organized into logical groups called VTP management domains. The VLAN information within a VTP domain is propagated through trunk links and is updated via the VTP, allowing all switches within a particular domain to maintain identical VLAN databases.

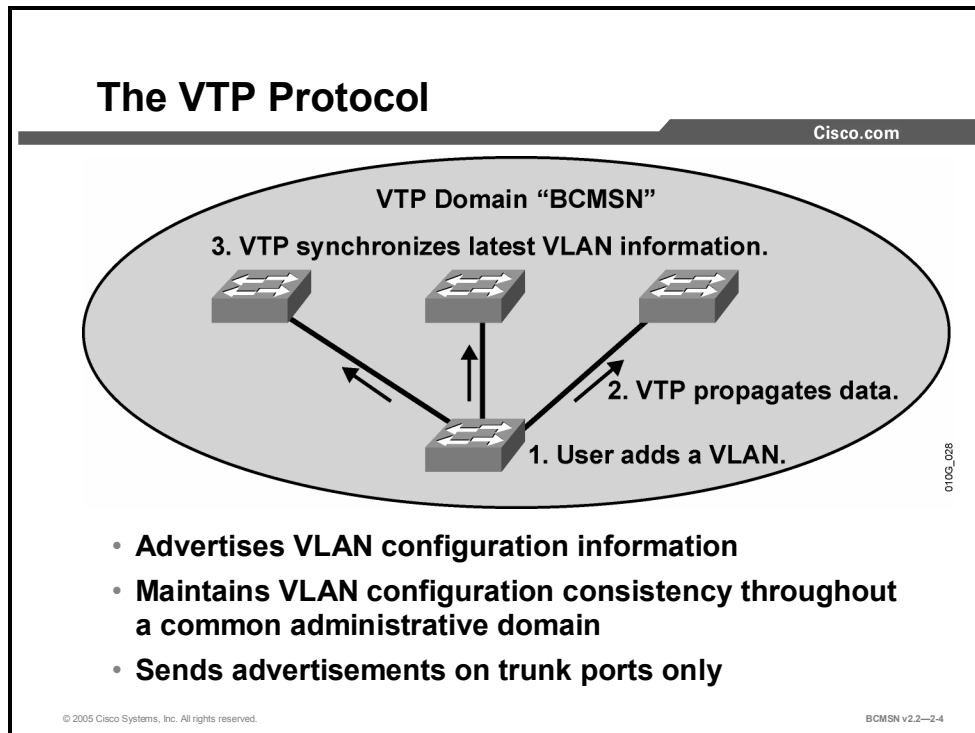
Only “global” VLAN information regarding VLAN number, name, and description is exchanged. Information on how ports are assigned to VLANs on a given switch is kept local to the switch and is not part of a VTP advertisement.

These are the attributes of a VTP domain:

- A switch may be in only one VTP domain.
- A VTP domain may be as small as only one switch.
- VTP updates will be exchanged only with other switches in the same domain.
- The way VLAN information is exchanged between switches in the same domain depends upon the VTP mode of the switch.
- 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 a management domain is configured.

What Is the VTP Protocol?

This topic describes the VTP protocol.



Switches in a single VTP domain exchange VTP updates to distribute and synchronize VLAN information. VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the additions, deletions, and name changes of VLANs on all switches in a VTP domain.

VTP runs over trunk links allowing interconnected switches to exchange Layer 2 frames, synchronizing a single list of configured VLANs. This reduces the manual configuration required at each switch; VLANs can be created on one switch and then propagated to others.

These are the attributes of VTP:

- VTP is a Cisco proprietary protocol.
- VTP will advertise VLANs 1–1005.
- VTP updates are exchanged only across trunk links.
- Each switch operates in a given VTP mode that determines how VTP updates are sent from and received by that switch.
- There are three VTP versions that support different features.

VTP in the Campus Infrastructure Module

There are some benefits to using VTP within the guidelines of the Campus Infrastructure module.

- The VTP domain is restricted to building switch blocks.

- VTP keeps VLAN information consistent between Building Distribution layer and Building Access layer switches.
- VLAN configuration errors or failures will be confined to a switch block.
- Knowledge of all VLANs does not need to exist on all switches within the Campus Infrastructure module. Use of VTP is optional, and in high-availability environments it is best practice to set all switches to ignore VTP updates.

Caution VLANs deleted on one switch may be deleted on all switches in the VTP domain, and thus all ports removed from that VLAN. Delete VLANs with caution on a switch participating in a VTP domain with other switches.

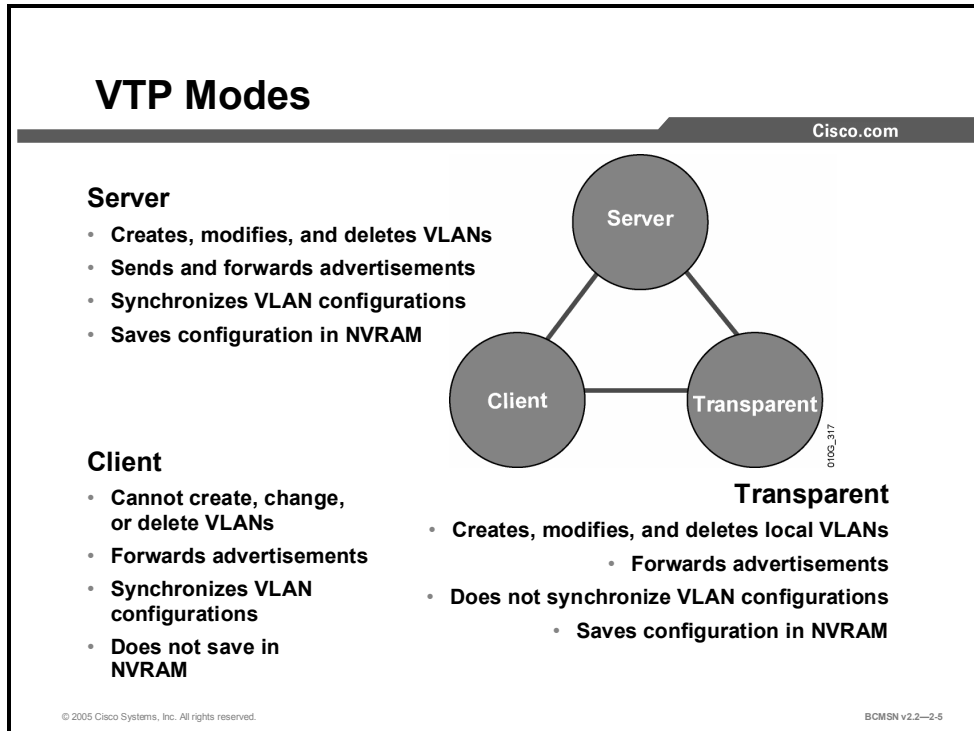
References

For additional information, refer to this resource:

http://cisco.com/en/US/tech/tk389/tk689/technologies_tech_note09186a0080094c52.shtml

VTP Modes

This topic describes the three modes in which VTP operates.



On each switch, VTP can be configured to operate in one of three modes: server, client, or transparent. The default VTP mode is server. The mode will determine if VLANs can be created on the switch and how the switch will participate in sending and receiving VTP advertisements. The number of VLANs that can be configured on a switch will vary by mode.

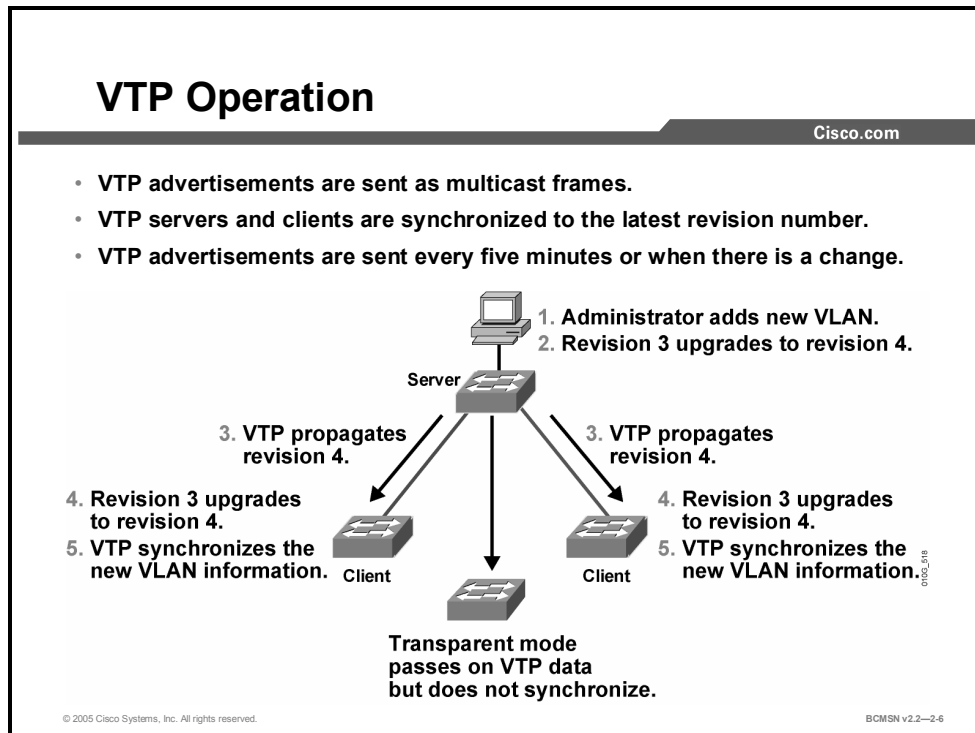
The table describes the features of the VTP client, server, and transparent modes.

VTP Mode	Features
Server	<ul style="list-style-type: none"> ■ Creates, modifies, and deletes VLANs at the command-line interface (CLI) ■ Generates VTP advertisements and forwards advertisements from other switches ■ May update its own VLAN database with information received from other servers in the management domain ■ Saves VLAN configuration information in “vlan.dat” file in Flash memory
Client	<ul style="list-style-type: none"> ■ Cannot create, modify, or delete VLANs at the CLI. ■ Forwards VTP advertisements received. ■ Synchronizes its own VLAN database with latest information received from VTP servers in the management domain. ■ VLAN information is in RAM only, not stored in NVRAM or Flash memory; must be repopulated from a VTP server if switch is powered off
Transparent	<ul style="list-style-type: none"> ■ Creates, modifies, and deletes VLANs for the VLAN database on the local switch only ■ Does not generate VTP advertisements ■ Does not update its VLAN database with information received from VTP servers in

VTP Mode	Features
	<p>the same management domain</p> <ul style="list-style-type: none"><li data-bbox="480 205 1409 233">■ Forwards VTP advertisements received from VTP servers in the same VTP domain<li data-bbox="480 254 1036 281">■ Always has a configuration revision number of 0<li data-bbox="480 302 927 329">■ Saves VLAN configuration in NVRAM

Describing VTP Operation

This topic describes VTP operation.



Switches within a VTP management domain synchronize their VLAN databases by sending and receiving VTP advertisements over trunk links. VTP advertisements are flooded throughout a management domain by switches running in specific modes of operation. Advertisements are sent every 5 minutes or whenever there is a change in VLAN configuration. VTP advertisements are transmitted over VLAN1, using a Layer 2 multicast frame. VLAN advertisements are not propagated from a switch until a management domain name is specified or learned.

VLAN synchronization over VTP follows this general order:

- Step 1** Configure the VTP domain, VTP mode, and VTP password (optional) on each switch. This proactively determines which switches will send updates.
- Step 2** Switches running VTP server mode then send VTP updates across trunk links.
- Step 3** A device that receives a VTP advertisement will check that the VTP management domain name and password in the advertisement match those configured in the local switch.
- Step 4** If a match is found, a switch further inspects the VTP update to see the configuration revision number.
- Step 5** If the configuration revision number of the message is greater than the number currently in use and the switch is running in VTP server or client mode, the switch overwrites its current VLAN information with that in the received update.
- Step 6** The switch may also request more information.

Configuration Revision Number

One of the most critical components of VTP is the configuration revision number. When initially configured, the VTP configuration revision number is set to 0. Each time a VTP server modifies its VLAN information, it increments the VTP configuration revision number by one. It then sends out a VTP advertisement referencing the new configuration revision number. If the configuration revision number being advertised is higher than the number stored on other switches in the VTP domain, they will overwrite their VLAN configurations with the new information.

Caution Given this overwrite process, if the VTP server sending the advertisement deletes all VLANs and sends an advertisement with a higher revision number, the receiving devices in the VTP domain will delete those VLANs as well.

VTP Advertisement Types

Three types of VTP advertisements are exchanged between switches.

Summary advertisements: An update sent by VTP servers every 300 seconds or when a VLAN database change occurs. Among other things, this advertisement lists the management domain, VTP version, domain name, configuration revision number, time stamp, and number of subset advertisements. If the advertisement results from a VLAN database change, one or more subset advertisements will follow.

Subset advertisements: An update that follows a summary advertisement resulting from a change in the VLAN database. A subset advertisement cites the specific change that was made to a specific VLAN entry in the VLAN database. One subset advertisement will be sent for each VLAN ID (VID) that encountered a change.

Advertisement requests from clients: An update sent by a switch requesting information in order to update its VLAN database. If a client hears a VTP summary advertisement with a configuration revision number higher than its own, the switch may send an advertisement request. A switch operating in VTP server mode then responds with summary and subset advertisements.

Note VTP advertisements are associated with VLAN database information only, not with VLAN information configured on specific switch ports. Likewise, on a receiving switch, the receipt of new VLAN information does not change the VLAN associations of trunk or access ports on that switch.

VTP Versions

Currently, Catalyst switches run VTP versions 1, 2, or 3. Version 2 is the most prevalent and provides these features:

- Forwarding of VTP updates from transparent mode switches without checking the version number
- Consistency checks on new VTP and VLAN configuration parameters
- Support for Token Ring switches
- Propagation of VTP updates that have an unrecognized type, length, or value

VTP version 3 is available on some switches now using the Catalyst software operating system version.

References

For additional information, refer to this resource:

http://www.cisco.com/en/US/partner/products/hw/switches/ps708/products_configuration_guide_chapter09186a008019f048.html#wp1017196

VTP Configuration Commands

This topic describes the commands used to configure VTP.

VTP Configuration Commands

Cisco.com

Configuring VTP

- **vtp domain**
- **vtp mode**
- **vtp password**

Verifying VTP

- **show vtp status**
- **show vtp counters**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-2.7

The **VTP** configuration command is used to configure VTP characteristics for a switch. All switches in the same VTP domain will share the same VTP domain name and VTP password, if one is configured. It is a good idea to set the VTP mode to “client” if switches are being added to an existing switched network.

The **Show VTP** commands are used to verify the current VTP parameter values.

VTP Commands

Command	Description
Switch(config)# vtp domain <i>domain_name</i>	Sets the VTP domain name. Enter an ASCII string from 1 to 32 characters to identify the VTP administrative domain for the switch. The domain name is case sensitive.
Switch(config)# vtp password <i>password</i>	The 16-byte secret value used in Message Digest 5 (MD5) digest calculation to be sent in VTP advertisements and to validate received VTP advertisements. The password can be an ASCII string from 1 to 32 characters and is case sensitive.
Switch(config)# vtp v2-mode	Enables VTP version 2 in the administrative domain.
Switch(config)# vtp mode client	Places the switch in VTP client mode. VLANs cannot be configured on a switch configured in this mode. When a VTP client starts up, it does not send VTP advertisements until it receives advertisements to initialize its VLAN database.

Command	Description
Switch(config)# vtp mode server	Places the switch in VTP server mode. The switch is enabled for VTP and sends advertisements. VLANs can be configured on the switch. The switch can recover all the VLAN information in the current VTP database from nonvolatile storage after reboot.
Switch(config)# vtp mode transparent	Places the switch in transparent mode. It cannot affect VLAN configurations on other devices in the network. The switch receives VTP advertisements and forwards them on all trunk ports except the one on which the advertisement was received.
Switch# show vtp status	Verifies the VTP name, mode, revision number, and other information.
Switch# Show vtp counters	Indicates if VTP updates are being sent and received by the switch.

How to Configure a VTP Management Domain

This topic describes configuration of a VTP management domain.

VTP Considerations for Configuration

Cisco.com

Configure the following to avoid dynamic learning:

- **VTP domain name**
- **VTP mode (server mode is the default)**
- **VTP password**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-2-8

Default VTP configuration values depend on the switch model and the software version. The default values for the Catalyst 2900, 4000, and 6000 Series switches are as follows:

- **VTP domain name:** None
- **VTP mode:** Server
- **VTP password:** None
- **VTP trap:** Disabled (Simple Network Management Protocol [SNMP] traps communicating VTP status)

The VTP domain name can be specified or learned from VTP updates seen from other switches. By default, the domain name is not set.

A password can be set for the VTP management domain. The password must be the same for all switches in the domain in order for the VLAN database to be synchronized among switches.

Configuring VTP on a Switch

This subtopic lists the steps used to configure VTP.

Configuring and Verifying VTP

Cisco.com

`Switch#show vlan brief`

- Displays a list of current VLANs

`Switch(config)#vtp mode`

- Sets the VTP mode to server, client, or transparent

`Switch(config)#vtp domain domain_name`

- Sets the VTP domain name

`Switch(config)#vtp password password_string`

- Sets the VTP password

`Switch# show vtp status`

- Displays the current settings for VTP

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—2-9

The steps for configuring VTP will vary per design and switch mode, but the general steps for configuring a switch are as follows:

- Step 1** Establish a design specifying what switches will be server, client, or transparent, and what the boundaries are for the VTP domain.
- Step 2** Verify the current VLAN information on any switch that will be configured as server.
- Step 3** Configure the switch mode.
- Step 4** Specify the VTP domain name.
- Step 5** Specify the version number, if other than default.
- Step 6** Specify the VTP password (optional).
- Step 7** Verify the configuration.
- Step 8** Verify that updates are being sent from or received by the switch as intended.

To configure a switch to become part of a VTP domain, follow these steps from privileged EXEC mode.

Step	Action	Notes
1	Display list of VLANs. Switch# show vlan brief	Determine if the list of VLANs displays before configuration. If this switch is about to be configured as a server, this list will overwrite VLAN information on client switches.
2.	Enter global configuration mode. Switch# configure terminal	
3.	Configure the VTP mode. Switch(config)# vtp mode	Enter server, client, or transparent. To revert to the default (server), enter no vtp mode .
4.	Configure the domain name. Switch(config)# vtp domain <i>domain_name</i>	Defines the VTP domain name, which can be up to 32 characters long. Domain name case must match other switches in the domain for updates to occur properly.
5.	Enable VTP version 2. Switch(config)# vtp v2-mode	To revert to VTP version 1, enter vtp v1-mode .
6.	Specify a VTP password. Switch(config)# vtp password <i>password_string</i>	Sets a password for the VTP domain, which can be from 1–34 characters and is case sensitive. The password can only be set on switches operating in server or client mode. Use no vtp password to clear the password.
7.	Exit global configuration mode. Switch(config)# exit	
8	Display list of VLANs. Switch# show vlan brief	Determine if the list of VLANs shown is as anticipated, given the mode of the switch.

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

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#
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—2-10

Use the **show vtp status** command to verify the VTP configuration.

When initially configuring switches in a VTP domain, pay close attention to the configuration revision number. Check to see that it increments only when changes are made at intended VTP servers.

In the example above, “Configuration last modified by 0.0.0.0” specifies the IP address of the switch that updated the VLAN database of this switch.

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).

Verifying the VTP Configuration (Cont.)

Cisco.com

```
Switch#show vtp counters
```

```
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
-----          -----          -----          -----
Fa5/8          43071          42766          5
```

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—2-11

VTP Counters

Use the **show vtp counters** command to display statistics about VTP operation.

Output from this command verifies if VTP updates are being sent and are being received by the switch, and records the number of updates that have been seen.

Common Problems with VTP Configuration

The following unexpected results can occur after VTP configuration:

VTP Configuration Problems

Problem	Possible Causes
Updates not received as expected	<ul style="list-style-type: none">■ VTP domain name and password must match on a given switch in order to receive updates from a VTP server.■ VTP version must be compatible with other switches in the domain.■ Ensure that there is at least one server in the domain.■ Check that a trunk link exists to VTP server.

Problem	Possible Causes
Missing VLANs	<ul style="list-style-type: none"> <li data-bbox="764 170 1435 247">■ Upon initial configuration, the VTP server may have had a partial VLAN database, and it overwrote the existing, more complete, database on the switch. <li data-bbox="764 268 1435 373">■ VLANs were deleted individually at the VTP server, and those deletions will be propagated in the domain. (To avoid this, ensure that any switch becoming a VTP server has a complete VLAN list.) <li data-bbox="764 394 1458 499">■ Not all Cisco switches support the same extended-range VLANs (those numbered higher than 1005). This information is not learned or propagated through VTP, so it may vary in a switched network.
Too many VLANs	<ul style="list-style-type: none"> <li data-bbox="764 527 1442 575">■ The VTP server has a VLAN list that is more complete than the list needed by other switches in the domain.

Best Practices: Configuring Switches in a VTP Domain

This topic describes best practices for implementing VTP.

Implementing VTP in the Enterprise Composite Network Model

Cisco.com

- **Plan VTP domain boundaries.**
- **Have only one or two VTP servers.**
- **Manually configure VTP domain name.**
- **Configure a VTP password.**
- **Consider order of device configuration:**
 - **New domain**
 - **Existing domain cleanup**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—2-12

Below are some general best practices with regard to configuring VTP in the Enterprise Composite Network Model:

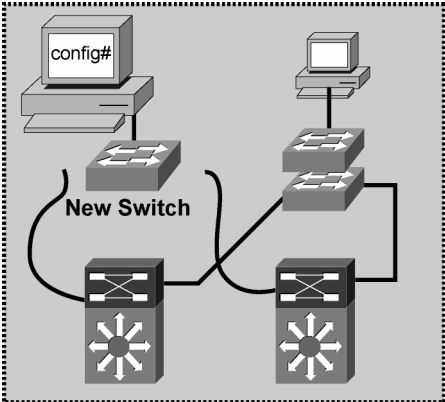
- Plan boundaries for the VTP domain. Not all switches in the network need information on all VLANs in the network. In the enterprise composite model, the VTP domain should be restricted to redundant distribution switches and the access switches they serve.
- Have only one or two switches specifically configured as VTP servers and the remainder as clients.
- Manually configure VTP on all switches installed in the network so the mode can be specified and the default mode of server on all switches can be overwritten.
- Configure a password so that no switch can join the VTP domain with a domain name only (which can be derived dynamically).
- When setting up a new domain, configure VTP client switches first so they participate passively; then configure servers to update client devices.
- In an existing domain, if performing VTP cleanup, configure passwords on servers first. Clients may need to maintain current VLAN information until the server contains a complete VLAN database. Once the VLAN database on the server is verified as complete, client passwords can be configured to be the same as the servers. Clients will then accept updates from the server.

How to Add a New Switch to an Existing VLAN

This topic describes adding a new switch to an existing VLAN.

Adding a Switch to an Existing VLAN

Cisco.com



Configure VTP before adding a new switch to a network.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-2-13

The configuration revision number is the only criterion used when determining if a switch should keep its existing VLAN database or overwrite it with the VTP update sent by another switch in the same domain with the same password. Therefore, when a switch is added to a network it is important that it does not inject spurious information into the domain.

Caution If a new switch was at one time attached to another network, it may contain a vlan.dat file, and its configuration revision number may be higher than that of other VTP servers in the VTP domain to which it is being added. If no VTP domain was explicitly configured on the switch, when it is connected to the network it is able to derive the VTP domain name from any VTP update it sees. If there is no password on the domain and if the new switch is in server mode (default), its VLAN information can overwrite the VLAN database on other switches in the VTP domain.

The “How to Add a Switch to an Existing Network” table explains the steps for adding a new switch to the network. It is critical to VLAN stability to add a switch in this manner.

How to Add a Switch to an Existing Network

Step	Action	Notes
1.	Ensure that there is no connectivity between the new switch and the network, and power the switch on.	This keeps updates from this switch from overwriting the VLAN databases of other switches in the domain before the switch is configured properly.
2.	Change the switch VTP mode to transparent.	This will set the configuration revision number to 0 and ensure that no updates are received if the switch is connected to the network out of sequence in subsequent steps.
3.	Delete vlan.dat.	This will remove any VLAN information from the switch.
4.	Change the VTP domain name to something unconventional and the mode to client.	This will keep the switch from dynamically learning the domain name upon reload.
5.	Reload or power cycle the switch.	This will remove all VLAN information from RAM.
6.	Verify the switch VTP and VLAN configuration.	Use Show commands to verify that the switch is in client mode and that it has the anticipated VLANs configured (or none configured).
7.	Configure the switch with a valid VTP domain name and password.	Switch is configured to participate in network as intended.
8.	Connect the switch to the network.	Switch should receive accurate, current VLAN information.
9.	Verify VLAN database.	Use Show commands to verify that the switch has received VLAN information as intended.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Switches in a VTP domain share VLAN information.**
- **The VTP protocol advertises VLAN information.**
- **VTP operates in one of three modes.**
- **VTP uses a specific process to distribute and synchronize VLAN information between switches.**
- **Various commands are used to configure and verify VTP operation on a switch.**
- **Best practices should be followed when configuring VTP in the Enterprise Composite Network Model.**
- **Specific steps should be followed when adding a new switch to an existing VTP domain.**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—2-14

Module Summary

This topic summarizes the key points discussed in this module.

Module Summary

Cisco.com

- **Access VLAN configuration isolates VLAN traffic to appropriate ports within the Building Access layer of the Campus Infrastructure module.**
- **Multiple VLANs can be carried on a single access to distribution link by configuring VLAN trunking.**
- **VLAN configuration information can be sent between switches using VTP.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—2-1

This module examined the function of VLANs and how they are implemented in a switched campus network. Depending on its configuration as an access or trunk port, each switch port can be associated with one or many VLANs. The ISL and 802.1Q protocols are used to establish trunk links carrying traffic for multiple VLANs. Trunk links between switches can also carry VTP information, which allows VLAN names and descriptions contained in a VLAN database to be shared between switches.

References

For additional information, refer to these resources:

- Cisco Systems, Inc., *VLAN Trunking Protocols: InterSwitch Link and IEEE 802.1Q Frame Format*:
http://www.cisco.com/en/US/tech/tk389/tk390/technologies_tech_note09186a0080094665.shtml
- Cisco Systems, Inc., *Catalyst 3550 Multilayer Switch Software Configuration Guide, 12.1(13)EA1: Configuring VTP*:
http://cisco.com/en/US/products/hw/switches/ps646/products_configuration_guide_chapter09186a008014f376.html

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which two statements identify network benefits provided by VLANs? (Choose two.) (Source: Implementing VLANs)
- 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.
- Q2) Which **show** command is used to verify a VLAN port configuration? (Source: Implementing VLANs)
- A) **running-config interface fa0/3**
 - B) **VLAN id 3**
 - C) **switchport access VLAN 200**
 - D) **interfaces**
- Q3) Which two features belong to the 802.1Q trunking protocol? (Choose two.) (Source: Supporting Multiple VLANs on a Single Trunk)
- A) It encapsulates Ethernet frames.
 - B) It alters the existing Ethernet frame.
 - C) It supports native VLANs.
 - D) It does not support native VLANs.
- Q4) Which two guidelines apply when adding a new switch to an existing VTP management domain? (Choose two.) (Source: Automating the Propagation of VLAN Information)
- A) Set the switch to client mode before connecting it to the network.
 - B) Set the switch to server mode before connecting it to the network.
 - C) Change the VTP domain information while the switch is in server mode.
 - D) Change the VTP domain information while the switch is in client mode.

Module Self-Check Answer Key

Q1) B, E

Q2) A

Q3) B, C

Q4) A, D

Implementing Spanning Tree

Overview

This module introduces the fundamentals of Spanning Tree Protocol (STP) operation in a switched network. The root bridge will be explained as well as how it and its backup are elected. Features for enhancing the performance of STP will be covered—namely, PortFast, UplinkFast, and BackboneFast. You will discover how EtherChannel is configured and how it interoperates with STP. The module also provides guidelines on improving STP resiliency when network faults occur.

Module Objectives

Upon completing this module, you will be able to implement the STP to accelerate network traffic convergence in Layer 2. This ability includes being able to meet these objectives:

- Define STP and explain how it can be used to accelerate network traffic convergence in Layer 2
- Use appropriate commands to configure and verify primary and backup root bridges
- Configure PortFast on Layer 2 access ports
- Configure BPDU guard, BPDU filtering, and root guard to prevent rogue Layer 2 devices from playing a key role in STP operation
- Use appropriate commands to configure and verify an UplinkFast group to create an alternative forwarding path if a link fails
- Configure BackboneFast on a port to create an alternative path through an intermediate switch
- Configure EtherChannel to bundle ports in order to support aggregated bandwidth between two switches

Defining the Spanning Tree Protocol

Overview

In a campus network where there are redundant links between switches, Spanning Tree Protocol (STP) manages which links will provide an active Layer 2 path, which ones will be inactive, and which ones will provide redundancy in the case of active path failure. This lesson will examine the general components and operation of STP in a switched network.

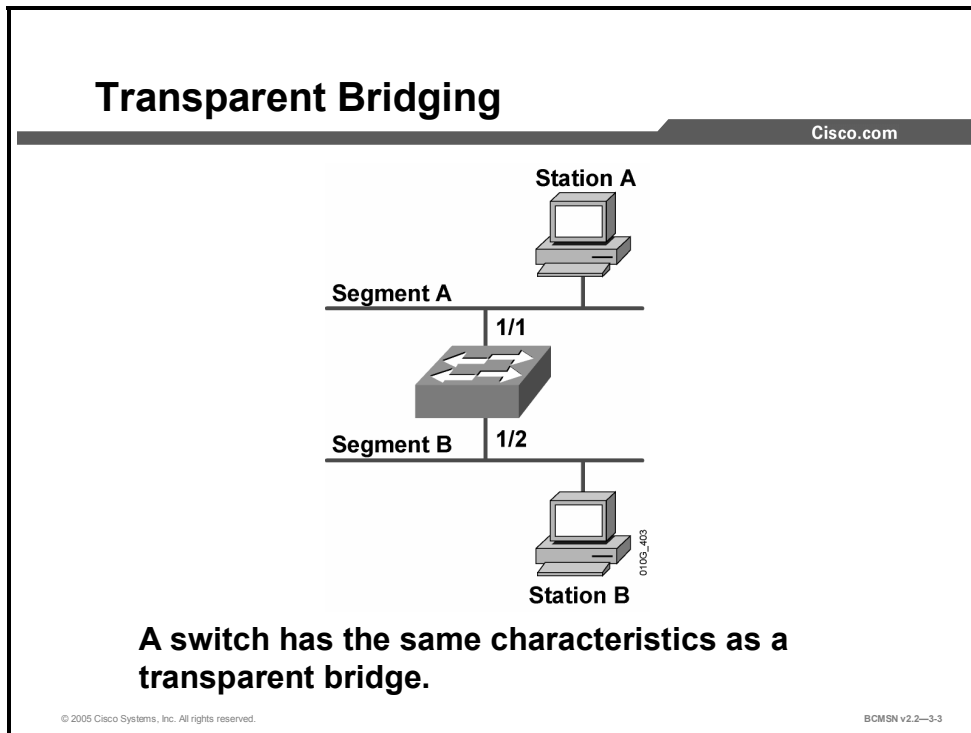
Objectives

Upon completing this lesson, you will be able to understand the purpose and functionality of STP. This ability includes being able to meet these objectives:

- Describe transparent bridges
- Identify traffic loops
- Explain how to prevent loops in a Layer 2 network
- Define the purpose of the 802.1D STP
- Define the purpose of a root bridge
- Identify the root selection process
- List the 802.1D port roles
- Identify how associations are formed between the root bridge and nonroot bridges
- Identify the process of determining the active STP topology

Transparent Bridges

This topic identifies the features that apply to transparent bridging.



Because switches have replaced bridges as the network device for implementing transparent bridging in modern networks, the basic functionality of a switch is identical to that of a transparent bridge on a per-VLAN basis. To understand STP, it is important first to look at the behavior of a transparent bridge without spanning tree.

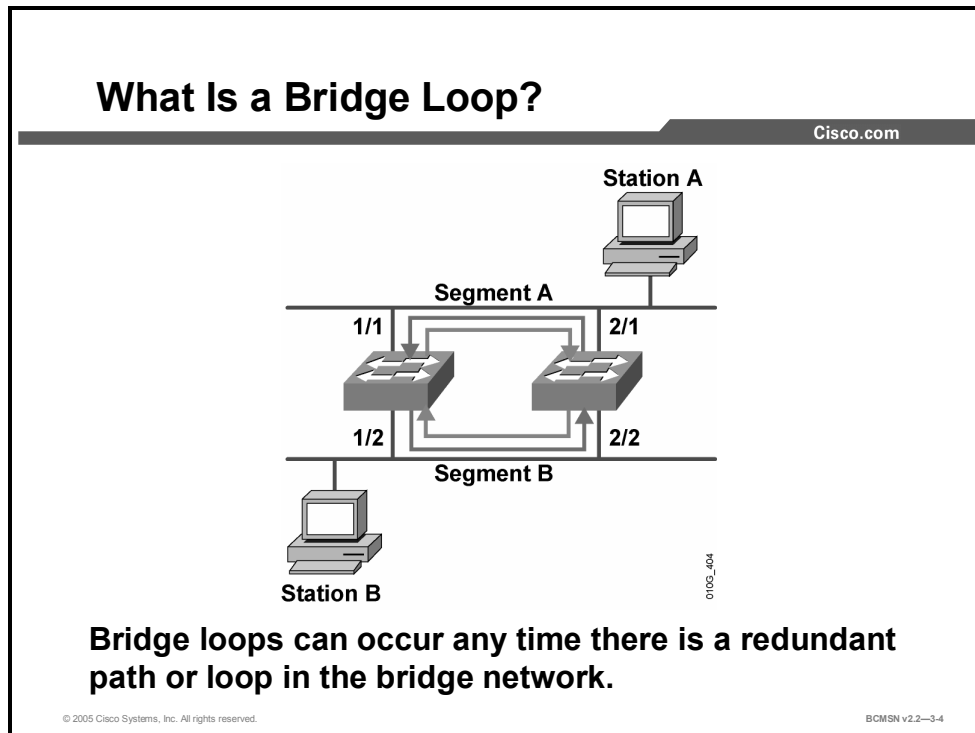
By definition, a transparent bridge has these characteristics:

- It must not modify the frames that are forwarded.
- It learns addresses by “listening” on a port for the source address of a device. If a source MAC address is read in frames coming in a specific port, the bridge assumes that frames destined for that MAC address can be sent out of that port. The bridge then builds a table that records what source addresses are seen on what port. A bridge is always listening and learning MAC addresses in this manner.
- It must forward all broadcasts out all ports, except for the port that initially received the broadcast.
- If a destination address is unknown to the bridge, it forwards the frame out all ports except for the port that initially received the frame. This is known as unicast flooding.

Transparent bridging, by definition, must be transparent to the devices on the network. End stations require no configuration. The existence of the bridging protocol operation is not directly visible to them—hence, the term transparent bridging.

Identifying Traffic Loops

This topic identifies the behavior of flooded unicast frames in a bridge loop.



A bridge loop is observed when a frame that is forwarded circulates cyclically and redundantly; this occurs where there is no mechanism to manage the redundant Layer 2 paths.

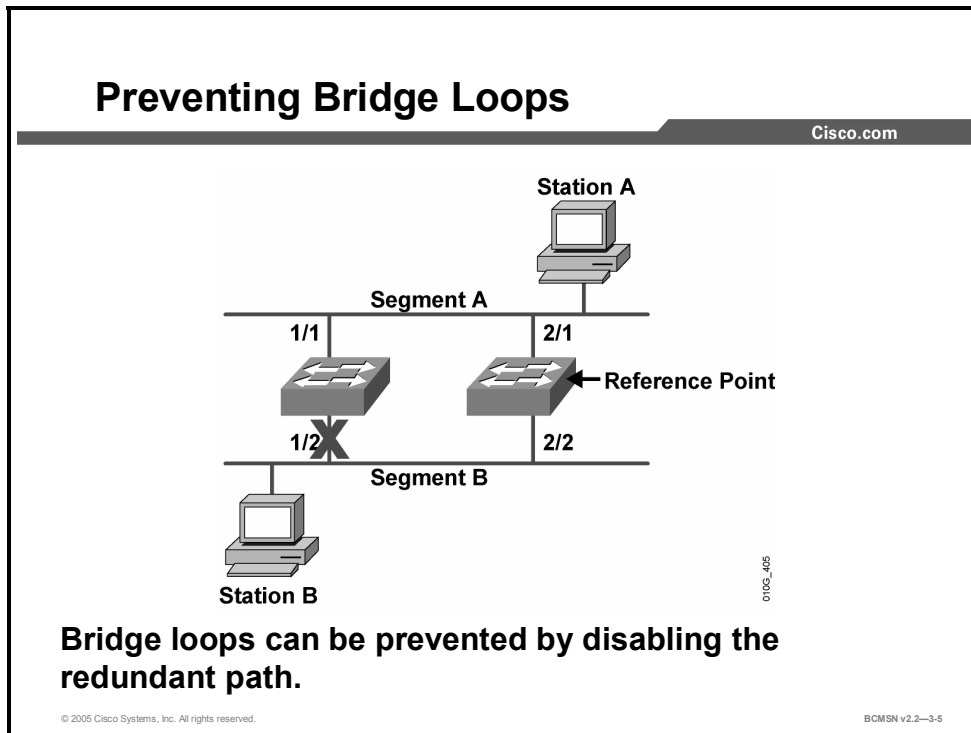
Example: Flooded Unicast Frames and Bridge Loops

Station A has two potential paths to station B by way of the two intermediate bridges. The following describes what happens if station A sends frames to station B, if there were no provisions enabled to deal with redundant paths.

- Station A transmits the frame destined for station B onto segment A. Both bridges on segment A pick up the frame on their bridge ports 1/1 and 2/1. Both bridges populate their respective MAC tables, indicating that station A resides on segment A, on bridge ports 1/1 and 2/1.
- Both bridges forward the frame to segment B. Station B receives the frame, and both bridges also see the same frame, with station A's MAC address in the Source Address (SA) field, coming from the *other* bridge. The bridges 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 bridges because it will be received on the same bridge ports that it considers the destination of station A.
- 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 bridge to forward the frame out all bridge ports except the bridge port upon which the frame was received. The broadcast frame would also be forwarded to the originating bridge, which would again forward the same broadcast out all bridge ports. This broadcast would continue until the loop is shut down or until the bridge could no longer handle the load.

Preventing Loops on a Layer 2 Network

This topic describes the characteristics of a loop free network and how loops can be avoided.



A loop free network is one in which no Layer 2 loops exist; therefore, the network cannot create Layer 2 broadcast storms or flooded unicast storms. A loop free network can be achieved manually by shutting down or disconnecting all redundant links between bridges. However, this leaves no redundancy in the network and requires manual intervention in the event of a link failure.

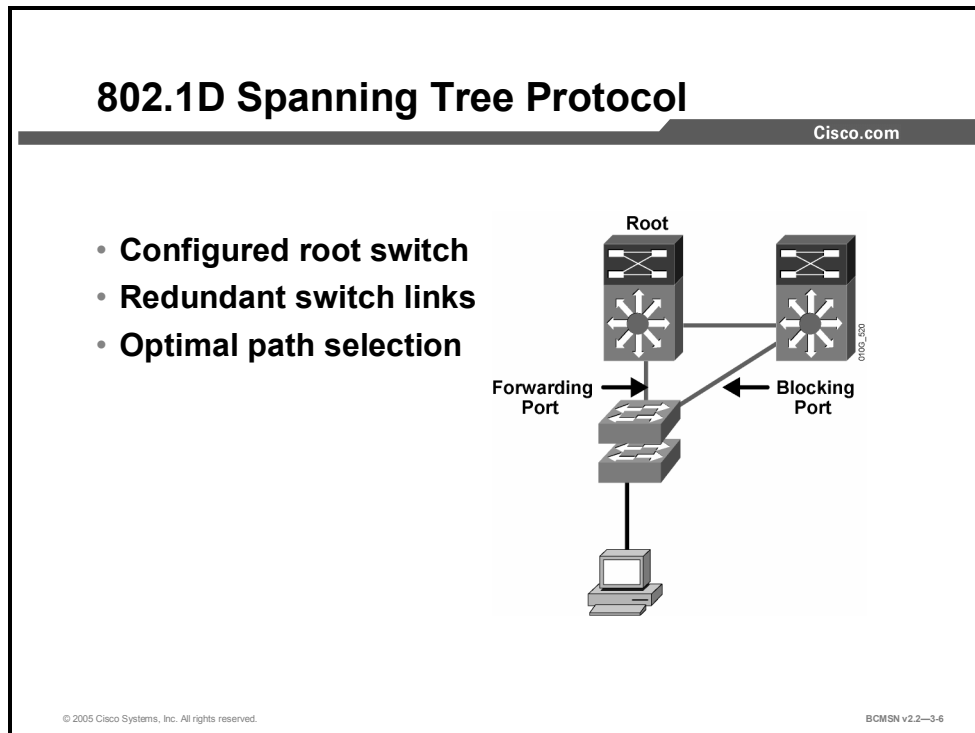
STP resolves this problem. Where there are alternative links to a destination on a switch, only one link will be used to forward data unless there is a failure on that link. The switch ports associated with alternative paths remain aware of the topology of the network and can be enabled if a failure occurs on a primary link. In the case of primary link failure, the switch will begin forwarding frames over an alternative link.

The spanning tree algorithm (STA) runs on each switch to activate or block redundant links. To find the redundant links, the STA chooses a reference point in the network and determines if there are redundant paths to that reference point. If the STA finds a redundant path, it chooses which path will forward frames and which redundant path or paths will be blocked. This effectively severs the redundant links within the network until they are needed when the primary link toward the reference point fails.

Spanning tree standards often refers to a “bridge,” but it is likely that all the devices exchanging spanning tree information will be Layer 2 switches.

802.1D Spanning Tree Protocol

This topic provides an overview of the Spanning Tree Protocol, what it does and how it works.



The 802.1D STP provides a mechanism for switches to reconfigure the paths over which they forward frames, making possible a loop free path when there are redundant switch paths through the network. This is accomplished by forwarding traffic over specific ports and by disabling other ports to prevent frames from being sent repeatedly or in a loop. STP prevents loops by using the following mechanisms:

- STP is implemented through the exchange of bridge protocol data unit (BPDU) messages between adjacent switches.
- A single "root bridge" is elected to serve as the reference point from which a loop free topology is built for all switches exchanging BPDUs.
- Each switch determines a "root port" that provides the best path to the root bridge.
- On a link between two nonroot switch ports, a port on one switch will become a designated port, and the port on the other switch will be in a blocking state, not forwarding frames. This effectively breaks any loop. Typically, the designated port will be on the switch with the best path to the root bridge.
- Any port state change on any switch is considered a network topology change (for example, if a port goes up or down and the spanning tree algorithm must be run on all switches to adapt to the new topology).

Spanning Tree Communication

This subtopic identifies the information contained in a BPDU that is used to send spanning tree information between switches.

Bridge Protocol Data Unit

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

BDUs provide for the exchange of information between switches.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-7

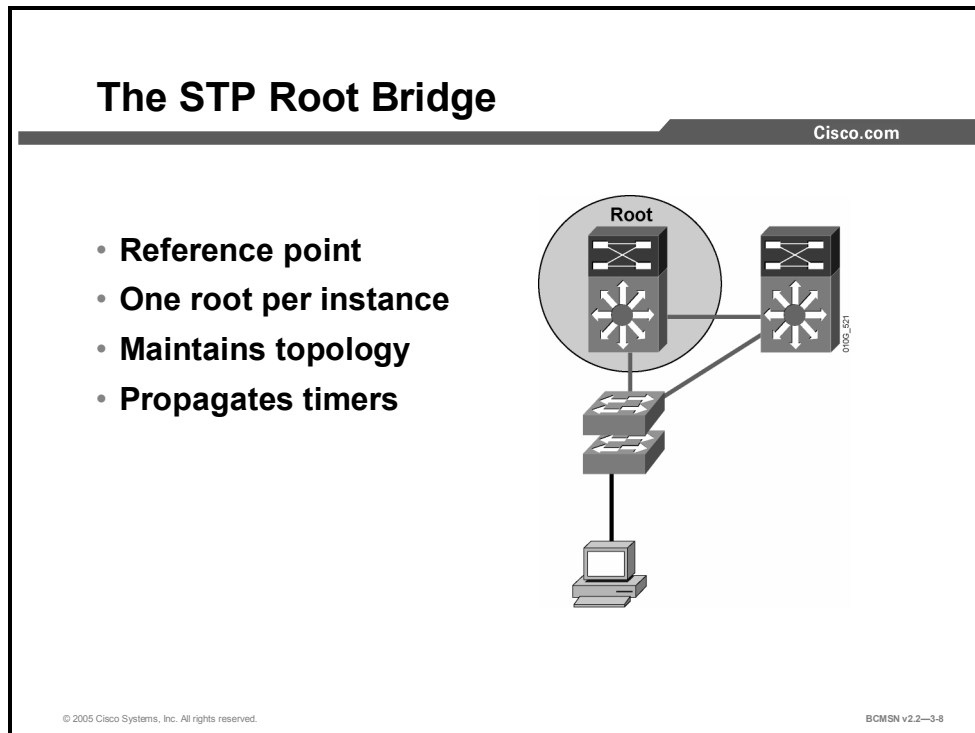
A BPDU is the unit of information sent by the STP between switches to establish and maintain a loop free Layer 2 topology over optimal network paths. The information provided in a BPDU includes the following:

- **Root ID:** The lowest bridge ID (BID) in the topology
- **Cost of path:** Cost of all links from the transmitting switch to the root bridge
- **BID:** of the transmitting switch
- **Port ID:** Transmitting switch port ID
- **STP timer values:** Max age, hello time, forward delay

The switch compares the BDUs received on all ports to its own values to determine what role the receiving switch and its ports will play in the STP topology.

What Is a Root Bridge?

This topic defines and describes a root bridge.



STP uses the concepts of root bridge, root ports, and designated ports to establish a loop free path through the network. The first step in creating the loop free spanning tree is to elect a root bridge. The root bridge is the reference point that all switches use to establish forwarding paths that will avoid loops in the Layer 2 network.

When a topology change occurs as a result of switch link state changes, the root will send messages throughout the tree regarding the topology change. This allows the content addressable memory (CAM) tables to adjust and to provide for a new path that may be used toward end host devices.

Timer information is also sent by the root bridge to nonroot bridges, informing them of the intervals to use as the ports transition through the spanning tree port states.

The root bridge maintains the stability of the forwarding paths between all switches for a single STP instance. A spanning tree instance is when all switches exchanging BPDUs and participating in spanning tree negotiation are associated with a single root. If this is done for all VLANs, it is called a Common Spanning Tree (CST) instance. There is also a Per VLAN Spanning Tree (PVST) implementation that provides one instance, and therefore one root bridge, for each VLAN.

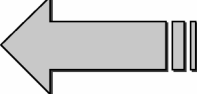
BPDUs Fields Associated with Root Bridge Selection

This subtopic describes the criteria used to determine which device will be elected as the root.

Root Bridge Selection Criteria

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.**

0106_150

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-9

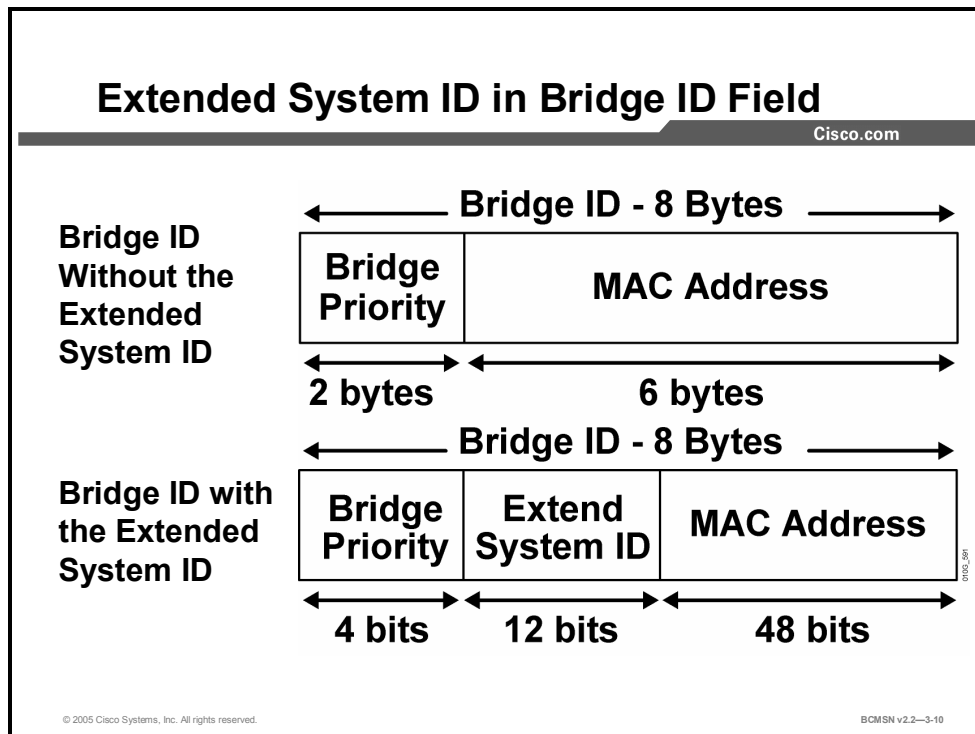
The BID and root ID are each 8-byte fields carried in a BPDU. These values are used to complete the root bridge election process. A switch identifies the root bridge by evaluating the Root ID field in the BPDUs it receives. The unique BID of the root bridge will be carried in the Root ID field of the BPDUs sent by each switch in the tree.

When a switch first boots and begins sending BPDUs, it has no knowledge of a root ID, so it will populate the Root ID field of outbound BPDUs with its own BID.

The switch with the lowest numerical BID will assume the role of root bridge for that spanning tree instance. Upon receipt of BPDUs with a lower BID than its own, a switch will place the lowest value seen in all BPDUs into the Root ID field information of its outbound BPDUs.

Bridge ID Field in the BPDU

This subtopic describes the bridge ID field content.



Spanning tree operation requires that each switch have a unique BID. In the original 802.1D standard, the bridge ID was composed of the Priority Field and the MAC address of the switch, and all VLANs were represented by a Common Spanning Tree. Because PVST requires that a separate instance of spanning tree run for each VLAN, the Bridge ID field is required to carry VLAN ID (VID) information. This is accomplished by reusing a portion of the Priority field as the extended system ID to carry a VID.

To accommodate the extended system ID, the original 802.1D 16-bit Bridge Priority field is split into two fields, resulting in these components in the BID:

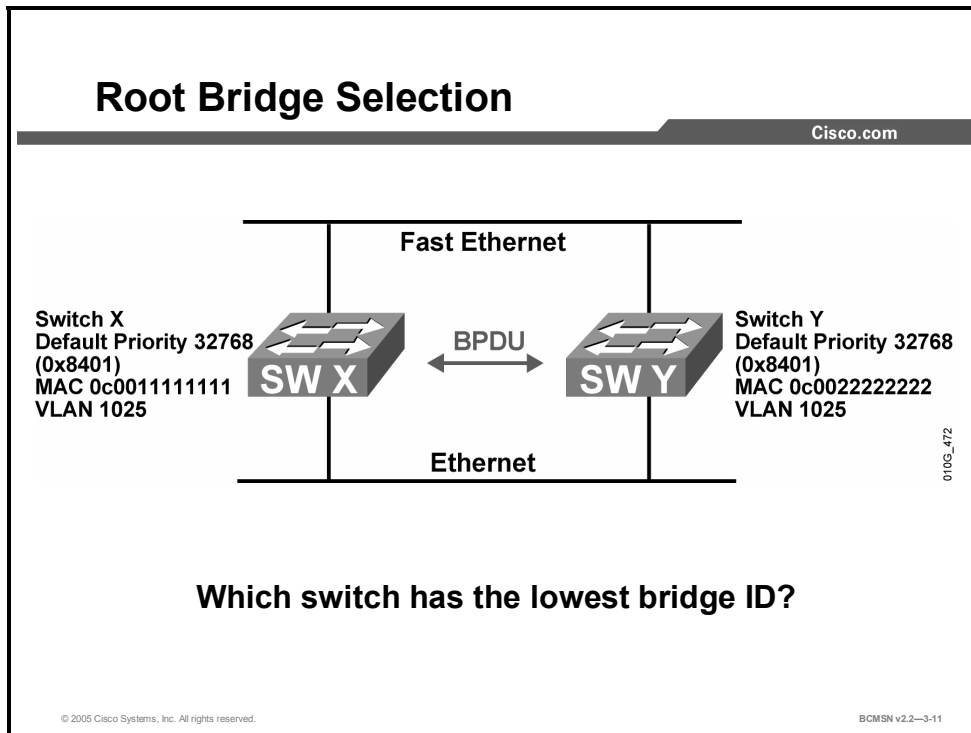
- **Bridge Priority:** A 4-bit field still used to carry bridge priority. Because of the limited bit count, priority is now conveyed in discreet values in increments of 4096 rather than discreet values in increments of 1, as they would be with the full 16-bit field available. The default priority, in accordance with IEEE 802.1D, is 32,768, which is the mid-range value.
- **Extended System ID:** A 12-bit field carrying, in this case, the VID for PVST.
- **MAC Address:** A 6-byte field with the MAC address of a single switch.

By virtue of the MAC address, a BID is always unique. When the priority and extended system ID are appended to the switch MAC address, each VLAN on the switch can be represented by a unique BID.

If no priority has been configured, every switch will have the same default priority and the election of the root for each VLAN will be based on the MAC address. This is a fairly random means of selecting the ideal root bridge and, for this reason, it is advisable to assign a lower priority to the switch that should serve as root bridge.

Identifying the Root Selection Process

This topic describes the process by which a root bridge is elected.



BPDUs are exchanged between switches, and the analysis of the bridge ID and root ID information from those BPDUs determines which bridge is selected as the root bridge.

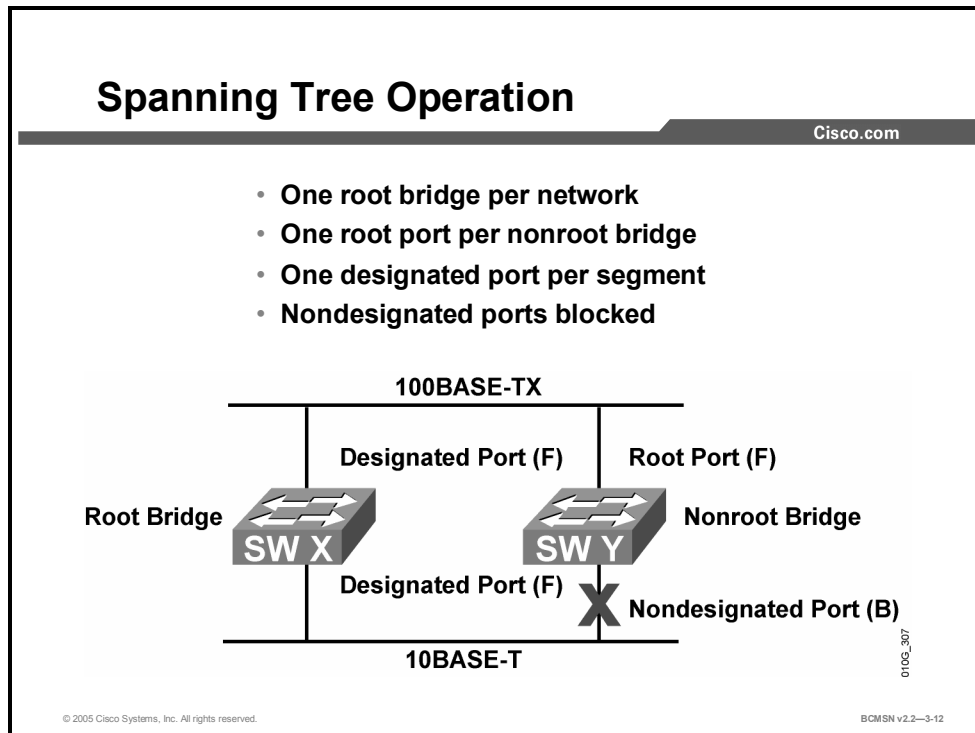
In the example shown, both switches have the same priority for the same VLAN. The switch with the lowest MAC address will, therefore, be elected root bridge. In the example, switch X is the root bridge, with a bridge ID of 0x8401:0c0011111111.

Election of a Root bridge

Step	Action
1.	Upon startup, each switch transmits BPDUs out all enabled interfaces on a per-VLAN basis. At startup, each switch sets the root ID equal to its own bridge ID. During this time, the switch ports are not used to forward standard data frames.
2.	As the BPDU goes out through the network, each switch compares the root BPDU it sent out to the one it received. The exact BPDU fields and how they are compared are outlined in the next topic.
3.	If the received root ID is superior, the switch will propagate it; otherwise, it will continue to send its own BID as the root BID in transmitted BPDUs.
4.	On the root bridge, all ports are designated ports in a forwarding state.
5.	Nonroot bridges must determine an optimal path to the root.

802.1D Port Roles

This topic lists the possible port roles associated with 802.1D.



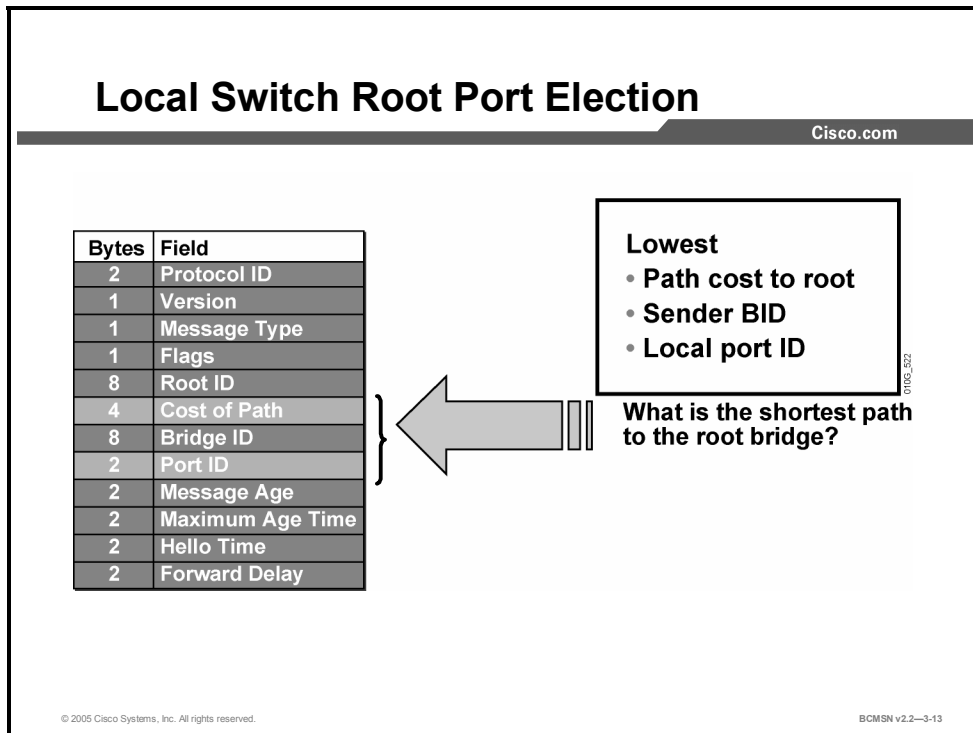
On a nonroot bridge, as spanning tree receives BPDUs on various ports, it determines the roles that each port will fill in the topology. There are four 802.1D port roles.

- **Root port:** This port exists on nonroot bridges and is the switch port with the best path to the root bridge. Root ports forward traffic toward the root bridge, and the source MAC address of frames received on the root port is capable of populating the MAC table. Only one root port is allowed per bridge.
- **Designated port:** This port exists on root and nonroot bridges. For root bridges, all switch ports are designated ports. For nonroot bridges, a designated port is the switch port that will receive and forward frames toward the root bridge as needed. Only one designated port is allowed per segment. 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. Designated ports are capable of populating the MAC table.
- **Nondesignated port:** This is a switch port not forwarding (blocking) data frames and not populating the MAC address table with the SAs of frames seen on that segment.
- **Disabled port:** This is a switch port that is shut down.

By examining the switch port roles on a switch, spanning tree can determine the most desirable forwarding path for data frames.

Forming an Association with the Root Bridge

This topic identifies methods by which switch ports determine their role in STP.



Nonroot bridges place various ports in their proper roles by listening to BPDUs as they come in on all ports. Receiving BPDUs on multiple ports indicates a redundant path to the root bridge.

The switch looks at these components in the BPDU to determine which switch ports will forward data and which switch ports will block data:

- Lowest path cost
- Lowest sender BID
- Lowest local port ID

The switch looks at the path cost first to determine which port is receiving the lowest cost path. The path is calculated on the basis of 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 and sender BID are equal, as with parallel links between two switches, the switch goes to the port ID as a “tiebreaker.” The port with the lowest port ID forwards data frames, and all other ports continue to block data frames.

Path Cost

This subtopic 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 Gbp	4	1
100 Mbps	19	10
10 Mbps	100	100

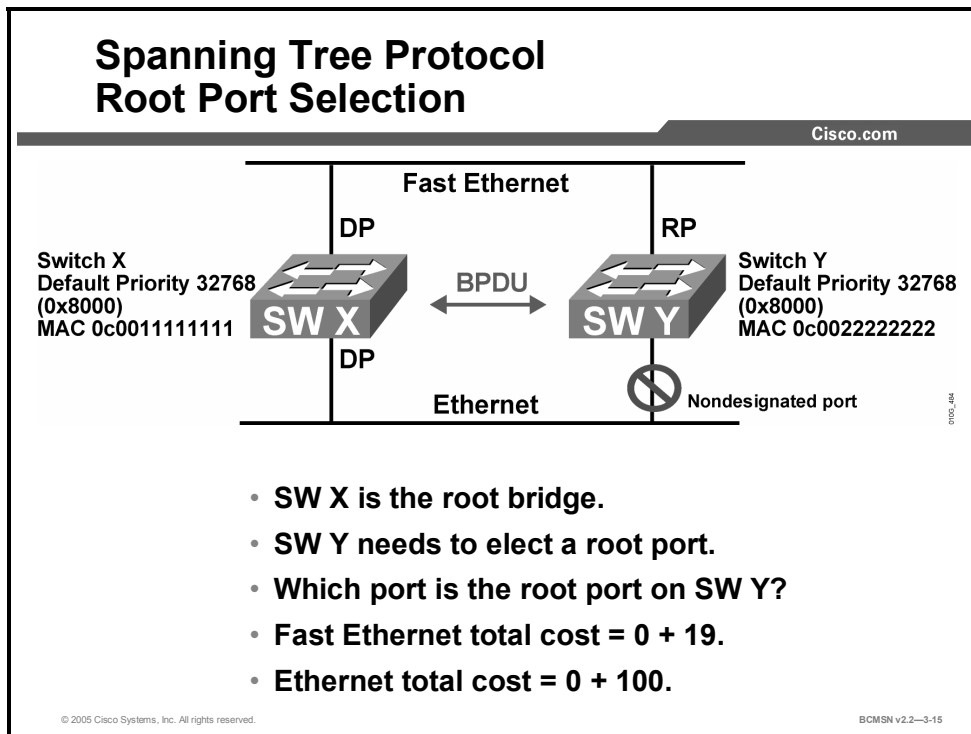
© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2--3-14

The spanning tree path cost is a value advertised in the BPDU by each bridge. This is a value representing the cumulative cost of all the links from the root bridge to the switch sending the BPDU. Path cost value is used by the receiving switch to determine the best path to the root bridge. The lowest cost is considered to be the best path.

Port cost values per link are shown in the table under the “Revised IEEE Spec” heading, with the lower values being associated with higher bandwidth and therefore being the more desirable paths. This new specification uses a nonlinear scale with port cost values as shown. In the previous IEEE specification, the cost value was calculated based on Gigabit Ethernet being the maximum Ethernet bandwidth, with an associated value of 1, from which all other values were derived in a linear manner.

Selecting the Root Port

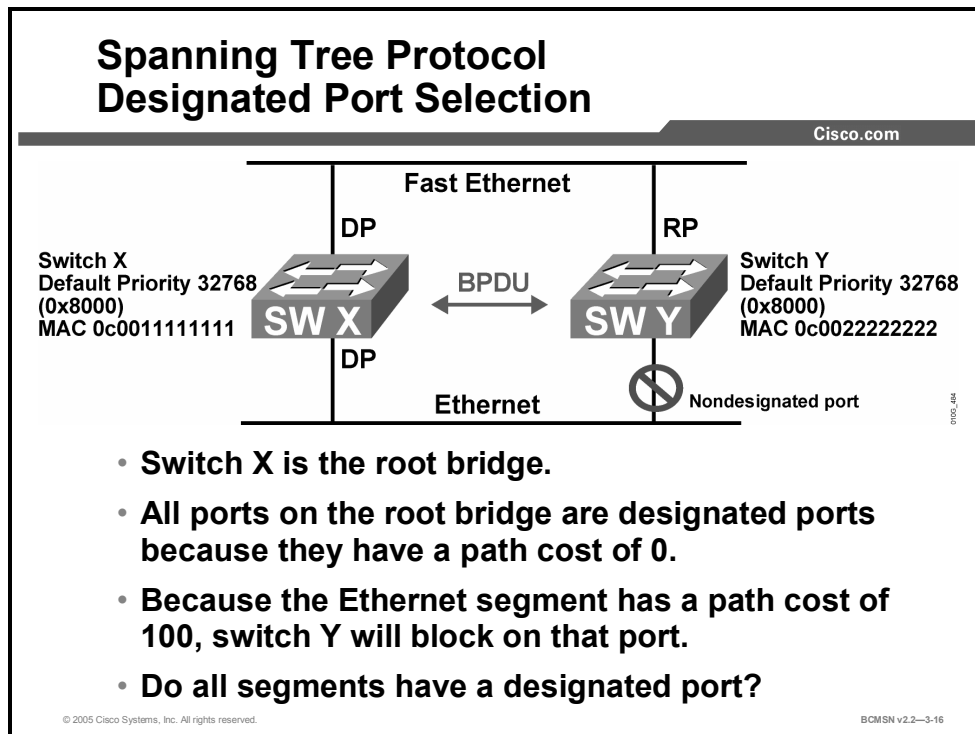
This topic identifies how a switch port is selected as the root port.



Switch Y receives a BPDU from the root bridge (switch X) on its switch port on the Fast Ethernet segment and another BPDU 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 bridge and is elected the root port for switch Y.

Selecting the Designated Port

This subtopic identifies the features that apply to designated switch ports.

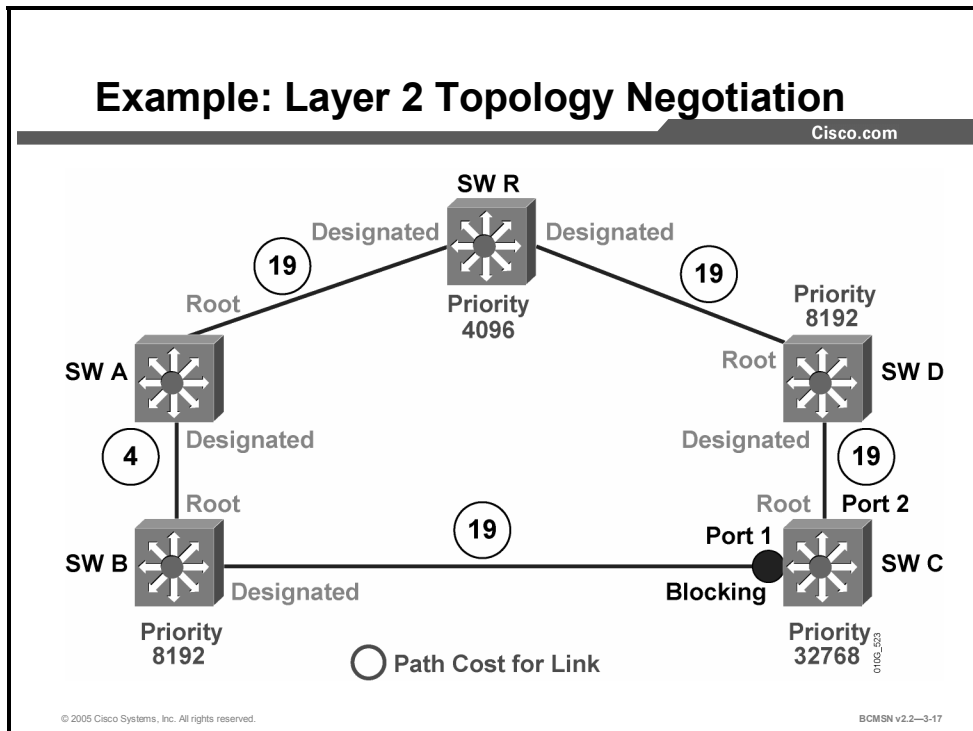


STP selects one designated port per segment to forward traffic. Other switch ports on the segment become nondesignated ports and continue blocking. The nondesignated ports receive BPDUs but do not forward data traffic to prevent loops. The switch port on the segment with the lowest path cost to the root bridge 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.

Example: Determining the Active Topology

This topic identifies the features that apply to BPDU messaging.



Here is a scenario with switches running STP and exchanging information as shown in the figure. From this information, exchange will yield the following final results:

- The election of a root bridge as a Layer 2 topology point of reference
- The determination of the best path to the root bridge from each 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 some switch links to a blocked state
- 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, using Topology Change Notifications, the switched network must reconfigure the activity topology if a link failure occurs.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Transparent bridges require no client configuration.**
- **A bridge loop may occur when there are redundant paths between switches.**
- **A loop free network eliminates redundant paths between switches.**
- **The 802.1D protocol establishes a loop free network.**
- **The root bridge is a reference point for STP.**
- **The root bridge is selected by examining BPDUs.**
- **Each STP port will host a specific port role.**
- **Each nonroot bridge will form an association with the root bridge.**
- **STP will form an active forwarding topology between switches.**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—3-18

Maintaining and Configuring STP

Overview

The root bridge plays an important role in spanning tree operation. This device is determined by an election process. Switches allow configuration of priority parameters that will force a specific device to win election as root or backup root bridge. Finally, the lesson discusses how topology changes may alter the path that data frames take as they traverse a switched network.

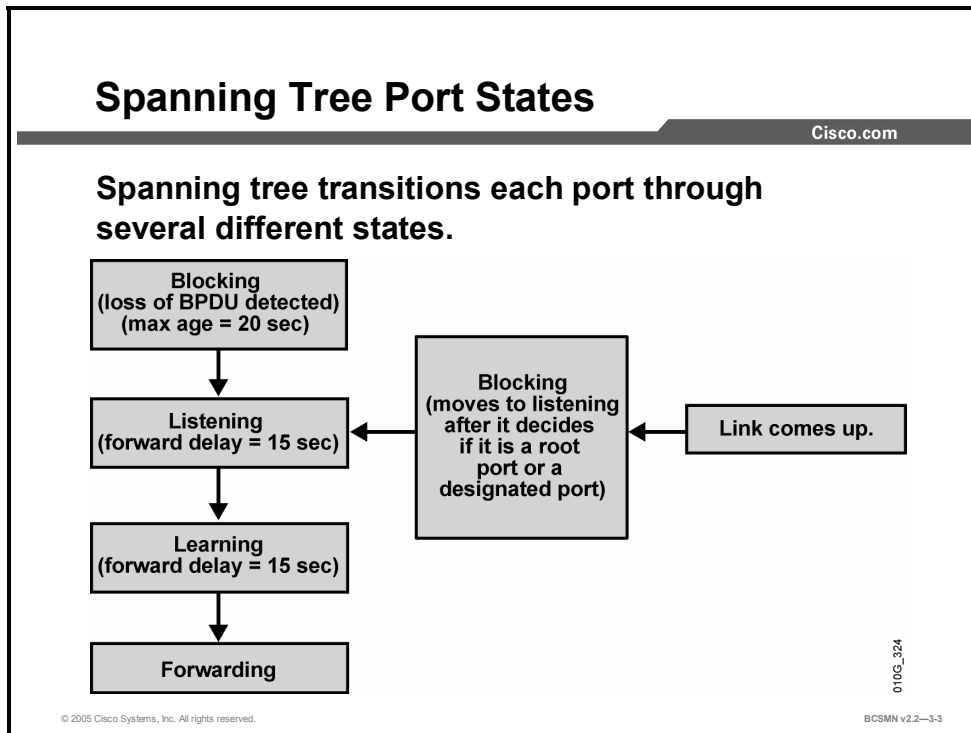
Objectives

Upon completing this lesson, you will be able to configure STP to force root bridge selection and paths toward the root bridge. This ability includes being able to meet these objectives:

- Describe STP port states
- Enumerate the spanning tree process that results from a Layer 2 topology change
- Identify the purpose of a backup root bridge
- Configure and verify the root and backup root bridge
- Use the appropriate commands to configure primary and backup root bridges
- Compare CST to PVST

Identifying Spanning Tree Port States and Timers

This topic identifies the features and functions that apply to each port state.



If a Layer 2 port were to transition directly from a Spanning Tree Protocol (STP) blocking to the forwarding state, the port could temporarily create a data loop if the switch was not aware of all topology information at the time. For this reason, STP introduces delay timers and port states that ensure that all STP information has been propagated to that switch before it will transition a port to a forwarding state.

Each Layer 2 port on a switch running STP exists in one of these five port states:

- **Blocking:** In this state, the Layer 2 port is a nondesignated port and does not participate in frame forwarding. The port receives bridge protocol data units (BPDUs) to determine the location and root ID of the root switch and what port roles (root, designated, or nondesignated) each switch port should assume in the final active STP topology.
- **Listening:** In this state, spanning tree has determined that the port can 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 adjacent switches that the switch port is preparing to participate in the active topology.
- **Learning:** In this state, the Layer 2 port prepares to participate in frame forwarding and begins to populate the content addressable memory (CAM) table.
- **Forwarding:** In this state, the Layer 2 port is considered part of the active topology; it forwards frames and also sends and receives BPDUs.
- **Disabled:** In this state, the Layer 2 port does not participate in spanning tree and does not forward frames.

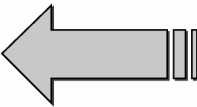
Spanning Tree Timers

This subtopic 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.

01002_7/1

- **Timers are used to prevent bridging loops.**
- **Timers determine how long it will take STP to converge after a failure.**

© 2005 Cisco Systems, Inc. All rights reserved.BCSMN v2.2-3-4

The amount of time that a port stays in the various port states is dependent upon the BPDU timers. Only the switch in the role of root bridge may send information through the tree to adjust the timers. The following three timers affect STP performance and state changes:

- **hello time:** 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 between 1 and 10 seconds.
- **forward delay:** The forward delay is the time spent in the listening and learning states. This is by default equal to 15 seconds for each state but can be tuned to between 4 and 30 seconds.
- **max age:** The max age timer controls the maximum length of time a switch port saves configuration BPDU information. This is 20 seconds by default but can be tuned to between 6 and 40 seconds.

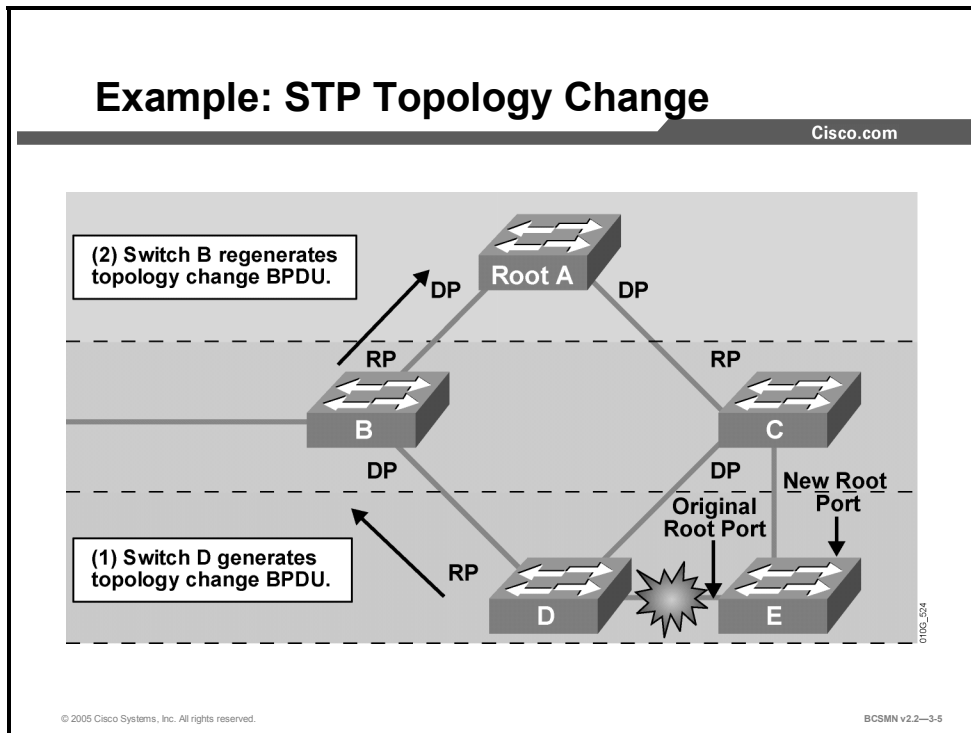
When STP is enabled, every switch port in the network goes through the blocking state and the transitory states of listening and learning at power up. The ports then stabilize in the forwarding or blocking state. During a topology change, a port temporarily implements the listening and learning states for a specified period called the “forward delay interval.”

These values allow adequate time for convergence in a network with switches seven Layer 2-hops from one another. This is referred to as an STP diameter, and a maximum of seven is permitted. The STP diameter value can be adjusted to a lower value that will automatically adjust the forward delay and max age timers proportionally for the new diameter.

Caution Best practice suggests not to alter the spanning tree timers individually but to adjust them indirectly by configuring the diameter to reflect the actual network topology.

Identifying Topology Changes

This topic identifies the steps involved in a topology change.

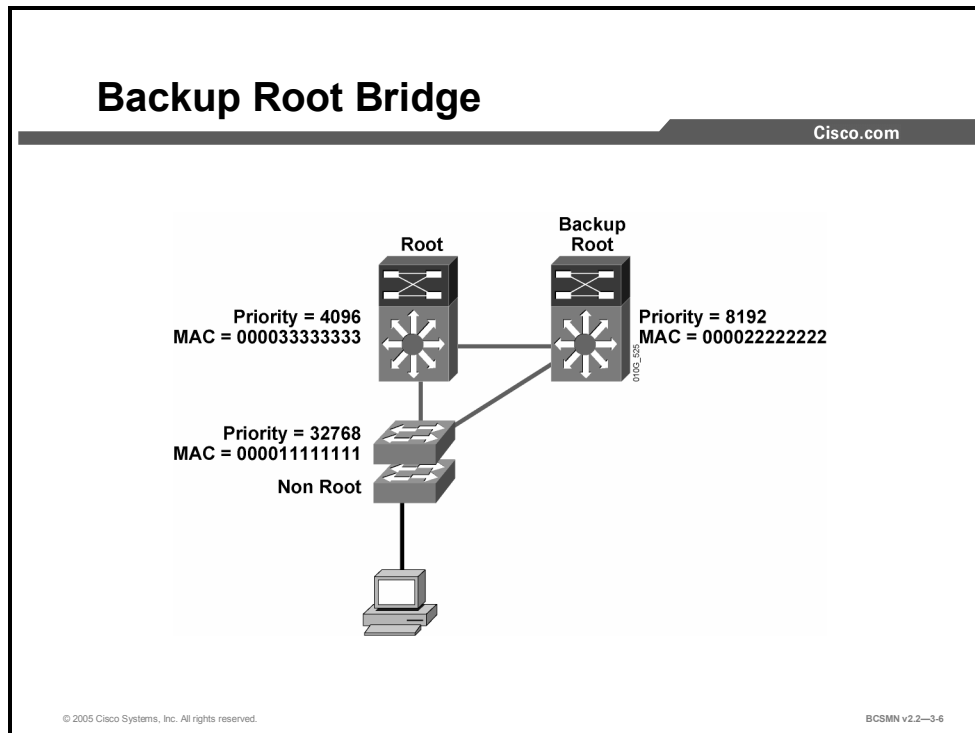


The steps that occur in a topology change are as follows:

- Step 1** Switch D notices that a change to a link has occurred.
- Step 2** Switch D sends a topology change notification (TCN) BPDU out the root port destined ultimately for the root bridge. 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 bridge. (This is a propagation TCN.)
- Step 4** When the root bridge receives the topology change message, it changes the flag portion of outbound BPDUs to indicate that a topology change is occurring. The root bridge 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 5-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.

What Is a Backup Root Bridge?

This topic describes a backup root bridge.



A backup (or secondary) root bridge is a switch that is preferentially configured to assume the role of the root bridge if the primary root bridge fails.

If no backup root bridge is configured and the root bridge fails, some other switch will be automatically chosen as the root bridge by the STP. However, it is likely that this automatic choice will not be optimal for network performance and stability.

The backup root bridge is configured to have a priority value set lower than the default but higher than the primary root bridge. In normal operation, when the primary root bridge is functioning, the backup root bridge behaves like any other nonroot bridge. When the primary root bridge fails, the backup root bridge then has the lowest priority in the network and so is selected to be the root bridge.

Configuration of an appropriate backup root bridge ensures optimal network forwarding and stability in the event of the primary root bridge failure.

Priority Commands

This topic describes the commands used to set switch priority on various Catalyst models.

Priority Commands

Cisco.com

Configuring Spanning Tree

- **spanning-tree vlan 10 root primary**
- **spanning-tree vlan 10 root secondary**

© 2005 Cisco Systems, Inc. All rights reserved. BCSMN v2.2-3-7

The priority of a switch can be altered to change the probability of its selection as a root bridge.

Note Commands to change priority vary by switch model, but typical commands are shown here.

Spanning Tree Priority Commands

Command	Description
Switch(config)# spanning-tree vlan <i>vlan_id</i> root primary	Sets STP priority to a value lower than the lowest priority value currently seen in BPDUs received by this switch.
Switch(config)# spanning-tree vlan <i>vlan_ID</i> root secondary	Sets STP priority to a fixed value lower than the default of 32768.
Switch(config)# spanning-tree vlan <i>vlan_ID</i> priority <i>bridge-priority</i>	Manually sets bridge priority to a specific value.
Switch(config)# spanning-tree vlan <i>vlan_ID</i> root primary diameter [2-7]	Configured at the root bridge to set the diameter for spanning tree, which in turn modifies all three spanning tree timers. The configured diameter value will not display in the configuration; only the altered timer values will be displayed.

How to Configure a Root Bridge

This topic identifies the commands for configuring a switch as the root bridge.

Configuring the Root Bridge

Cisco.com

```
Switch(config)#spanning-tree vlan 1 root primary
```

- **This command forces this switch to be the root.**

```
Switch(config)#spanning-tree vlan 1 root secondary
```

- **This command configures this switch to be the secondary root.**

© 2005 Cisco Systems, Inc. All rights reserved.BCSMN v2.2—3-8

The switch with the lowest bridge ID (BID) will become the root bridge for a VLAN. Specific configuration commands are used to determine which switch will become the root bridge.

A Catalyst switch running Per VLAN Spanning Tree (PVST) maintains an instance of spanning tree for each active VLAN configured on the switch. A unique BID 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 results in the recomputation of the root bridge for the VLAN.

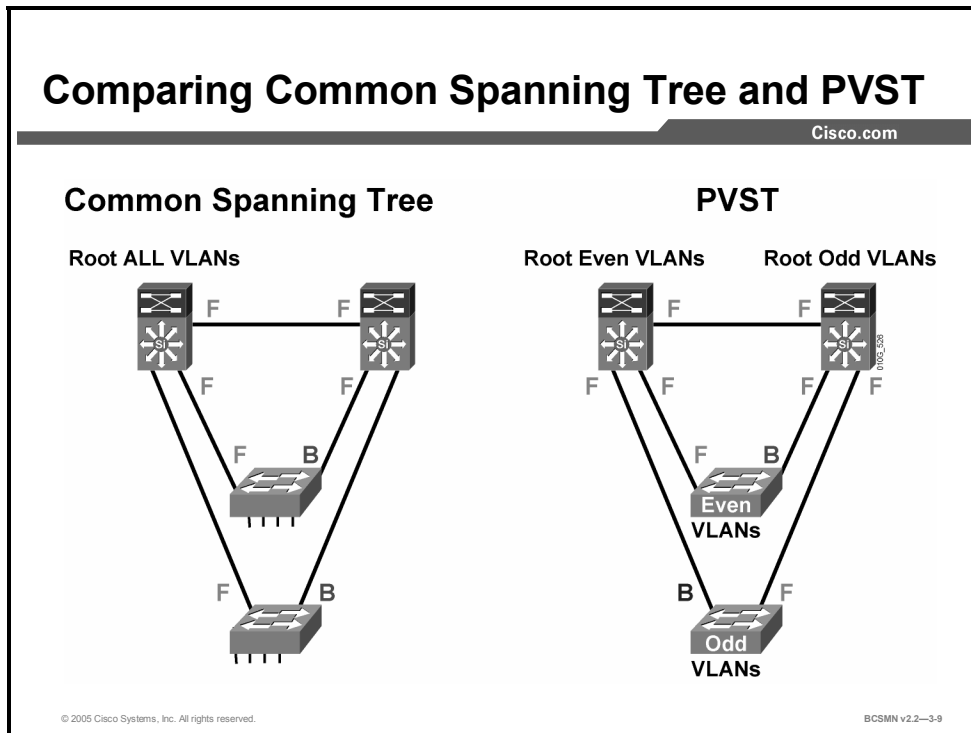
To configure a switch to become the root bridge for a specified VLAN, use the command **spanning-tree vlan *vlan-ID* root primary**. Assuming that the other bridges in the VLAN retain their default STP priority, this switch will become the root bridge.

Note The root bridge 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 the backbone switches should be the root bridge.

A secondary root is a switch that may become the root bridge for a VLAN if the primary root bridge fails. To configure a switch as the secondary root bridge for the VLAN, use the command **spanning-tree vlan *vlan-ID* root secondary**. Assuming that the other bridges in the VLAN retain their default STP priority, this switch will become the root bridge in the event that the primary root bridge fails. This command can be executed on more than one switch to configure multiple backup root bridges.

Comparing CST and PVST

This topic identifies the features of Per VLAN Spanning Tree.



The 802.1D standard for spanning tree assumes that a single instance of spanning tree runs for all VLANs. This is known as Common Spanning Tree (CST). Per VLAN Spanning Tree (PVST) provides the option of potentially establishing a root bridge on a VLAN-by-VLAN basis.

PVST is fully compatible with the 802.1Q trunking protocol and with Inter-Switch Link (ISL). PVST runs the same spanning tree algorithm (STA) as 802.1D 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 bridge, a separate set of designated switches, and associated port roles and states are calculated.

Example: Comparing CST and PVST

In the example, the CST operation causes all the uplinks from the access layer to go to one of the distribution switches. This is undesirable because these uplinks are fully loaded with traffic, while alternative path uplinks are inactive. Further, should you wish to configure the default gateway for each VLAN to be on alternative distribution switches, traffic will have to traverse the distribution-to-distribution link for half of the VLANs.

When PVST is used, each distribution switch is configured to be the primary root for each alternative VLAN. Dual uplinks from the access switches are both active but are forwarding for different VLANs. Also, given that access switches usually have a data and a voice VLAN, the root of each VLAN can be configured to be on alternative distribution switches.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **STP port states are used to transition switch ports to forwarding while maintaining a consistent topology between switches.**
- **Switches identify changes in the network topology and use the root bridge to establish an alternative traffic path.**
- **A backup root bridge is configured to assume the role of the root bridge in the event of primary root bridge failure.**
- **Primary and secondary root bridges are configured to provide Layer 2 fault tolerance.**
- **Common Spanning Tree has one instance for all VLANs, whereas PVST has one STP instance per VLAN.**

© 2005 Cisco Systems, Inc. All rights reserved.BCSMN v2.2—3-10

References

For additional information, refer to this resource:

- Cisco Systems, Inc., *Understanding SpanningTree Protocol Topology Changes*:
http://www.cisco.com/en/US/partner/tech/tk389/tk621/technologies_tech_note09186a0080094797.shtml

Configuring PortFast

Overview

A switch port will transition through a number of states to determine what action it should take when data frames are detected on the media of that port. These transitions are time-consuming and unnecessary for a port that is not connected to another switch. The PortFast feature allows this time to be reduced on ports to which end devices are connected.

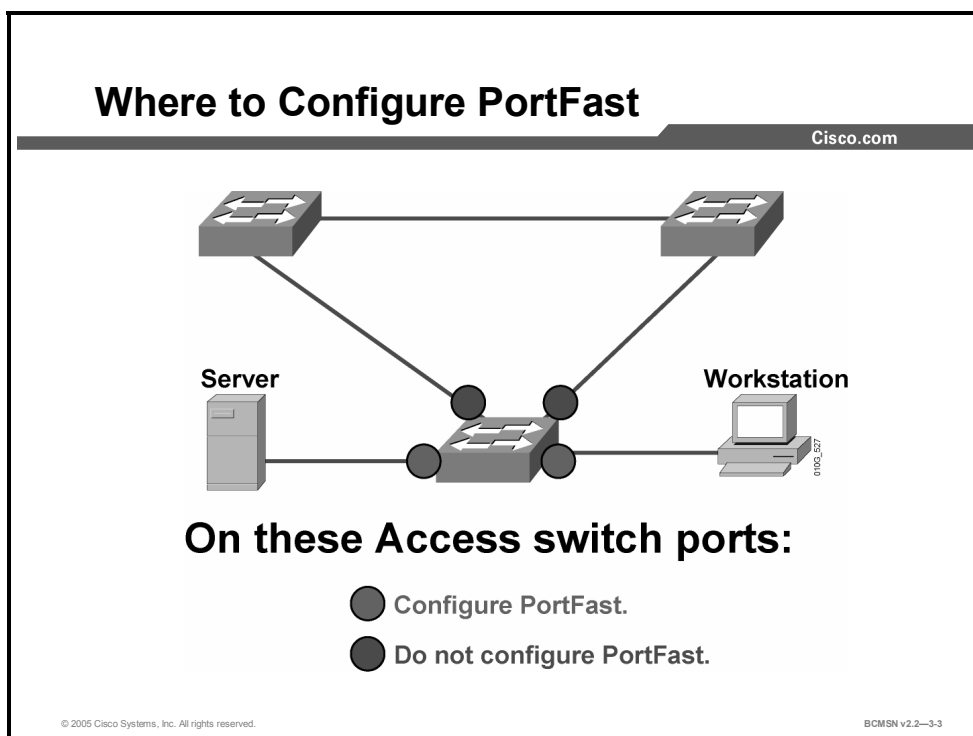
Objectives

Upon completing this lesson, you will be able to configure PortFast on Layer 2 access ports so that ports connected to a single workstation or server connect to the network immediately rather than waiting for spanning tree to transition the port to a forwarding state. This ability includes being able to meet these objectives:

- Describe PortFast
- List the commands used to configure PortFast
- Use the appropriate commands to configure and verify PortFast

What Is PortFast

This topic identifies the features of PortFast.



Spanning tree PortFast causes an interface configured as a Layer 2 access port to transition from blocking to forwarding state immediately, bypassing the listening and learning states. You can use PortFast on Layer 2 access ports that are connected to a single workstation or to a server to allow those devices to connect to the network immediately, rather than waiting for spanning tree to converge. If an interface configured with PortFast receives a bridge protocol data unit (BPDU), then spanning tree can put the port into the blocking state by using a feature called BPDU guard.

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.

PortFast Configuration Commands

This topic identifies the commands used to configure port-level tuning with PortFast.

Configuring PortFast

Cisco.com

Configuring

- **spanning-tree portfast**

Verifying

- **show running-config interface fastethernet 1/1**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—3-4

This table lists the commands used to implement and verify PortFast on an interface.

PortFast Commands

Argument	Description
Switch(config-if) # spanning-tree portfast	Enables PortFast on a Layer 2 access port and forces it to enter the forwarding state immediately.
Switch(config-if) # no spanning-tree portfast	Disables PortFast on a Layer 2 access port. PortFast is disabled by default.
Switch# show running-config interface type slot/port	Indicates if PortFast has been configured on a port. It can also be used to show if configuration has occurred on an EtherChannel link by specifying <i>port-channel</i> and <i>channel_number</i> in the place of <i>type mod/port</i> .

How to Configure PortFast

This topic explains how to apply PortFast in a campus network.

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
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-5

Spanning tree PortFast causes an interface configured as a Layer 2 access port to enter the forwarding state immediately, bypassing the STP listening and learning states. It is to be enabled only on access ports connected to end devices.

Configure PortFast

To enable PortFast on a Layer 2 access port, use this command:

```
Switch(config-if)#[no] spanning-tree portfast
```

Verify PortFast

To display interface information, including PortFast configuration, use this command:

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

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **PortFast allows a switch port to move directly from blocking to forwarding state.**
- **PortFast should only be configured on interfaces directly connected to end devices.**
- **PortFast is configured and verified on individual switch ports.**

© 2005 Cisco Systems, Inc. All rights reserved. BCM5N v2.2—3-6

Guarding Against Rogue STP Root Bridges

Overview

Once Spanning Tree Protocol (STP) operations are stable in a switched network, the administrator may want to guard against rogue switches being attached to the network, because these switches may usurp the role of the root or backup root bridge. Bridge protocol data unit (BPDU) guard, BPDU filtering, and root guard are features that attempt to contain the points at which switches and root bridges can be attached to the network.

Objectives

Upon completing this lesson, you will be able to configure BPDU guard, BPDU filtering, and root guard to prevent rogue Layer 2 switches from playing a key role in STP operations when placed on specific switch ports. This ability includes being able to meet these objectives:

- Describe the methods available to protect the operation of STP
- Identify the commands used to configure BPDU guard
- Identify the commands that are used to configure BPDU filtering
- Describe how root guard is used to improve the stability of Layer 2 networks
- List the commands used to configure root guard
- Use the appropriate commands to configure and verify root guard

Protecting Spanning Tree

This topic identifies the features of BPDU guard and BPDU filtering.

Protecting Spanning Tree

Cisco.com

Protection against switches being added on PortFast ports

- BPDU guard shuts ports down.
- BPDU filter specifies action to be taken when BPDUs are received.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-3

Cisco provides two features to protect spanning tree from loops being created on ports where PortFast has been enabled. In a proper configuration, PortFast would be enabled only on ports supporting end devices such as servers and workstations. It is anticipated that BPDUs from a switch device should not be received on a PortFast interface. However, should this happen, BPDU guard and BPDU filtering provide protection. Both BPDU guard and BPDU filtering can be configured globally on all PortFast-configured ports or on individual ports.

BPDU Guard

BPDU guard is used to protect the switched network from the problems that may be caused by the receipt of BPDUs on ports that should not be receiving them. The receipt of unexpected BPDUs may be accidental or may be part of an unauthorized attempt to add a switch to the network.

BPDU Filtering

PortFast BPDU filtering affects how the switch acknowledges BPDUs seen on PortFast-configured ports. Its functionality differs if it is configured globally or on a per-port basis and will be explained elsewhere in this course.

BPDU Root Guard

BPDU root guard protects against a switch outside the designated network attempting to become the root bridge by blocking its access until the receipt of its BPDUs ceases.

BPDU Guard Configuration Commands

This topic identifies the command for configuring BPDU guard on a switch.

Enabling and Verifying BPDU Guard

Cisco.com

```
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

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-4

BPDU guard protects the network from loops that might form if BPDUs are received on a PortFast-enabled switch port.

Note When the BPDU guard feature is enabled, spanning tree applies BPDU guard to all PortFast-configured interfaces.

BPDU Filtering Applied Globally Versus Per-Port

At the global level, you can enable BPDU guard on PortFast-enabled ports by using the **spanning-tree portfast bpduguard default** global configuration command. In a valid configuration, PortFast-enabled ports do not receive BPDUs. Receiving a BPDU on a PortFast-enabled port signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the port in the error-disabled state.

At the interface level, you can enable BPDU guard on any port by using the **spanning-tree bpduguard enable** interface configuration command without also enabling the PortFast feature. When the port receives a BPDU, it is put in the error-disabled state.

Configuring BPDU Guard

To enable BPDU guard, use this command:

```
Switch(config)# spanning-tree portfast bpduguard
```

The **no** form of the command will disable the feature on the switch.

Verifying BPDU Guard

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

BPDU Filtering Configuration Commands

This topic identifies the command used to enable BPDU filtering on a switch.

Enabling and Verifying BPDU Filtering

Cisco.com

```
Switch(config)#spanning-tree portfast bpdupfilter 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
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-6

BPDU Filtering Applied Globally Versus Per-Port

BPDU filtering can be configured globally or on individual PortFast-enabled ports.

When enabled globally, BPDU filtering has these attributes:

- It affects all operational PortFast ports on a switch that do not have BPDU filtering configured on the individual ports.
- If BPDUs are seen, the port loses its PortFast status, BPDU filtering is disabled, and STP sends and receives BPDUs on the port like any other STP port on the switch.
- Upon startup, the port transmits 10 BPDUs. If this port receives any BPDUs during that time, PortFast and PortFast BPDU filtering are disabled.

When enabled on an individual port, BPDU filtering has these attributes:

- It ignores all BPDUs received.
- It sends no BPDUs.

Caution Explicit configuration of PortFast BPDU filtering on a port not connected to a host station can result in bridging loops. The port ignores any incoming BPDUs and changes to the forwarding state. This does not occur when PortFast BPDU filtering is enabled globally.

The “BPDU Filtering Results” table lists the possible combinations that result from configuring BPDU filtering globally and on individual ports and on the same switch.

BPDU Filtering Results

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

Configuring BPDU Filtering

To enable PortFast BPDU filtering globally on the switch, enter this command:

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

To enable PortFast BPDU filtering on a specific switch port, enter this command:

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

Verifying BPDU Filtering

To verify the configuration on the switch, enter this command:

```
Switch#show spanning-tree summary totals
```

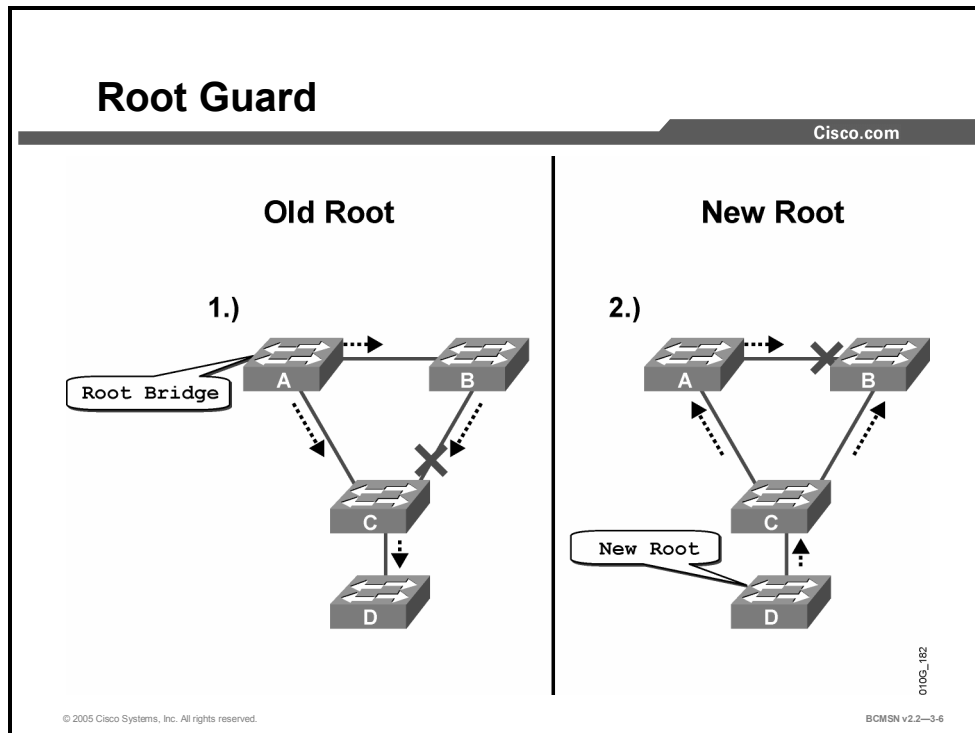
To verify the configuration on a specific port, enter this command to see the associated output:

```
Switch# show spanning-tree interface fastEthernet 4/4 detail
```

```
Port 196 (FastEthernet4/4) of VLAN0010 is forwarding
  Port path cost 1000, Port priority 160, Port Identifier 160.196.
  Designated root has priority 32768, address 00d0.00b8.140a
  Designated bridge has priority 32768, address 00d0.00b8.140a
  Designated port id is 160.196, designated path cost 0
  Timers:message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state:1
  The port is in the portfast mode by portfast trunk configuration
  Link type is point-to-point by default
  Bpdu filter is enabled
  BPDU:sent 0, received 0
```

Root Guard

This topic identifies the function of root guard.



Root guard limits the switch ports out of which the root bridge may be negotiated. If a root guard-enabled port receives BPDUs that are superior to those being sent by the current root bridge, then that port will be moved to a root-inconsistent state, which is effectively equal to an STP listening state. No data traffic will be forwarded across this port.

Example: Using Root Guard

In the example, switches A and B are the core of the network. Switch 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. If the priority of switch D were any value lower than that of the current root bridge, it would be a superior BPDU, and switch D would be elected the root bridge. This would cause the link connecting switches A and B to block, thus causing all traffic from switch B to flow through switch C in the access layer, which is clearly not advantageous. If root guard were configured on the port of switch C where switch D is attached, switch D would never have been elected the root bridge.

Root guard is configured on a per-port basis. If a superior BPDU is received on the port, root guard does not take the BPDU into account and so puts the port into root-inconsistent state. Once switch D stops sending superior BPDUs, the port will be unblocked again and will transition through STP states like any other port. Recovery requires no intervention. A root guard port is in an STP-designated state.

Root guard should be enabled on all ports where the root bridge is not anticipated. In the example, root guard should be enabled as follows:

- Switch A: port connecting to switch C
- Switch B: port connecting to switch C
- Switch C: port connecting to switch D

A root guard-enabled port is in an STP-designated port state.

The following console message appears 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
```

Root Guard Configuration Commands

This topic identifies the command used to configure root guard.

Root Guard Configuration Commands

Cisco.com

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

- **Configures root guard**

```
Switch#show running-config interface fa 0/1  
Switch#show spanning-tree inconsistentports
```

- **Verifies root guard**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—3-7

The following commands can be used to configure and verify root guard.

Root Guard Commands

Command	Description
Switch(config-if) # spanning-tree guard root	Enables root guard on an interface
Switch(config-if) # no spanning-tree guard root	Disables root guard on an interface
Switch# show running-config interface type mod/port	Indicates if root guard has been configured on an interface
Switch# show spanning-tree inconsistentports	Indicates if any ports are in a root-inconsistent state

How to Configure Root Guard

This topic discusses the steps used to configure and to verify root guard.

Enabling Root Guard

Cisco.com

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

- **Enables root guard on an interface**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-8

Configuring Root Guard

Here are the commands for configuring and verifying root guard.

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
```

Verifying Root Guard

This example demonstrates how to verify the root guard configuration.

Verifying Root Guard

Cisco.com

```
Switch#show running-config interface interface mod/port
```

- 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
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-9

To verify root guard use the following commands:

```
Switch#show running-config interface fastethernet 5/8
```

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

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **BPDU guard and BPDU filtering protect the operation of STP on PortFast-configured ports.**
- **When BPDU guard is configured globally it affects all PortFast configured ports.**
- **BPDU can be configured per port even on those ports not configured with PortFast.**
- **BPDU filtering can be configured globally or per port.**
- **The root switch cannot be elected via BPDUs received on a root guard–configured port.**
- **Root guard can be configured and verified using various commands.**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—3-10

Configuring UplinkFast

Overview

When a switch loses its connection to the root switch, determining an alternative path can be time consuming as the switch waits the requisite Spanning Tree Protocol (STP) time intervals. If a switch has a direct link to the root and a finite number of alternative paths to the root, this wait is not necessary. UplinkFast allows a switch to quickly failover to an alternative path to the root when the primary path fails due to a link fault.

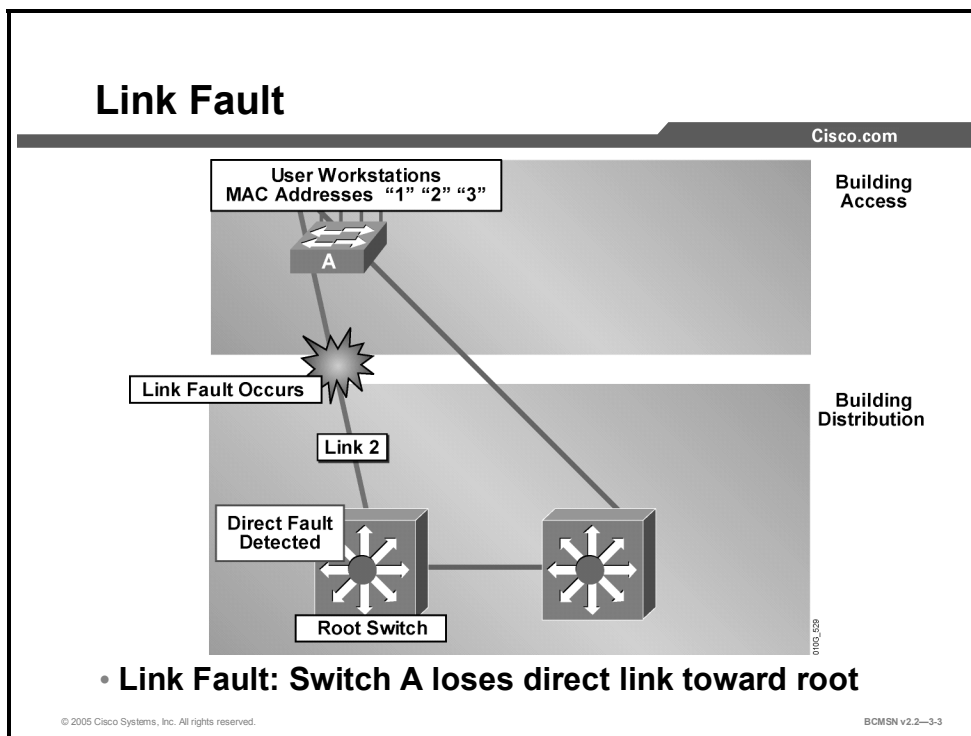
Objectives

Upon completing this lesson, you will be able to configure an UplinkFast group to create an alternative forwarding path if a link fails. This ability includes being able to meet these objectives:

- Define a link fault
- Describe UplinkFast
- Identify the commands that are used to configure UplinkFast
- List the steps for configuring UplinkFast

What Is a Link Fault?

This topic defines a link fault as it applies to STP.



In STP, a link fault is the loss of connectivity on a port directly connected to the switch. A link fault will occur on a switch root port. Loss of connectivity is determined by the absence of bridge protocol data units (BPDUs) on the link providing access to the root switch.

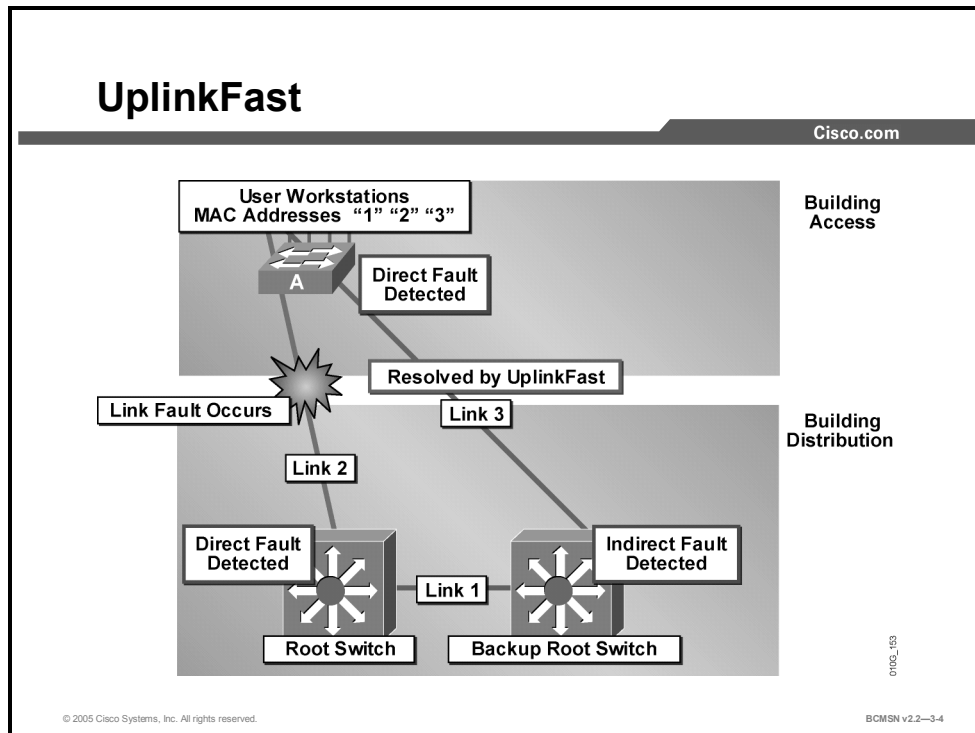
A link fault can occur as a result of the following:

- The media toward the root switch are physically disconnected.
- The media toward the root fail.
- The switch port at the other end of the root port fails, is disabled, or is shut down.
- Some event in software or hardware causes BPDUs not to arrive from the root switch on the port that a switch considers its root port.

In the example, switch A experiences a link fault. It is no longer receiving BPDUs over the link to the root switch.

UplinkFast

This topic identifies the features of UplinkFast.



Cisco's spanning tree UplinkFast provides fast convergence after a direct link failure. This immediate convergence is facilitated through the creation of an uplink group: a set of Layer 2 interfaces on a single switch, only one of which is forwarding at any given time. An uplink group consists of the root port (which is forwarding) and a set of blocked ports. The uplink group provides alternative failover paths if the root port link fails.

UplinkFast can failover to a backup link very quickly; therefore, the MAC address tables of other network switches must in turn be updated quickly to account for data traffic that should now traverse the backup path. To accomplish this, the UplinkFast switch will begin flooding frames with a source MAC address for all the entries in its content addressable memory (CAM) table to a destination Cisco proprietary multicast MAC address. These frames will be sent out the backup port. This will in turn populate the CAM table of switches on the backup path with MAC addresses that were previously learned through the failed link.

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 to the backup root switch over link 3. 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. UplinkFast is implemented on an access switch with at least one forwarding port and one blocked port toward the root.

UplinkFast Configuration Commands

This topic identifies the commands used to configure UplinkFast.

UplinkFast Configuration Commands

Cisco.com

Configure

- **spanning-tree uplinkfast**

Verify

- **show spanning-tree uplinkfast**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-3-5

The following commands enable and verify UplinkFast on a switch.

UplinkFast Commands

Argument	Description
Switch(config) # spanning-tree uplinkfast	Enables UplinkFast on a switch
Switch# show spanning-tree uplinkfast	Displays UplinkFast configuration information

How to Configure UplinkFast

This topic describes the steps for configuring UplinkFast.

Enabling UplinkFast

Cisco.com

```
Switch(config)#spanning-tree uplinkfast [max-update-rate  
max_update_rate]
```

- **Enables UplinkFast**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—3-6

UplinkFast is enabled on a switch rather than on a port. When enabled, it 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, which makes it unlikely that the switch will become the root switch. If bridge priority or port cost has been manually configured on a switch, UplinkFast will not alter those spanning tree values; it will only alter default values.

Enabling UplinkFast affects all VLANs on the switch. You cannot configure UplinkFast on an individual VLAN.

Configuring UplinkFast

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 dummy multicast packets transmitted per second when a failover occurs. The default value is 150 packets per second (pps).

For example, this is the command that would be used to enable UplinkFast with a maximum update rate of 400 pps:

```
Switch(config)#spanning-tree uplinkfast max-update-rate 400
```

Note UplinkFast should only be configured on access switches.

Verifying UplinkFast

This subtopic demonstrates how to verify that UplinkFast has been configured properly.

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)
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-7

Use the **show spanning-tree uplinkfast** command to verify UplinkFast configuration.

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
VLAN10
VLAN15
VLAN1002             Gi5/7(fwd)
VLAN1003             Gi5/7(fwd)
VLAN1004             Gi5/7(fwd)
```


Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **A link fault occurs when a link to the root switch fails.**
- **UplinkFast provides immediate convergence after a direct link failure.**
- **UplinkFast is configured on a per-switch basis.**
- **Specific commands are used to configure and verify UplinkFast.**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-8

Configuring BackboneFast

Overview

An indirect link fault is a fault that disrupts communication to the root on a link that is not directly connected to a switch. When this occurs, BackboneFast allows a switch to examine the information that is carried in inferior bridge protocol data units (BPDUs) prior to the expiration of the max age timer, allowing failover to an alternative root switch in an expedited manner.

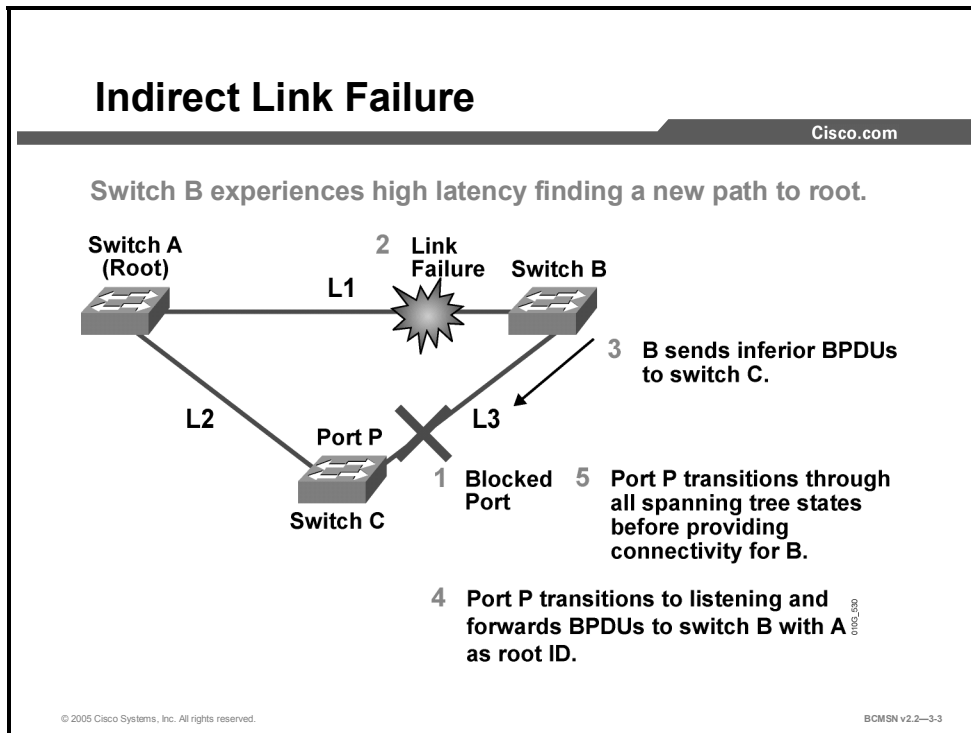
Objectives

Upon completing this lesson, you will be able to configure BackboneFast on a switch to create an alternative path through an intermediate switch when an indirect link failure causes a topology change. This ability includes being able to meet these objectives:

- Define an indirect link failure
- Describe BackboneFast
- Identify the commands that are used to configure BackboneFast
- List the steps for configuring BackboneFast

What Are Indirect Link Failures?

This topic defines an indirect link failure.



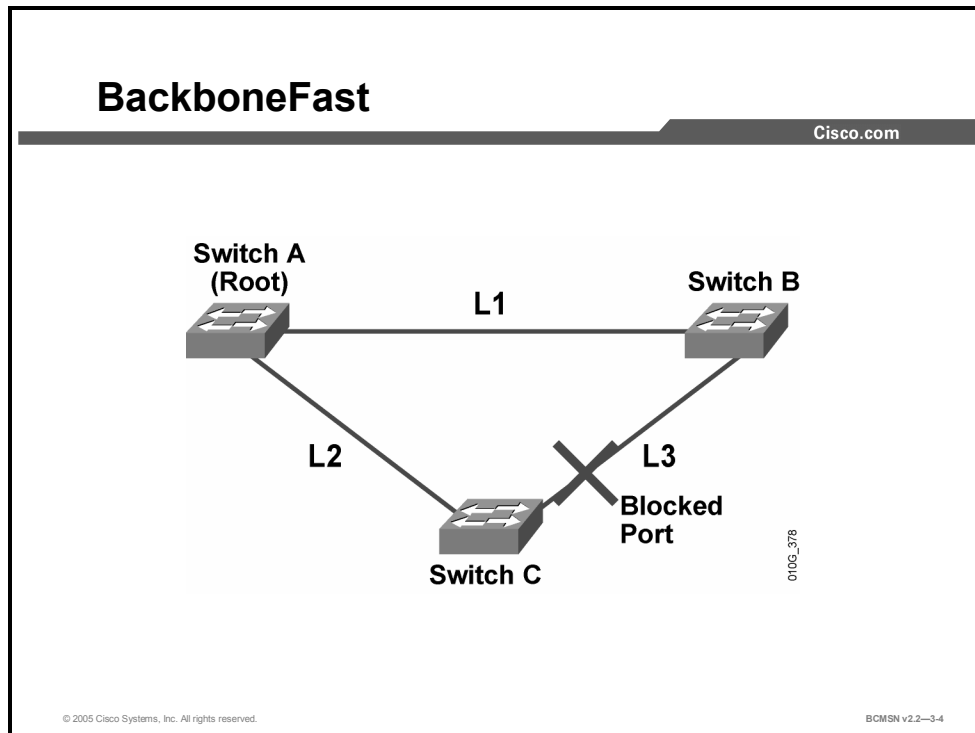
When a switch receives an inferior bridge protocol data unit (BPDU), it indicates that a link to which the switch is not directly connected (an indirect link) has failed; that is, the designated bridge has lost its connection to the root bridge.

Consider the example involving three switches in a fully meshed topology. Assume that switch A is the root bridge. Switch C is blocking on port P, and switch B is the designated bridge for link (L) 3.

- If link 1 goes down, switch B detects the failure and sends BPDUs to switch C, claiming itself as the new root.
- When C receives this new BPDU from B, it realizes that it is inferior to the one stored for port P and ignores it for the duration of the max age timer (20 seconds by default).
- After the max age timer has expired, the BPDU stored on switch C for port P ages out. The port goes into listening state, and switch C starts sending its BPDUs to switch B with switch A as the root ID.
- As soon as switch B receives a BPDU from switch C with the root ID of switch A, then switch B stops sending its own bridge ID (BID) as the root ID in its BPDUs.
- After the expiration of the listening and learning states (twice the forward delay value), port P moves to the forwarding state, and connectivity is restored for switch B.

BackboneFast

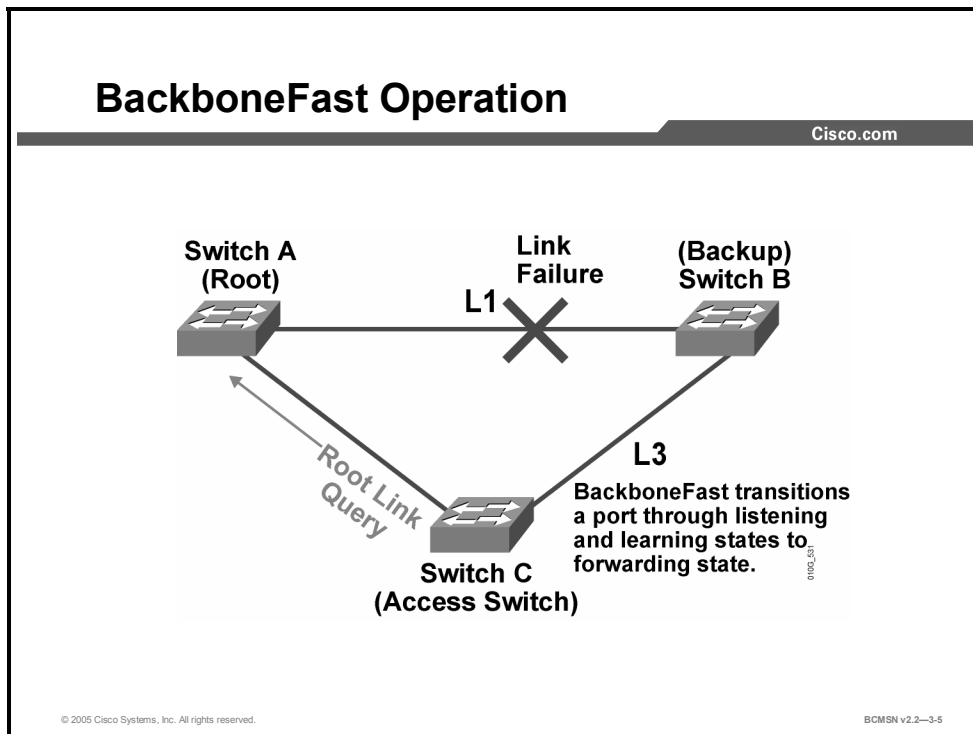
This topic identifies the features of BackboneFast.



BackboneFast addresses the situation in which an indirect failure causes a topology change, and therefore a switch must find an alternative path through an intermediate switch. BackboneFast is initiated when a root port or blocked port on a switch receives inferior BPDUs from its designated bridge. Under normal spanning tree rules, the switch ignores inferior BPDUs for the configured maximum aging time, as specified by the *agingtime* variable.

Example: BackboneFast Operation

This example addresses how BackboneFast operates.



BackboneFast operation is best illustrated by the failure of the link between the root and the backup root bridge. The backup root bridge does not assume that the root bridge is still available. The backup switch will immediately block all previously forwarding ports and then transmit configuration BPDUs, claiming root responsibility.

When the access switch receives the BPDUs of the backup root bridge, the access switch views the BPDUs as inferior because its own root port is still active, and the last indication it has is that the backup root bridge is the designated root bridge. If configured for BackboneFast, the access switch then transmits a special root query message to explicitly determine if the root bridge is still active. Upon receipt of a response, the access switch sends a BPDUs using its known root bridge parameters to the backup root bridge and cycles the port connected to the backup root bridge 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.

BackboneFast Configuration Commands

This topic identifies the commands used to configure BackboneFast.

BackboneFast Configuration Commands

Cisco.com

Configure

- **spanning-tree backbonefast**

Verify

- **show spanning-tree backbonefast**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—3-6

This table shows the commands used to configure and verify BackboneFast.

BackboneFast Commands

Argument	Description
Switch(config) # spanning-tree backbonefast	Enables BackboneFast on an entire switch. BackboneFast must be enabled on all switches in a spanning tree instance.
Switch(config) # no spanning-tree backbonefast	Disables BackboneFast on the switch.
Switch# show spanning-tree backbonefast	Displays current status of BackboneFast operation and statistics.

How to Configure BackboneFast

This topic describes the steps for configuring BackboneFast.

Enabling and Verifying BackboneFast

Cisco.com

```
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
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-7

BackboneFast must be enabled on all switches in the spanning tree instance. BackboneFast will interoperate with third-party switches. However, it is a Cisco proprietary feature and is not supported on Token Ring VLANs.

Configure BackboneFast

To enable or disable BackboneFast, use this command:

```
Switch(config)# [no] spanning-tree backbonefast
```

Verify BackboneFast

To verify BackboneFast configuration, use this command:

```
Switch#show spanning-tree backbonefast
```

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **An indirect link failure occurs on a link that is not directly connected to a device.**
- **BackboneFast allows immediate recovery from indirect link failure.**
- **BackboneFast must be configured on all switches in the spanning tree instance.**
- **Various commands are used to configure and verify BackboneFast.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—3-8

References

For additional information, refer to this resource:

- Cisco Systems, Inc., *Understanding and Configuring BackboneFast on Catalyst Switches*
http://www.cisco.com/en/US/partner/tech/tk389/tk621/technologies_tech_note09186a00800c2548.shtml

Configuring EtherChannel

Overview

When multiple physical links exist between two switches, these links can be bundled into a single logical link that provides high aggregate bandwidth and fault tolerance for inter-switch connectivity. This lesson will examine the specifics of EtherChannel.

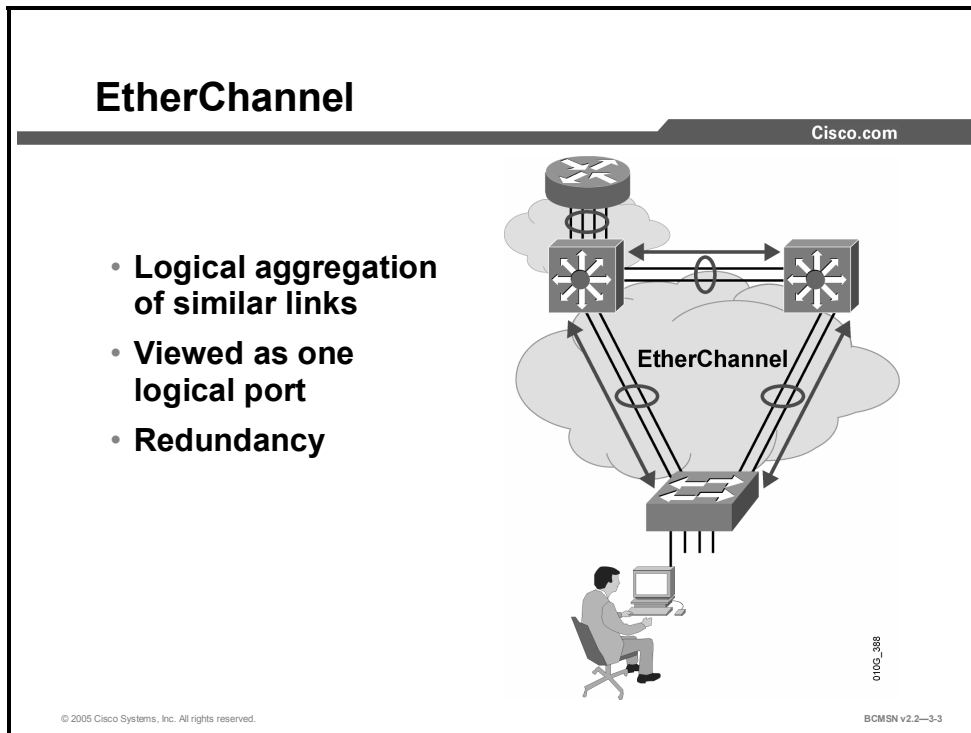
Objectives

Upon completing this lesson, you will be able to configure EtherChannel in order to support aggregated bandwidth and provide redundancy between two devices. This ability includes being able to meet these objectives:

- Describe EtherChannel
- Compare PAgP and LACP
- Configure and verify EtherChannel operation
- Describe load balancing of traffic over the links in an EtherChannel bundle
- Identify the guidelines and best practices for configuring EtherChannel

EtherChannel

This topic identifies the characteristics of EtherChannel.



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 duplex and must be configured as either Layer 2 or Layer 3 interfaces.

If a link within the EtherChannel bundle fails, traffic previously carried over the failed link will be carried over the remaining links within the EtherChannel

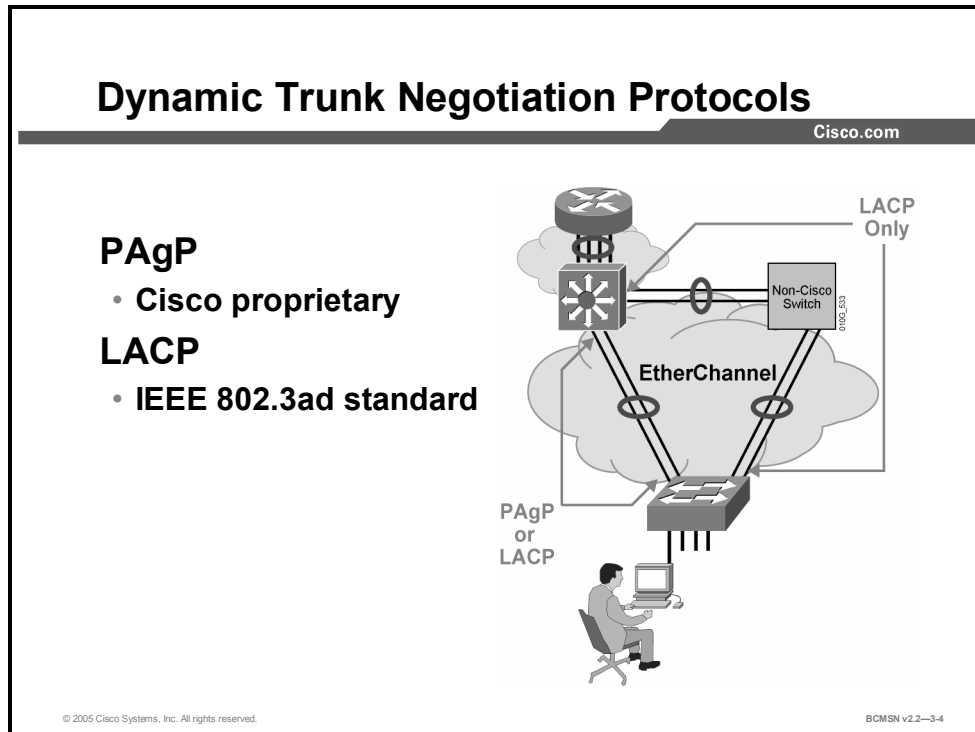
The configuration applied to the individual physical interfaces that are to be aggregated by EtherChannel affects only those interfaces. Each EtherChannel has a logical port channel interface. A configuration applied to the port channel interface affects all physical interfaces assigned to that interface. (Such commands can be Spanning Tree Protocol (STP) commands or commands to configure a Layer 2 EtherChannel as a trunk.)

Etherchannel Features and Benefits

- Allows for the creation of a very high-bandwidth logical link
- Load balances among the physical links involved
- Provides automatic failover
- Simplifies subsequent logical configuration (configuration is per logical link instead of per physical link)

PAgP and LACP Protocols

This topic discusses the features of Port Aggregation Protocol and Link Aggregation Control Protocol.



The Port Aggregation Protocol (PAgP) aids in the automatic creation of Fast EtherChannel links. PAgP packets are sent between Fast EtherChannel-capable ports to negotiate the forming of a channel. When PAgP identifies matched Ethernet links, it groups the links into an EtherChannel. The EtherChannel is then added to the spanning tree as a single bridge port.

Link Aggregation Control Protocol (LACP) is part of an IEEE specification (802.3ad) that allows several physical ports to be bundled together to form a single logical channel. LACP allows a switch to negotiate an automatic bundle by sending LACP packets to the peer. It performs a similar function as PAgP with Cisco EtherChannel. Because LACP is an IEEE standard, it can be used to facilitate EtherChannels in mixed-switch environments.

Interface Modes

Interfaces can be set in any of several modes to control EtherChannel formation.

This mode enables Etherchannel.

- **On:** This is the mode that forces the interface to channel without PAgP or LACP.

The next two modes enable 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 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.

The EtherChannel modes that use LACP are as follows:

- **Passive:** This LACP mode places a port in a passive negotiating state. In this state, the port responds to the LACP packets that it receives but does not initiate LACP packet negotiation (default).
- **Active:** This LACP mode places a port in an active negotiating state. In this state, the port initiates negotiations with other ports by sending LACP packets.

LACP Parameters

These 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:
 - The port physical characteristics, such as data rate, duplex capability, and point-to-point or shared medium
 - The configuration constraints that you establish

When enabled, LACP attempts to configure the maximum number of compatible ports in a channel. In some instances, LACP is not able to aggregate all the ports that are compatible; for example, the remote system might have more restrictive hardware limitations. When this occurs, 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.

EtherChannel Configuration Commands

This topic identifies the commands used to configure EtherChannel.

About EtherChannel Configuration Commands

Cisco.com

Configure PAgP

- interface port-channel 1
- channel-protocol pagp
- channel-group 1 mode auto

Verify

- show interfaces fastethernet 0/1 etherchannel
- show etherchannel 1 port-channel
- show etherchannel 1 summary

© 2005 Cisco Systems, Inc. All rights reserved.
BCMSN v2.2—3-5

These commands are used to configure and verify EtherChannel.

EtherChannel Configuration Commands

Command	Description
Switch(config)# interface port-channel port-channel-number	Creates a port-channel interface and moves to port-channel configuration mode, allowing the configuration of port-channel interface configuration parameters
Switch(config-if)# interface media-type slot/port	Moves to configure physical ports into EtherChannel bundles
Switch(config-if)# channel-group number mode mode_type	Associates an interface with a specific port-channel group and specifies if negotiation is to occur
Switch(config)# Port-channel load-balance load-balance-type	Instructs the switch how to load balance traffic over the individual links in the EtherChannel bundle
Switch# Show running-config interface port-channel channel_number	Shows the running configuration for a specific port-channel interface
Switch# show running-config interface type mod/port	Shows the running configuration for a specific physical interface
Switch# show interfaces type mod/port	Displays information on a physical interface that is specific to its role in an

Command	Description
etherchannel	EtherChannel bundle
Switch# show etherchannel num port-channel	Displays information on the current state of the port-channel interface
Switch# show etherchannel num summary	Displays a one-line summary per channel-group

Configuring Port Channels Using EtherChannel

This topic covers the guidelines and best practices for configuring port channels using EtherChannel.

Configuring Layer 2 EtherChannel

Cisco.com

```
Switch(config)#interface range interface slot/port - port
```

- Specifies the interfaces to configure in the bundle

```
Switch(config-if-range)#channel-protocol {pagp | lacp}
```

- Specifies the channel protocol—either pagp or lacp

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

- Creates the port-channel interface and places the interfaces as members

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—3-6

This table shows the steps for configuring and verifying an EtherChannel interface.

Configure a Layer 2 EtherChannel Bundle

Step	Action	Notes
1.	Switch(config)# interface range fastethernet [interface_range]	Specifies the interfaces that will comprise the EtherChannel group
2.	Switch(config-if-range)# channel-protocol {pagp lacp}	Specifies the channeling protocol to be used
3.	Switch(config-if-range)# channel-group 2 mode desirable	Creates the port-channel interface, if necessary, and assigns the specified interfaces to it

Configuring Layer 3 Etherchannel

This subtopic shows the commands for configuring a Layer 3 EtherChannel.

Configuring Layer 3 EtherChannel

Cisco.com

```
Switch(config)#interface port-channel port-channel-number
```

- **Creates a port-channel interface**

```
Switch(config-if)#no switchport  
Switch(config-if)#ip address address mask
```

- **Specifies L3 and assigns an IP address and subnet mask to the EtherChannel**

```
Switch(config)#interface interface slot/port
```

- **Specifies an interface to configure**

```
Switch(config-if)#no switchport  
Switch(config-if)#channel-group number mode {auto | desirable | on}
```

- **Configures the interface as L3 and specifies the port channel and the PAgP mode**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-7

This table shows the steps for configuring and verifying a Layer 3 EtherChannel interface.

Configure EtherChannel

Configure a Layer 3 EtherChannel Bundle

Step	Action	Notes
1.	Create the port-channel interface. Switch(config)# interface port-channel 1	Creates a virtual Layer 2 interface
2.	Switch(config-if)# no switchport	Changes interface to Layer 3 to enable the use of the IP address command.
3.	Assign an IP address to the port-channel interface, as this will be a Layer 3 interface. Switch(config-if)# ip address 172.32.52.10 255.255.255.0	Assigns an IP address to the port-channel interface.
4.	Navigate to the interface that is to be associated with the EtherChannel bundle. Switch(config)# interface range fastethernet 5/4 - 5	This example shows navigation to a range of interfaces with the port-channel. Individual interfaces can be used as well.

Step	Action	Notes
5.	Prepare interface. <pre>Switch(config-if-range) # no switchport Switch(config-if-range) # channel-protocol pagp</pre>	The independent Layer 2 and Layer 3 functionality of the port must be removed so the port can function as part of a group. Optionally, can specify the channel protocol.
6.	Associate physical interfaces with the port-channel. <pre>Switch(config-if-range) # channel-group 1 mode desirable</pre>	Assigns all of the physical interfaces in the range to the EtherChannel group.

Verifying EtherChannel

This subtopic addresses the commands that can be used to verify the EtherChannel configuration.

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 : 66 bytes
!
interface Port-channel1
 switchport mode dynamic desirable
end
```

```
Switch#show run interface gig 0/9
Building configuration...

Current configuration : 127 bytes
!
interface GigabitEthernet 0/9
 switchport mode dynamic desirable
 channel-group 2 mode desirable
 channel-protocol pagp
end
```

© 2005 Cisco Systems, Inc. All rights reserved.
BCMSN v2.2-3-8

Use the **show interfaces [interface] [num] etherchannel** command to display information about the port channel and the specific EtherChannel interfaces.

Example: Verifying the Configuration of a Layer 3 EtherChannel

```
Switch#show interfaces fastethernet 5/4 etherchannel

Port state      = EC-Enbl'd 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:

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 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
```

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  = 2                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:

Group	Flags	State	Timers	Hello Interval	Partner Count	PAGP Priority	Learning Method
Fa5/6 56	SC	U6/S7		30s	1	128	Any

Partner's information:

Partner Group	Partner Name	Partner Device ID	Partner Port	Age	Flags
Fa5/6 2F	JAB031301	0050.0f10.230c	2/47	18s	SAC

Age of the port in the current state: 00h:10m:57s

Example: Verifying Port-Channel Configuration

Verifying EtherChannel (Cont.)

Cisco.com

```
Switch#show interfaces gigabitethernet 0/9 etherchannel
Port state      = Up Mstr In-Bndl
Channel group   = 1          Mode = Desirable-S1      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:
Port      Flags State  Timers  Hello  Partner  PAGP   Learning  Group
Gi0/9    SC   U6/S7  H       30s   1        128    Any       15

Partner's information:
Port      Partner      Partner      Partner      Partner Group
Gi0/9    Name         Device ID    Port         Age  Flags  Cap.
        DSW122      0005.313e.4780  Gi0/9       18s SC    20001

Age of the port in the current state: 00d:20h:00m:49s
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-9

Use the **show etherchannel** command to display port-channel information after 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
```

Load Balancing over EtherChannel

This topic identifies the commands for configuring 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
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-3-10

EtherChannel balances traffic load across the links in a channel. Load balancing can be based upon the variables listed. The default load balancing method varies among the Catalyst models.

Load balancing is applied globally for all EtherChannel bundles in the switch. To configure EtherChannel load balancing, use the **port-channel load-balance** 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

Configuring and Verifying EtherChannel Load Balancing

This example shows how to configure and verify EtherChannel load balancing.

```
Switch(config)# port-channel load-balance src-dst-ip
Switch(config)# exit
Switch# show etherchannel load-balance
```

Source XOR Destination IP address

Guidelines and Best Practices for Configuring EtherChannel

This topic addresses guidelines and best practices for 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 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.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—3-11

Follow these guidelines and restrictions when configuring EtherChannel interfaces.

- **EtherChannel support:** 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.
- **Speed and duplex:** Configure all interfaces in an EtherChannel to operate at the same speed and in the same duplex mode. Also, if one interface in the bundle is shut down, it is treated as a link failure and traffic will traverse other links in the bundle.
- **SPAN and 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.
- **VLAN match:** All interfaces in the EtherChannel bundle must be assigned to the same VLAN or be configured as a trunk.

Guidelines for Configuring EtherChannel (Cont.)

Cisco.com

All interfaces must support the same allowed range of VLANs.

Interfaces in the same bundle can support varying port costs.

Port-channel interface configuration changes affect the EtherChannel.

Physical interfaces configuration changes affect the interface only.

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—3-12

- **Range of VLANs:** 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 auto or desirable mode. For Layer 2 EtherChannels, either assign all interfaces in the EtherChannel to the same VLAN or configure them as trunks.
- **STP path cost:** 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.
- **Port channel versus interface configuration:** 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.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **EtherChannel increases bandwidth and provides redundancy by aggregating of individual links between switches.**
- **EtherChannel can be dynamically configured between switches using either PAgP or LACP.**
- **Etherchannel is configured and verified using a variety of show commands.**
- **EtherChannel load balances traffic over all the links in the bundle.**
- **Best Practices should be followed for EtherChannel configuration.**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—3-13

Module Summary

This topic summarizes the key points discussed in this module.

Module Summary

Cisco.com

- **The Spanning Tree Protocol protects the network from Layer 2 frames that might loop.**
- **Primary and secondary switches should be configured to affect what switches will be selected as the STP root.**
- **PortFast is configured on end-device switch ports that need to transition quickly to a forwarding state.**
- **BPDU guard, BPDU filtering, and root guard protect the network from rogue switches entering STP negotiation.**
- **UplinkFast can quickly select an alternative link to the root should it experience a link fault on its link toward the root.**
- **BackboneFast can quickly select an alternative link to the root should an indirect link fault occur.**
- **EtherChannel adds redundancy and creates high-bandwidth connections between switches by bundling physical links together into one logical interface.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—3-1

This module reviewed the fundamentals of the Spanning Tree Protocol (STP) operation in a switched network. We discovered the purpose of the primary and root bridges and how they are elected. PortFast, UplinkFast, and BackboneFast were discussed in their various roles in enhancing the performance of STP. EtherChannel was examined as well as its interoperability with STP. This module provided guidelines for enhancing the resiliency of STP when a network fault occurs.

References

For additional information, refer to these resources:

- Cisco Systems, Inc., *Understanding Spanning-Tree Protocol Topology Changes*:
http://www.cisco.com/en/US/partner/tech/tk389/tk621/technologies_tech_note09186a0080094797.shtml
- Cisco Systems, Inc., *Understanding and Configuring BackboneFast on Catalyst Switches*:
http://www.cisco.com/en/US/partner/tech/tk389/tk621/technologies_tech_note09186a00800c2548.shtml

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which two statements best describe how spanning tree uses the spanning tree algorithm to prevent bridging loops in a redundant network? (Choose two.) (Source: Defining the Spanning Tree Protocol)
- A) Some ports are blocked.
 - B) Some ports send out a broadcast message.
 - C) Alternative switches are introduced.
 - D) Some redundant paths are blocked.
- Q2) Which two of the following are results of BPDUs being exchanged between switches? (Choose two.) (Source: Maintaining and Configuring STP)
- A) lowest cost path to the root switch determined by each 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
 - D) removal of loops in the switched network by disabling ports
- Q3) What is the purpose of STP in switched networks? (Source: Defining the Spanning Tree Protocol)
- A) to maintain redundant switching paths in case of failure
 - B) to automate VLAN creation between switches
 - C) to provide a path for all VLANs to traverse
 - D) to allow multiple switches to act like one
- Q4) When a switch port is configured in the PortFast mode, which two STP states does it skip? (Source: Configuring PortFast)
- A) blocking and learning
 - B) listening and learning
 - C) blocking and listening
 - D) blocking and forwarding
- Q5) The UplinkFast feature should be configured on which type of switch? (Source: Configuring UplinkFast)
- A) access switches
 - B) distribution switches with one forwarding and one blocked port toward the root
 - C) any core switch with one forwarding and one blocked port toward the root
 - D) any access switch with one forwarding and one blocked port toward the root
- Q6) The BackboneFast feature is designed to remove which delay in standard STP convergence? (Source: Configuring BackboneFast)
- A) max aging time
 - B) listening
 - C) blocking
 - D) forward delay
- Q7) Which command will configure a bridge as the root bridge for VLAN11? (Source: Maintaining and Configuring STP)
- A) **spanning-tree vlan 11 root**
 - B) **spanning-tree 11 root primary**
 - C) **spanning-tree root vlan 11 primary**

D) **spanning-tree vlan 11 root primary**

Q8) What are the two protocol choices you have when implementing an EtherChannel bundle? (Choose 2.) (Source: Configuring EtherChannel)

- A) PAgP
- B) PAgD
- C) LACP
- D) LAPD

Module Self-Check Answer Key

Q1) A, D

Q2) A, C

Q3) A

Q4) B

Q5) D

Q6) A

Q7) D

Q8) A, C

Enhancing Spanning Tree

Overview

Identifying where issues occur in a spanning tree environment and defining why these issues occur can make problem isolation easier. Enhancements have been made to the Spanning Tree Protocol (STP) to help reduce the likelihood of anomalous STP conditions. Error conditions can be detected and prevented through the execution of specific troubleshooting and configuration commands. Two major enhancements have been made to original 802.1D specifications for spanning tree. Rapid Spanning Tree Protocol RSTP, 802.1w, provides a faster and more resilient means of adapting to topology change in a switched environment. Multiple Spanning Tree provides an alternative to Per VLAN and Common Spanning Tree by requiring only a minimal number of STP instances in a switched environment.

Module Objectives

Upon completing this module, you will be able to troubleshoot spanning tree and identify enhancements to 802.1D provided by Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree (MST). This ability includes being able to meet these objectives:

- Identify problems common to Spanning Tree Protocol and how to resolve these problems
- Configure UDLD and loop guard to mitigate STP failures due to unidirectional links
- Implement RSTP to increase the speed of the recalculation of spanning tree when a network topology changes
- Implement MST to reduce the number of instances of spanning tree running in a network with multiple VLANs

Troubleshooting Spanning Tree

Overview

Spanning tree issues manifest themselves through a number of symptoms. Properly identifying these symptoms and taking appropriate action is critical to maintaining reliable spanning tree operation. A series of debug commands specific to spanning tree will assist you in identifying problems and suggest a proper troubleshooting approach.

Objectives

Upon completing this lesson, you will be able to identify problems common to Spanning Tree Protocol (STP) and resolve these problems to restore optimal network performance. This ability includes being able to meet these objectives:

- Identify common STP problems and their associated solutions
- Identify **debug** commands specific to spanning tree
- Identify the steps for detecting and resolving common STP problems

STP Problems

This topic identifies the most common STP problems.

STP Problems

Cisco.com

- **Duplex mismatch**
- **Unidirectional link failure**
- **Frame corruption**
- **Resource errors**
- **PortFast configuration error**
- **EtherChannel issues**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—4-3

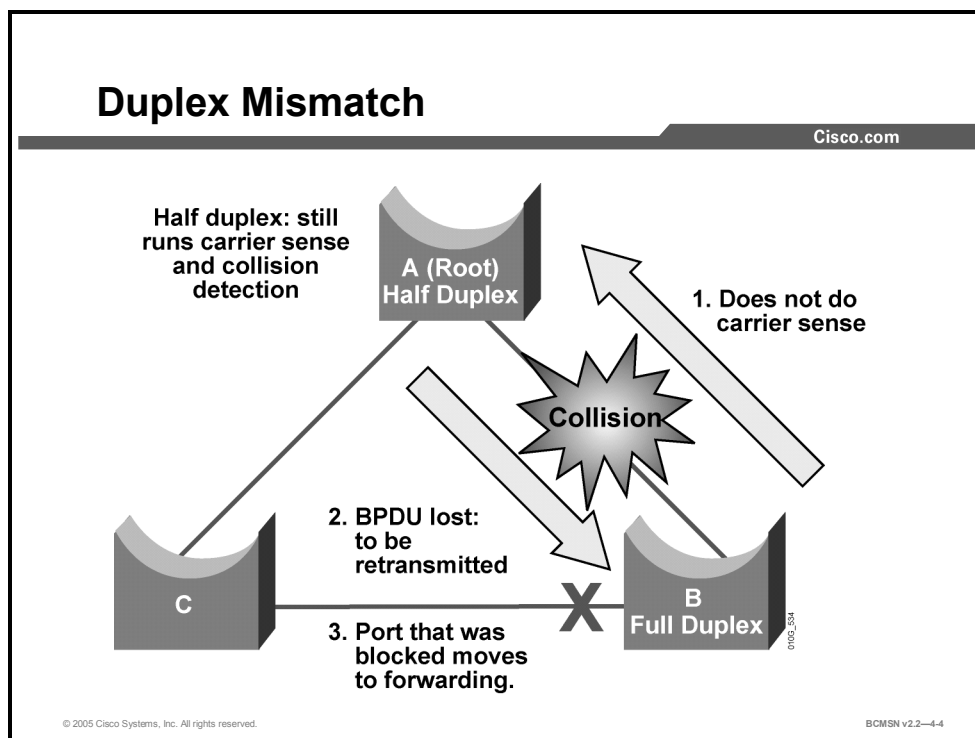
STP problems are most often evidenced by the existence of a bridge loop. Troubleshooting STP involves the identification and prevention of such loops.

The primary function of the spanning tree algorithm (STA) is to remove loops created by redundant links in bridged networks. The STP operates at Layer 2 of the Open Systems Interconnection (OSI) mode, exchanging bridge protocol data units (BPDUs) between bridges and selecting the ports that will eventually forward or block traffic. If BPDUs are not being sent or received over a link between switches, the role of the protocol in preventing loops may fail. Troubleshooting the resulting problems can be difficult in a complex network.

Any condition that prevents BPDUs from being sent or received can result in a bridge loop. Here is an explanation of how that condition may occur.

Duplex Mismatch

This subtopic describes how improper transmission type configurations can create problems for the STP.



Duplex mismatch on a point-to-point link is a common configuration error and can have specific implications for STP. The results of the mismatch may vary by platform.

Following are two common mismatch scenarios between switches and their resulting STP problems.

Switch configured for full duplex connected to a host in autonegotiation mode: The rule of autonegotiation is that, upon negotiation failure, a port is required to assume half-duplex operation. This creates a situation in which there is either no connectivity or inconsistent connectivity between the two devices as one side of the connection defaults to half-duplex mode and the other side is set to full-duplex operation. In many cases, this condition will allow traffic to flow at low-data rates, but, as the traffic level increases on the link, the half-duplex side of the link will be overwhelmed, causing data and link integrity errors. As the error rate goes up, BPDUs may not successfully negotiate the link.

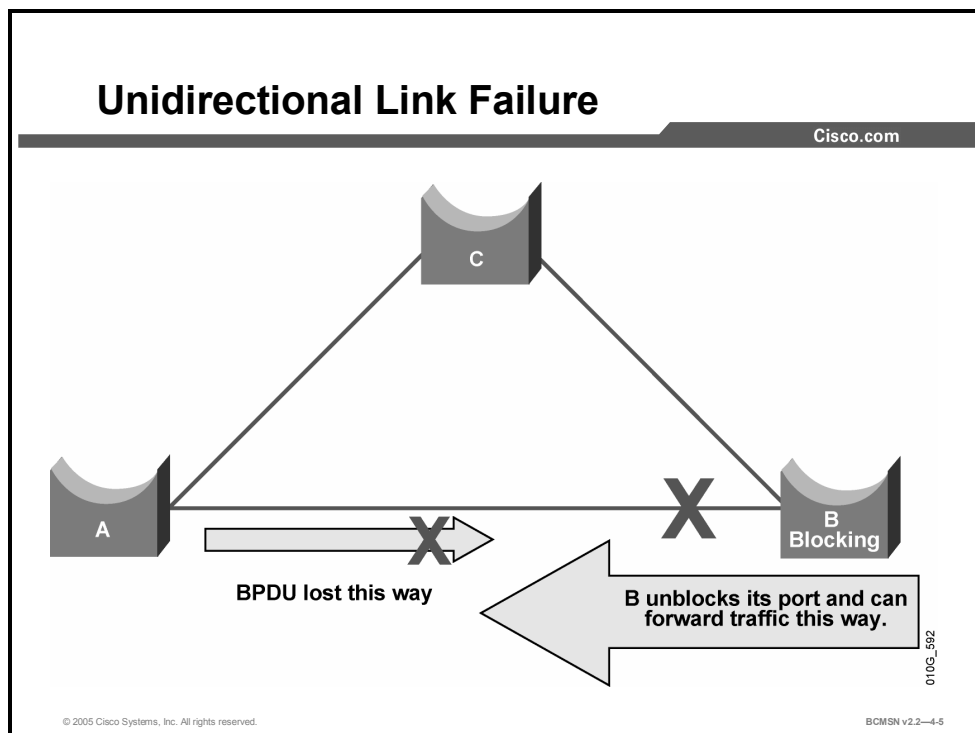
Switch is configured for half duplex on a link; the peer switch is configured for full duplex: In the example, the duplex mismatch on the link between bridge A and bridge B can lead to a bridge loop. Because B is configured for full duplex, it does not perform carrier sense when accessing the link. B 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 frames sent by A, including BPDUs, may be deferred or collide and eventually be dropped. Because it does not receive BPDUs from A, bridge B may lose its connection to the root. This will cause B to unblock its port to bridge C, thereby creating the loop.

To mitigate transmission type mismatches, the best practice is to establish a standard within the organization regarding how each interface is configured prior to attaching it to the network. It is not always possible to disable autonegotiation on all attached client devices, but, where possible, network infrastructure devices and servers should have matching transmission type

settings, and no switch ports should be set to autonegotiate. This will make it easier to troubleshoot these types of issues.

Unidirectional Link Failure

This subtopic identifies the scenario that best represents unidirectional link failure.



A unidirectional link stays up while providing only one-way communication. Unidirectional links cause specific STP problems. 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 the port on bridge B was blocking. A port will block only if it receives BPDUs from a bridge with a higher bridge ID (BID). In this case, all the BPDUs coming from bridge A are lost, so bridge B will never see the BPDU with the higher BID. Bridge B will unblock the port and eventually forward traffic, potentially creating a loop when other switches are in the scenario. If the unidirectional failure exists at startup, the STP will not converge correctly.

Frame Corruption

Frame corruption can result from duplex mismatch, bad cables, or incorrect cable length, and can lead to an STP failure. If a link is receiving a high number of frame errors, BPDUs can be lost. This may cause a port in blocking state to transition to forwarding. In 802.1D, if a blocking port does not see any BPDUs for 50 seconds, it will transition to the forwarding state. If a single BPDU were successfully transmitted, it would break the loop. This problem would be most likely if STP timing parameters, such as the max age value setting, had been set too low.

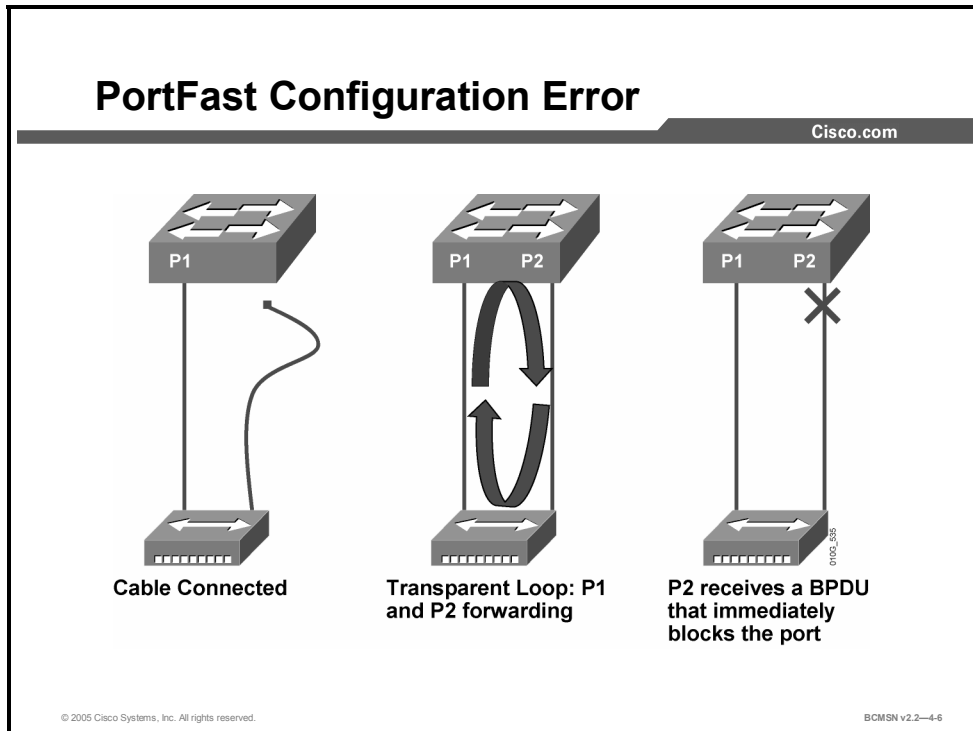
Resource Errors

STP is implemented in software. This means that, if the CPU of the bridge is over-utilized, the switch may lack the resources to send out or receive BPDUs in a timely manner. Lack of BPDUs can cause ports to transition from blocking to forwarding when they should not transition. This can cause loops to form in the network.

The STA, however, is not processor intensive and has priority over other processes. Therefore, a CPU utilization problem is unlikely on current Catalyst switch platforms.

PortFast Configuration Error

This subtopic identifies the issues that can arise if PortFast is used incorrectly.



PortFast is intended for configuration on a port connected to a single 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. If a switch is inadvertently attached to a PortFast port, a loop may occur, or this rogue switch may be elected as the STP root bridge. Furthermore, if a hub is attached to a PortFast port with redundant connections to the switch, a loop will be introduced that will go unchecked by STP.

In the example, A is a bridge with port 1 (P1) forwarding and port 2 (P2) configured for PortFast. B is a hub. As soon as the second cable is plugged into A, P2 goes to the forwarding state and creates a loop between P1 and P2, given that both ports are in forwarding state. As soon as P1 or P2 receives a BPDU, one of these two ports will transition to a blocking state. The traffic generated by this kind of loop may occur at such a high rate that the bridge may have trouble in successfully sending the BPDU to stop the loop. Implementing BPDU guard will prevent this problem.

EtherChannel Issues

The challenges for EtherChannel can be divided into two main areas: Troubleshooting during the configuration phase and troubleshooting during the execution phase. Configuration errors usually occur because of mismatched parameters on the ports involved (different speeds, different duplex, different spanning tree port values, mismatched native VLAN settings, and so forth). However, you can also generate errors during the configuration by setting the channel on one side to “on” and waiting too long before configuring the channel on the other side. This causes temporary spanning tree loops, which can generate an error and shut down the port.

Depending on the version of operating system and platform being configured for EtherChannel, the ports on one side of the link may remain in disabled state even after the configuration issue has been resolved. Be sure to verify that both sides of the link are operational after changing any port parameters on an EtherChannel link.

Spanning Tree debug Commands

This topic identifies the commands used to display STP debug information.

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

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-4-7

The **debug spanning-tree** command is used to troubleshoot spanning tree activities.

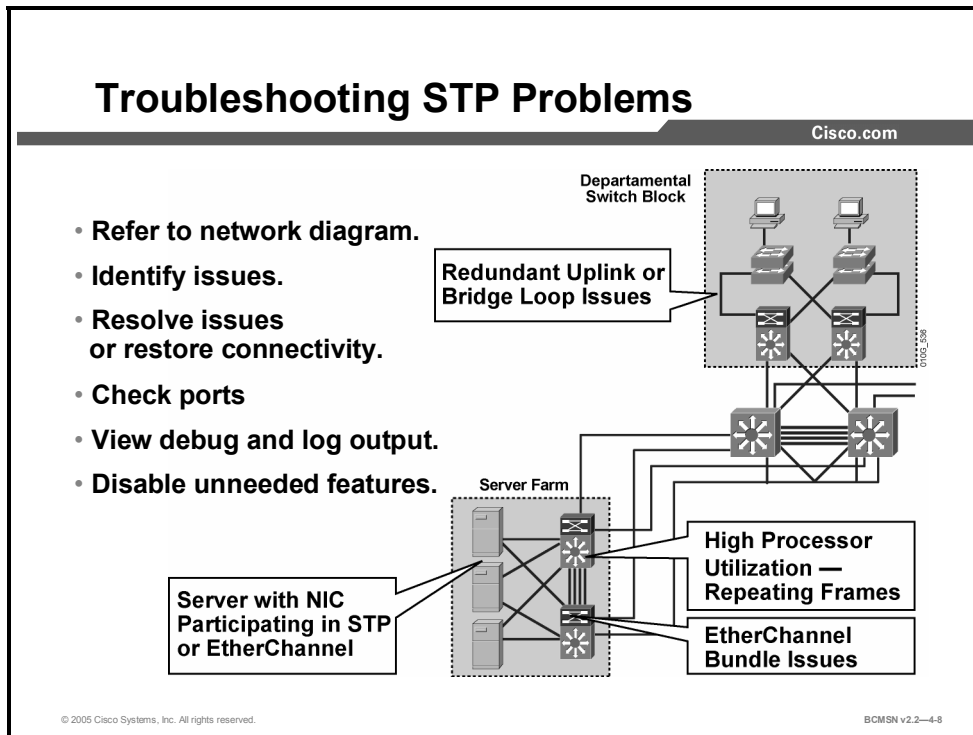
Spanning Tree debug Commands

Command	Description
Switch# debug spanning-tree all	Displays all spanning tree debug messages
Switch# debug spanning-tree events	Displays spanning tree topology debug event messages
Switch# debug spanning-tree backbonefast	Enables debugging of spanning tree BackboneFast events
Switch# debug spanning-tree uplinkfast	Enables debugging of spanning tree UplinkFast events

Caution Debugging, especially to a console or auxiliary port, can cause excessive processor utilization on an infrastructure device. If debugging commands must be issued, consider disabling console logging and send the output to a terminal session. Additionally, remember to turn off all debugging when the appropriate data has been collected.

How to Troubleshoot STP Problems

This topic identifies an approach for troubleshooting STP so that the problem can be resolved effectively.



Refer to a Network Diagram

Collect the following network information before troubleshooting a bridging loop. Knowledge of the following items in your environment is critical:

- The physical and logical topology of the bridged network
- Where the root bridge is located (for all VLANs if Per VLAN Spanning Tree [PVST] is in use)
- Where the redundant links and blocked ports are to be located

Identify Issues

This knowledge is essential for at least the following two reasons.

- To identify a problem, you need to know how the STP network should be laid out when it is operating correctly.
- The STP troubleshooting steps use **show** commands to display error conditions. Knowledge of the network helps focus your attention on the critical portions of these displays.

Identify a Bridge Loop

The best way to identify a bridge loop is to capture the traffic on a saturated link and check whether identical frames are traversing multiple links. Bridge loops often result in high port utilization due to excessive frames. Check the port utilization on your devices and look for abnormal values.

You can monitor STP operations using the **debug spanning-tree** command. This command is helpful in verifying correct bridging operation as well as in identifying loops.

Restore Connectivity Versus Resolve Issues

Bridge loops have severe consequences in a switched network. When one occurs, administrators generally do not have time to identify the reason for the loop during working hours and will often take temporary measures to stabilize the network but never resolve the actual problem. It is important to recreate and correct the original problem at a planned network downtime.

Break the Loop Disabling Ports

A simple troubleshooting approach is to manually disable ports providing Layer 2 redundancy. Begin 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, it is a good indication that the failure was located on a redundant path where this port was located.

Log STP Events on Devices Hosting Blocked Ports

If you cannot identify precisely the source of an STP problem, or if the problem is only transient, enable logging of STP events on the bridges and the switches of the network that is experiencing the failure. At a minimum, enable logging on devices hosting blocked ports, because it is typically the transition of a blocked port to forwarding that creates a 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 in the buffers of the device.

Check Ports

The ports to be investigated first are the blocking ports. Here is a list of what to check for on the various 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 and root ports.

If you are running Cisco IOS Release 12.0 or a later release, the command **show spanning-tree <bridge-group #>** displays a field named BPDU, which displays the number of BPDUs received on each interface. Issuing the command several times will indicate if the device is receiving BPDUs.

Check for Duplex Mismatch

To look for a duplex mismatch, check each side of a point-to-point link. 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.

STP debug Command

The command **debug spanning-tree** is very useful for troubleshooting STP issues. It accepts 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
```

General Recommendations

In general, it is difficult to troubleshoot spanning tree problems in a very large, flat switched network. If the network is being restructured, it is advisable to implement a hierarchical network structure that is designed around the Campus Infrastructure module. This would create manageable failure domains and reduce the overall network complexity.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Problems with STP can result from configuration errors, link failures, unidirectional links, frame corruption, or resource limitations.**
- **A series of debug commands specific to spanning tree is provided to properly identify problems and a troubleshooting approach.**
- **Taking the proper steps to identify and resolve STP problems is critical to maintaining reliable spanning tree operation.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—4-9

References

For additional information, refer to this resource:

- Cisco Systems, Inc., *Spanning Tree Protocol Problems and Related Design Considerations*,
http://www.cisco.com/en/US/partner/tech/tk389/tk621/technologies_tech_note09186a00800951ac.shtml

Preventing STP Forwarding Loops

Overview

Spanning tree operations can be severely disrupted by links that pass traffic in one direction and not in the other direction. The Cisco Catalyst platform provides features to guard against this condition. Unidirectional Link Detection (UDLD) and loop guard protect the network from anomalous conditions that result from unidirectional link conditions.

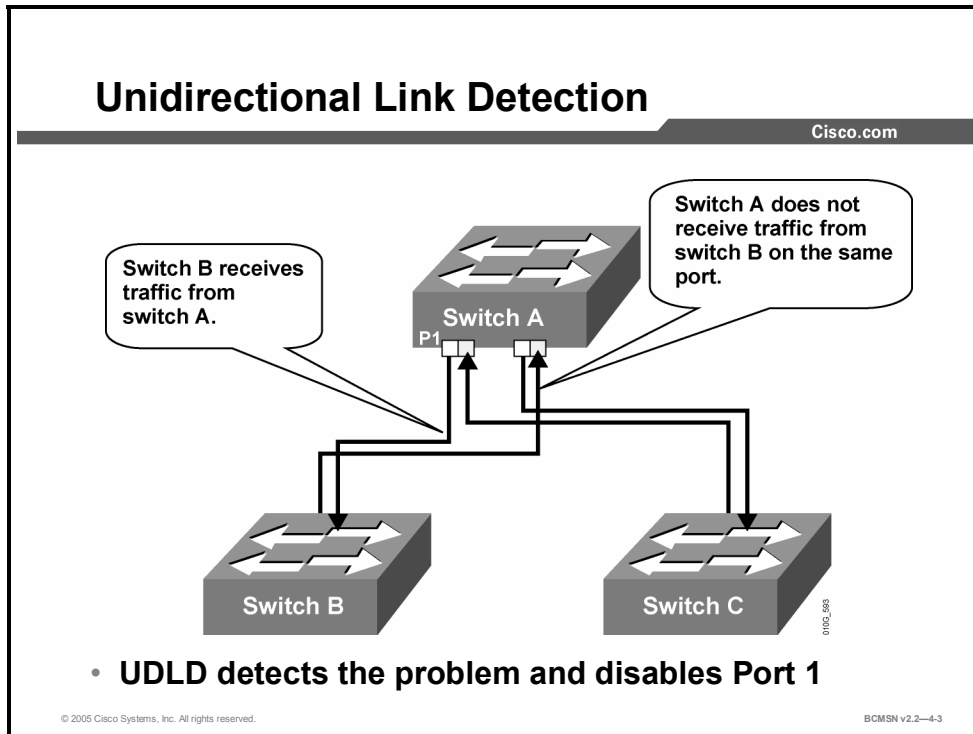
Objectives

Upon completing this lesson, you will be able to configure UDLD and loop guard to mitigate the adverse effects that unidirectional links have on spanning tree. This ability includes being able to meet these objectives:

- Describe how UDLD is used to detect and shut down unidirectional links
- Describe how loop guard is used to protect against Layer 2 forwarding loops
- Compare the features of loop guard and UDLD as they protect against unidirectional links
- Configure and verify UDLD and loop guard to protect STP from unidirectional link failures

Unidirectional Link Detection

This topic identifies the features associated with Unidirectional Link Detection (UDLD) protocol.



A unidirectional link occurs when traffic is transmitted between neighbors in one direction only. Unidirectional links can cause spanning tree topology loops. UDLD allows devices to detect when a unidirectional link exists and also to shut down the affected interface.

UDLD is a Layer 2 protocol that works with the Layer 1 mechanisms to determine the physical status of a link. If one fiber strand in a pair is disconnected, autonegotiation would not allow the link to become active or stay up. If both fiber strands are operant from a Layer 1 perspective, UDLD determines if traffic is flowing bi-directionally between the correct neighbors.

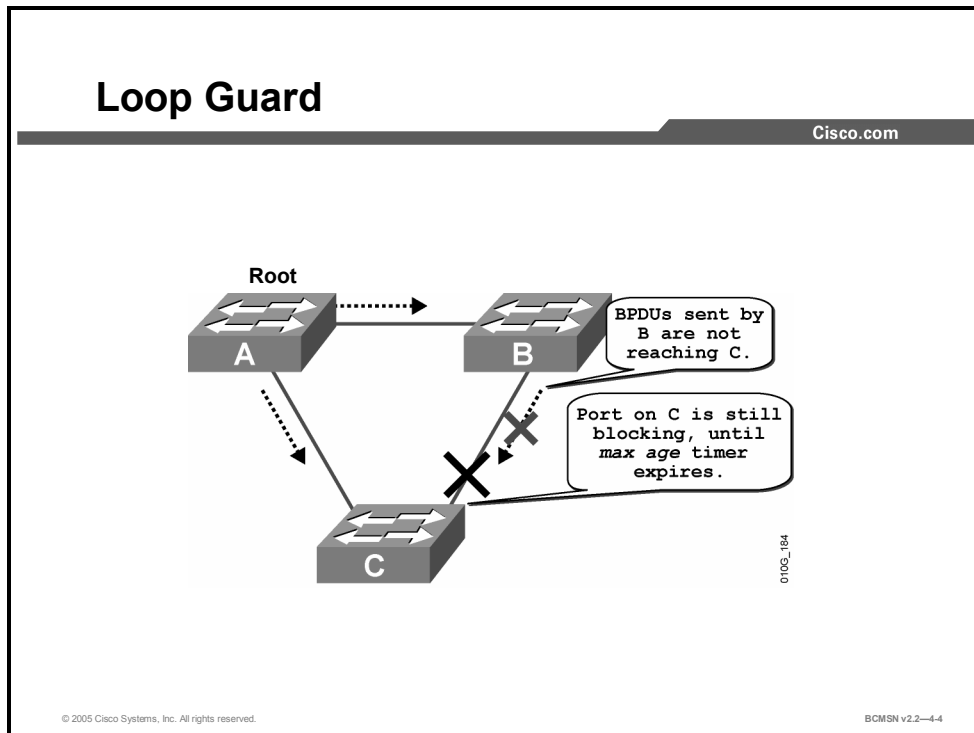
The switch periodically transmits UDLD packets on an interface with UDLD enabled. If the packets are not echoed back within a specific time frame, the link is flagged as unidirectional, and the interface is shut down. Devices on both ends of the link must support UDLD for the protocol to successfully identify and disable unidirectional links.

The table describes the default status for the UDLD on a global and an interface basis.

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

Loop Guard

This topic describes how loop guard protects against Layer 2 forwarding loops.



Like UDLD, loop guard provides protection for Spanning Tree Protocol (STP) when a link is unidirectional and bridge protocol data units (BPDUs) are being sent, but not received, on a link that is considered operational. Without loop guard, a blocking port will transition to forwarding if it stops receiving BPDUs. If loop guard is enabled and the link is not receiving BPDUs, the interface will move into the STP loop-inconsistent blocking state. When loop guard blocks a port, this message is generated to the console or log file if allowed:

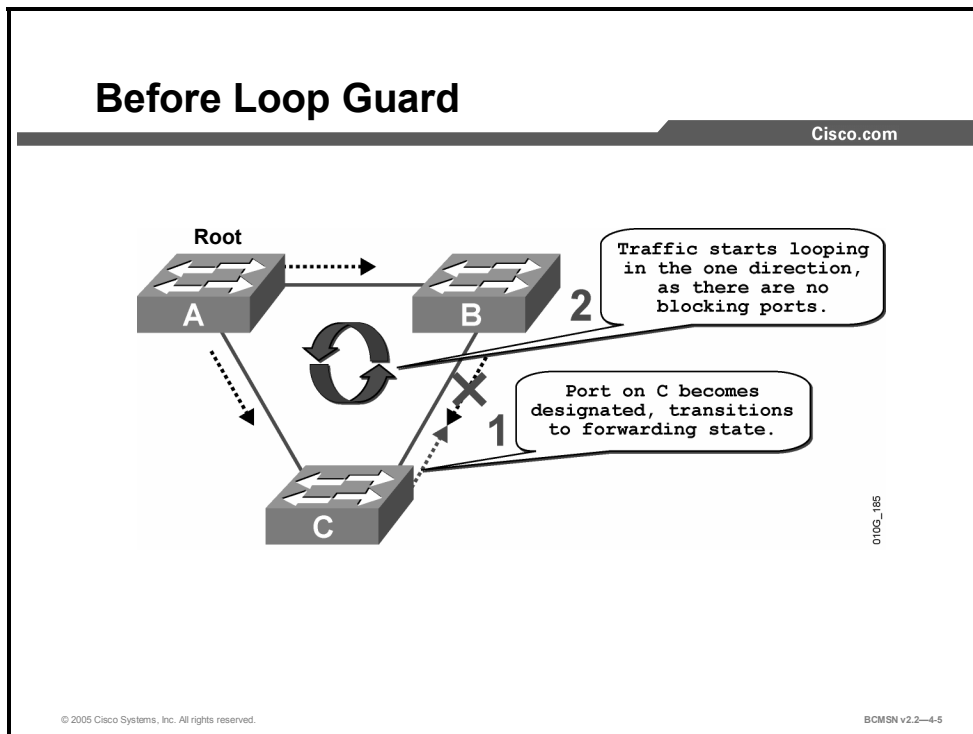
```
SPANTREE-2-LOOPGUARDBLOCK: No BPDUs were received on port 3/2 in vlan 3. Moved to loop-inconsistent state.
```

Once a BPDU is received on a loop guard port that is in a loop-inconsistent state, the port will transition to the appropriate state as determined by the normal functioning of spanning tree. The recovery requires no user intervention. After the recovery, this message is logged:

```
SPANTREE-2-LOOPGUARDUNBLOCK: port 3/2 restored in vlan 3.
```

Example: Before Loop Guard

This topic demonstrates how loops can occur due to unidirectional link failure.

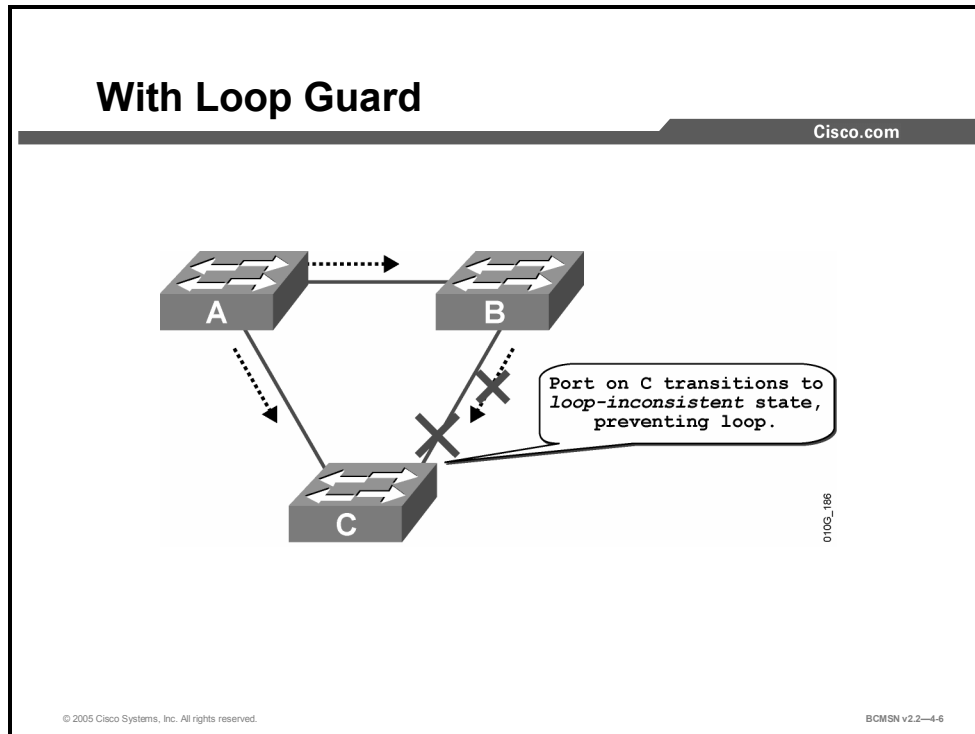


In this example, switch A is the root bridge. Due to unidirectional link failure on the link between switch B and switch C, switch C is not receiving BPDUs from B.

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.

Example: With Loop Guard

This example demonstrates how loop guard works to prevent loops during a unidirectional link failure.



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.

References

For additional information, refer to this resource:

http://www.cisco.com/en/US/partner/tech/tk389/tk621/technologies_tech_note09186a0080094640.shtml

How to Prevent STP Failures Due to Unidirectional Links

This topic discusses loop guard and UDLD features.

Comparing Loop Guard and UDLD		
<small>Cisco.com</small>		
	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 alternative 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

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-4-7

The functions of UDLD and loop guard 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 the way they function. The figure identifies the key differences.

Depending on various design considerations, you can choose either UDLD or loop guard. UDLD provides no protection against STP failures caused by software that result in the designated switch not sending BPDUs. This type of failure, however, is less common than those caused by hardware failure.

On an EtherChannel bundle, UDLD will disable individual failed links. The channel itself remains functional if other links are available. Loop guard will put the entire channel in loop-inconsistent state if any physical link in the bundle fails.

Loop guard does not work on shared links or a link that has been unidirectional since its initial setup. Enabling both UDLD and loop guard provides the highest level of protection.

Configuring UDLD and Loop Guard

This topic identifies the commands to configure UDLD and loop guard.

UDLD and Loop Guard Configuration Commands

Cisco.com

Configuring and verifying UDLD

- `udld enable`
- `show udld interface fa0/1`

Configuring and verifying loop guard

- `spantree global-default loopguard enable`
- `show spantree guard fa0/1`

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—4-8

To enable or disable UDLD and loop guard, use these commands.

Command	Description
Switch(config-if) # udld enable	Enables UDLD on fiber and non-fiber interfaces
Switch(config) # udld enable	Enables UDLD globally on all fiber-optic switch interfaces
Switch(config-if) # no udld enable	Disables UDLD on individual non-fiber-optic interfaces
Switch(config-if) # udld disable	Disables UDLD on individual fiber-optic interfaces
Switch# udld reset	Resets all interfaces that have been shut down by UDLD
Switch# show udld interface <i>type mod/port</i>	Verifies the UDLD configuration for an interface
Switch(config) # spantree global-default loopguard enable	Globally enables loop guard
Switch(config) # spantree global-default loopguard disable	Globally disables loop guard
Switch# show spantree guard <i>type mod/port</i> vlan	Verifies loop guard status

Configuring UDLD

This subtopic identifies the command options for configuring 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 non-fiber-optic interface**

`Switch(config-if)#udld disable`

- **Disables UDLD on an individual fiber-optic interface**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—4-9

UDLD is used when a link should be shut down because of a hardware failure that is causing unidirectional communication. In an EtherChannel bundle, UDLD will shut down only the physical link that has failed.

UDLD can be enabled globally for all fiber interfaces or on a per-interface basis.

Enable UDLD on an Interface

To enable UDLD on an interface use the following command:

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

Enable UDLD Globally

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

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


Verifying and Resetting UDLD

This subtopic identifies the command options for resetting UDLD and verifying UDLD configuration.

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

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—4-10

Interfaces will be shut down by UDLD. 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
```

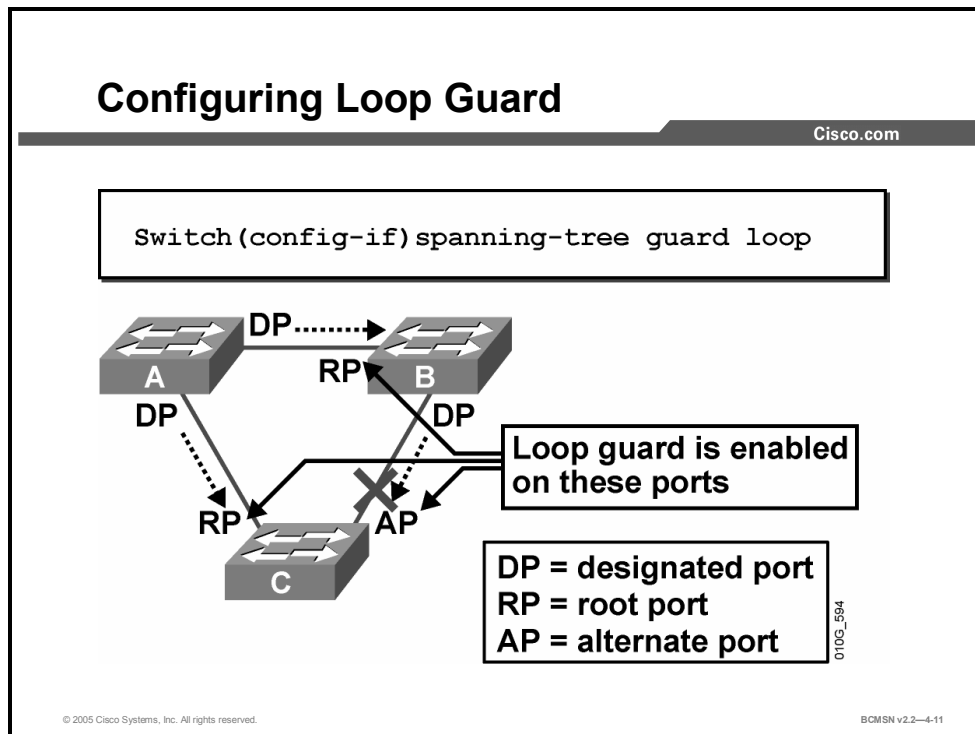
```
No multiple neighbors detected
```

```
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
```

Configuring Loop Guard

This subtopic identifies how to configure loop guard.



Loop guard is enabled on a per-port basis. When loop guard is enabled, it is automatically applied to all of the active VLAN instances to which that port belongs. When you disable loop guard, it is disabled for the specified ports. Disabling loop guard moves all loop-inconsistent ports to the listening state. 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 should be enabled on the root port and the alternative ports on access switches.

Enable Loop Guard on an Interface

To enable loop guard on a specific interface, issue this command:

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

To disable loop guard, issue this command:

```
Switch(config-if) # no spanning-tree guard loop
```

Enabling loop guard will disable root guard, if root guard is currently enabled on the ports.

Enable Loop Guard Globally

Loop guard can be enabled globally on a switch for all point-to-point links. A full-duplex link is considered to be a point-to-point link. The status of loop guard can be changed on an interface even if the feature has been enabled globally.

To enable loop guard globally, issue this command:

```
Switch(config)#spantree global-default loopguard enable
```

To globally disable loop guard, issue this command:

```
Switch(config)#spantree global-default loopguard disable
```

Verifying the Loop Guard Status

To verify the loop guard status, issue this command:

```
Switch#show spantree guard mod/port | vlan
```

For example,

```
Switch#show spantree guard 3/13
```

Port	VLAN	Port-State	Guard Type
3/13	2	forwarding	loop

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **UDLD detects and disables an interface with unidirectional connectivity, protecting the network from anomalous STP conditions.**
- **Loop guard detects and disables an interface with Layer 2 unidirectional connectivity, protecting the network from anomalous STP conditions.**
- **Implementation of UDLD and loop guard will protect spanning tree operations from being disrupted due to unidirectional links.**
- **UDLD and loop guard are configured and verified using specific commands.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—4-12

References

For additional information, refer to this resource:

- Cisco Systems, Inc., *Spanning-Tree Protocol Enhancements using loop guard and BPDU Skew Detection Features*,
http://www.cisco.com/en/US/partner/tech/tk389/tk621/technologies_tech_note09186a0080094640.shtml

Implementing RSTP

Overview

Rapid Spanning Tree Protocol (RSTP) is an improvement on the original 802.1D Spanning Tree Protocol (STP) standard. RSTP provides much faster convergence when topology changes occur in a switched network. Through the use of specific port states, port roles, and link types, RSTP very quickly adapts to network topology transitions. A proposal and agreement process between neighbor switches is unique to RSTP. Also, Topology Change Notifications (TCNs) are transferred in a very different manner than they are in 802.1D STP operation. Configuration of RSTP is much the same as in 802.1D except for a few variations and identifiable characteristics in the spanning tree verification commands.

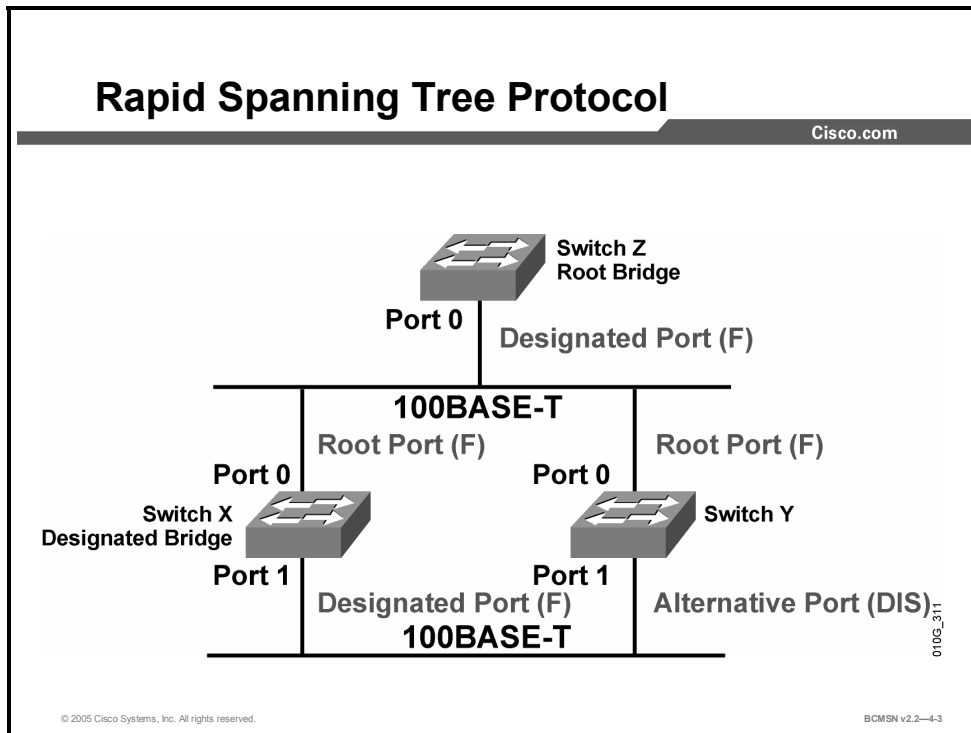
Objectives

Upon completing this lesson, you will be able to implement RSTP to increase the speed of STP recalculation when a Layer 2 topology change occurs. This ability includes being able to meet these objectives:

- Describe RSTP
- Identify the three RSTP port states
- Identify the five different RSTP port roles
- Define an edge port
- Identify the function of the different RSTP link types
- Distinguish the 802.1w use of the BPDU
- Identify the stages of the RSTP proposal and agreement process
- Identify the process RSTP uses to notify all switches of a topology change
- Identify the commands used to implement RSTP
- Implement RSTP in a switched network

Rapid Spanning Tree Protocol

This topic identifies the features of RSTP.



Rapid Spanning Tree Protocol (RSTP) speeds the recalculation of the spanning tree when the Layer 2 network topology changes. It is an IEEE standard that redefines STP port roles and states, and the bridge protocol data units (BPDUs).

RSTP is proactive and therefore negates the need for the 802.1D delay timers. RSTP (802.1w) supersedes 802.1D, while still remaining backward compatible. Much of the 802.1D terminology remains, and most parameters are unchanged. In addition, 802.1w is capable of reverting back to 802.1D to interoperate with legacy switches on a per-port basis.

In a switched domain, there can be only one forwarding path toward a single reference point; this is the root bridge. The RSTP spanning tree algorithm (STA) elects a root bridge in exactly the same way as 802.1D elects a root.

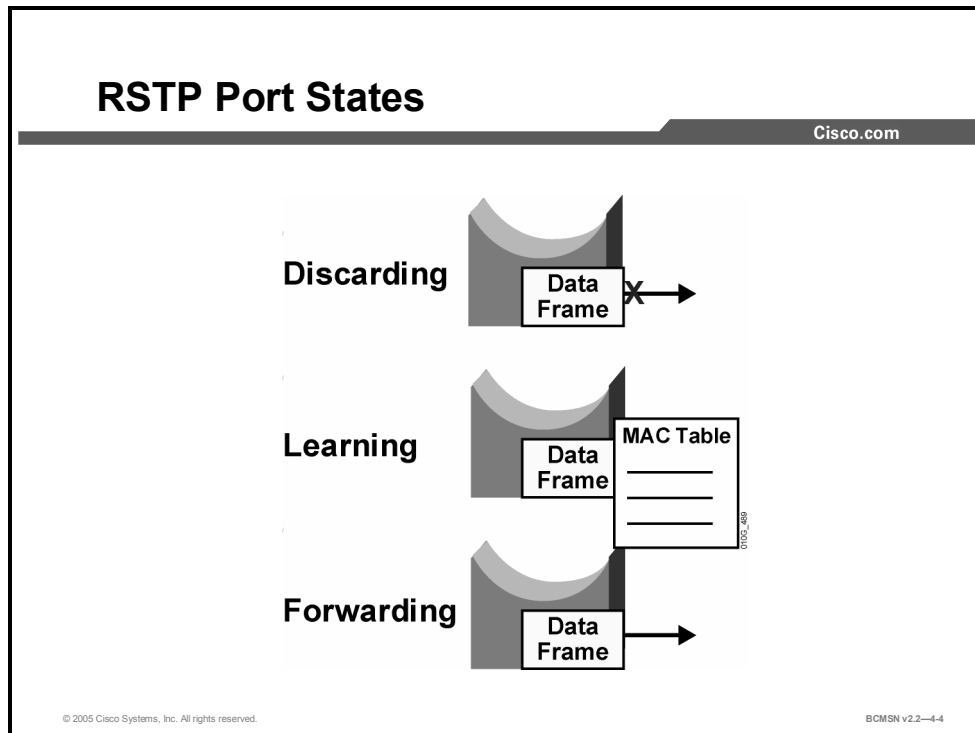
However, there are critical differences that make RSTP the preferred protocol for preventing Layer 2 loops in a switched network environment. Many of the differences stem from the Cisco proprietary enhancements. The Cisco-based RSTP enhancements have these characteristics:

- 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 802.1D enhancements.
- BPDUs carry information about port roles and is sent to neighbor switches only.

Because the RSTP and the Cisco-proprietary enhancements are functionally similar, features such as UplinkFast and BackboneFast are not compatible with RSTP.

RSTP Port States

This topic discusses RSTP port states and their appropriate functions.



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 either 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 could be in the discarding state temporarily, even though its final state is to be forwarding.

The RSTP port states correspond to the three basic operations of a switch port: discarding, learning, and forwarding.

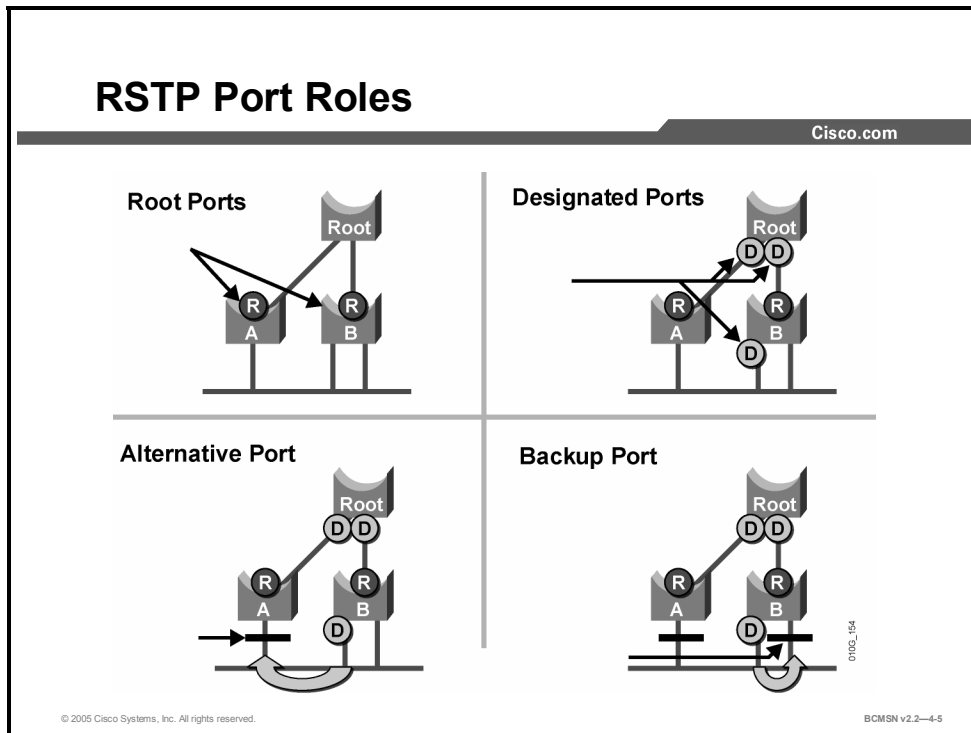
The port states have these characteristics:

- **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, thus “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 occurs only after a proposal and agreement process.

In all port states, a port will accept and process BPDU frames.

RSTP Port Roles

This topic discusses the RSTP port role functions.



The port role defines the ultimate purpose of a switch port and the way it handles data frames. Port roles and port states 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 bridge that is the chosen path to the root bridge. There can only be one root port on every switch. The root port assumes the forwarding state in a stable active topology.
- **Designated port:** Each segment will have at least one switch port that is the designated port for that segment. In a stable, active topology, the switch with the designated port will receive frames on the segment that are destined for the root bridge. There can only be one designated port per segment. The designated port assumes the forwarding state. All switches connected to a given segment listen to all BPDUs and determine the switch that will be the designated switch for a particular segment.
- **Alternative port:** This is a switch port that offers an alternative path toward the root bridge. The alternative port assumes a discarding state in a stable, active topology. An alternative 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.

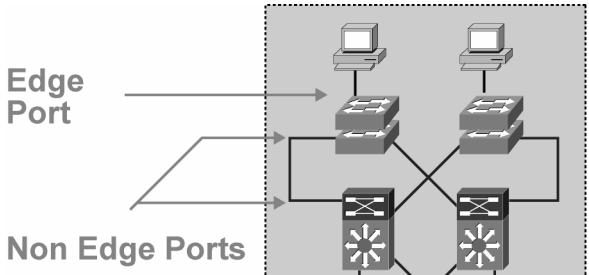
Establishing the additional port roles allows RSTP to define a standby switch port before a failure or topology change. The alternative port moves to the forwarding state if there is a failure on the designated port for the segment.

What Are Edge Ports?

This topic describes edge ports.

What Are Edge Ports?

Cisco.com



The diagram shows a network topology with four switches. Two switches at the top are connected to two desktop computers. Two switches at the bottom are connected to each other and to the top switches. Arrows point from the text 'Edge Port' to the ports on the top switches that are connected to the computers. Arrows point from the text 'Non Edge Ports' to the ports on the bottom switches that are connected to other switches.

- **Will never have a switch connected to it**
- **Immediately transitions to forwarding**
- **Functions similarly to PortFast**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—4-6

An RSTP edge port is a switch port that is never intended to be connected to another switch device. It immediately transitioned to the forwarding state when enabled.

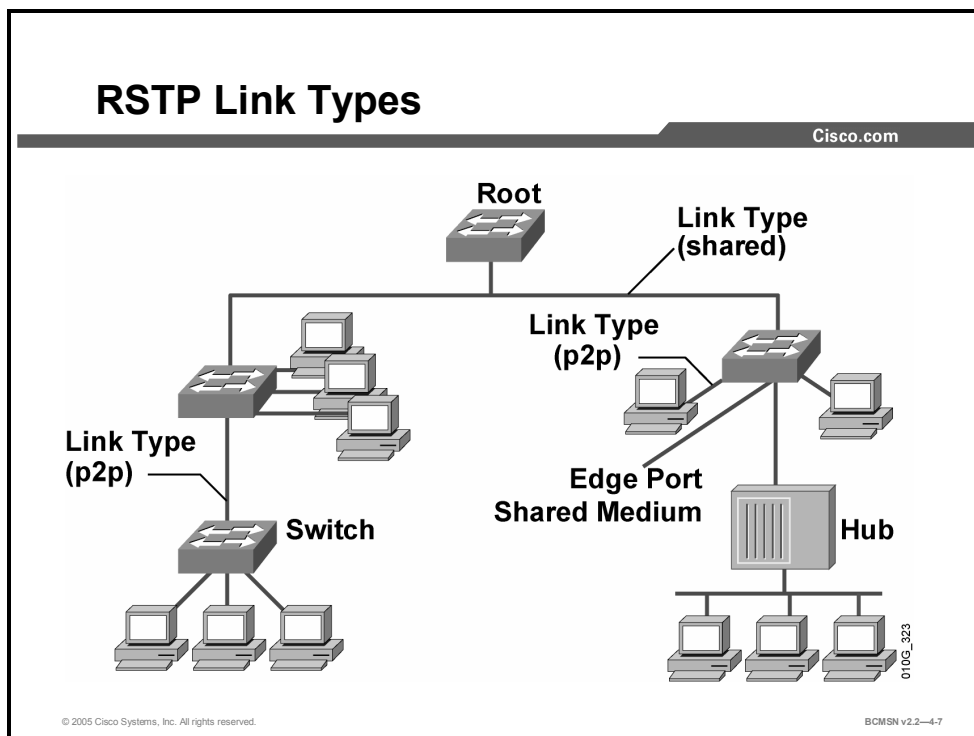
The edge port concept is well-known to Cisco spanning tree users as it corresponds to the PortFast feature. All ports directly connected to end stations anticipate that no switch device will be connected to them and immediately transition to the STP forwarding state, thereby skipping the time-consuming listening and learning stages. Neither edge ports nor PortFast-enabled ports generate topology changes when the port transitions to a disabled or enabled status.

Unlike PortFast, an edge port that receives a BPDU immediately loses its edge port status and becomes a normal spanning tree port. A switch with the edge port receiving a BPDU generates a topology change notification (TCN).

Cisco's RSTP implementation maintains the PortFast keyword for edge port configuration, thus making an overall network transition to RSTP more seamless. Configuring an edge port where the port will be attached to another switch can have negative implications for RSTP when it is in the "sync" state.

RSTP Link Types

This topic explains RSTP link types.



Link type provides a categorization for each port participating in RSTP. The link type can predetermine the active role that the port plays as it stands by for immediate transition to a forwarding state, if certain parameters are met. These parameters are different for edge ports and non-edge ports. Non-edge ports are categorized into two link types. Link type is automatically determined but can be overwritten with an explicit port configuration.

RSTP Link Types

Link Type	Description
Point-to-point	Port operating in full-duplex mode. It is assumed that the port is connected to a single switch device at the other end of the link.
Shared	Port operating in half-duplex mode. It is assumed that the port is connected to shared media where multiple switches might exist.

Edge ports, the equivalent of PortFast-enabled ports, and point-to-point links are candidates for rapid transition to a forwarding state. Before the link type parameter can be considered for the purpose of expedient port transition, RSTP must determine the port role.

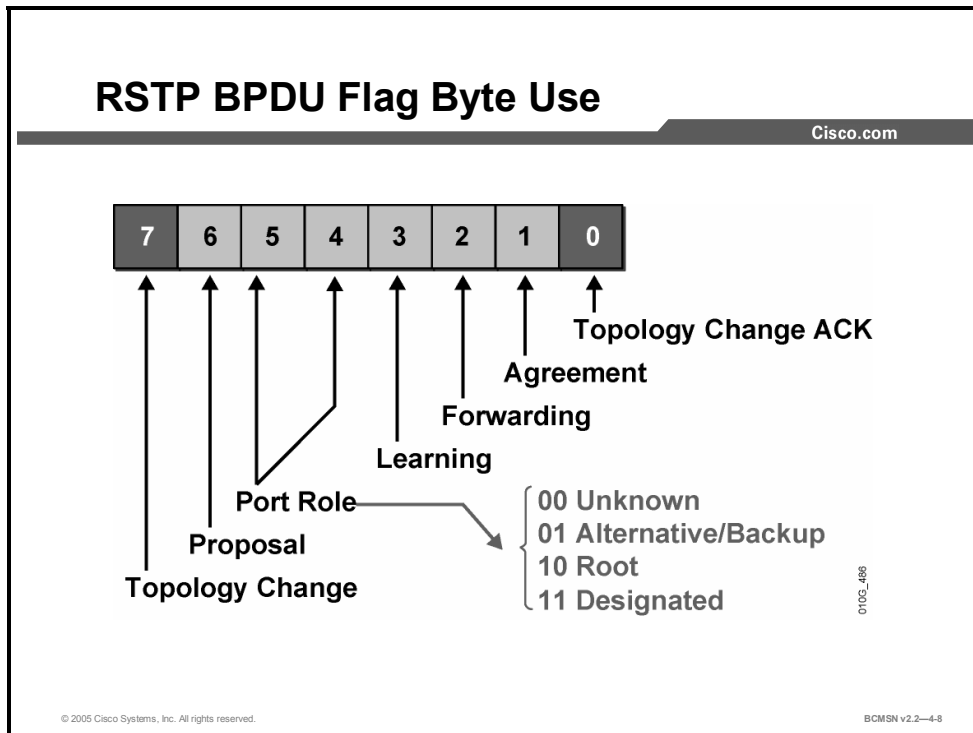
Root ports: Do not use the link type parameter. Root ports are able to make a rapid transition to the forwarding state as soon as the port is in “sync.”

Alternative and backup ports: Do not use the link type parameter in most cases.

Designated ports: Make the most use of the link type parameter. Rapid transition to the forwarding state for the designated port occurs only if the link type parameter indicates a point-to-point link.

Examining the RSTP BPDUs

This topic identifies the features of RSTP BPDUs.



RSTP (802.1w) uses type 2, version 2 BPDUs so an RSTP bridge can communicate with 802.1D on any shared link or with any switch running 802.1D. RSTP sends BPDUs and populates the flag byte in a slightly different manner than the manner used by 802.1D.

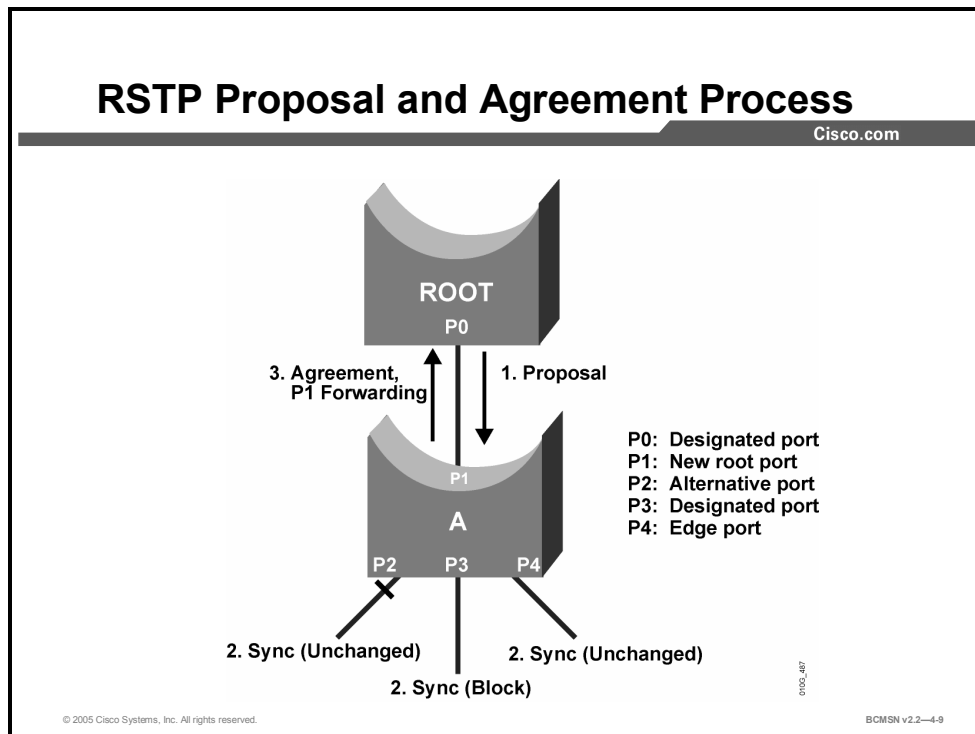
- An RSTP bridge sends a BPDUs with its current information every hello time period (2 seconds by default), even if it does not receive any BPDUs from the root bridge.
- Protocol information can be immediately aged on a port if hellos are not received for three consecutive hello times or if the max age 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.

RSTP uses the flag byte of version 2 BPDUs as shown in the figure.

- Bits 0 and 7 are used for topology change notification and acknowledgement (ACK), as they are in 802.1D.
- Bits 1 and 6 are used for the proposal agreement process.
- Bits 2–5 encode the role and state of the port originating the BPDUs.

Identifying the RSTP Proposal and Agreement Process

This topic identifies the process of proposal and agreement between RSTP switches.

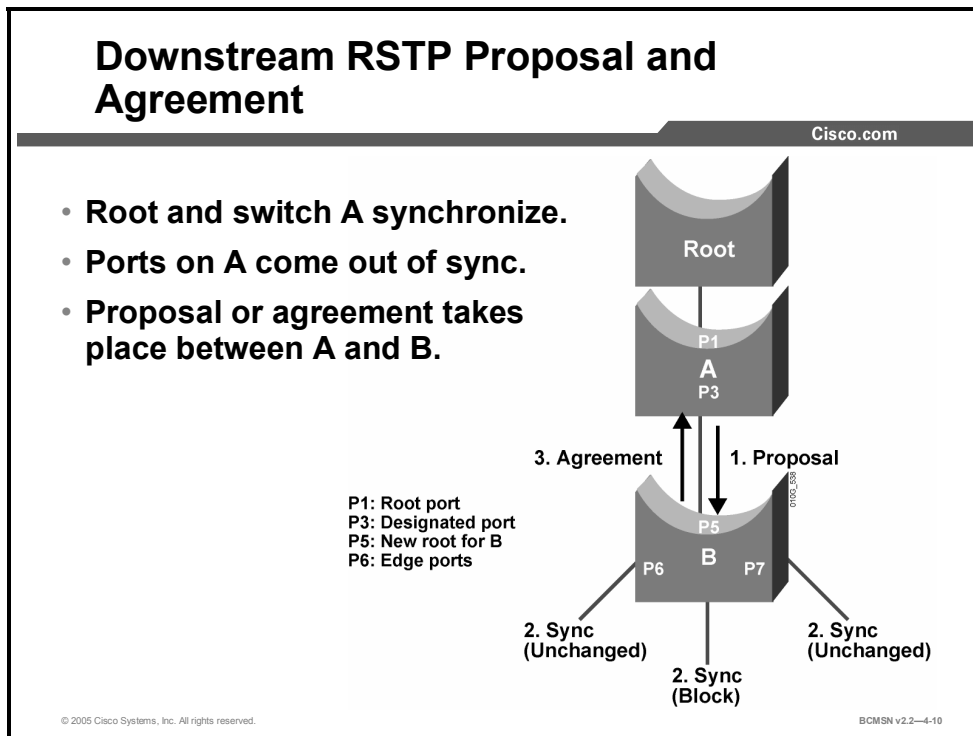


In 802.1D, when a port has been selected by spanning tree to become a designated port, it must wait two times the forward delay before transitioning the port to a forwarding state. RSTP significantly speeds up the recalculation process after a topology change in the network, as it converges on a link-by-link basis and does not rely on timers expiring before ports can transition. Rapid transition to forwarding state can only be achieved on edge ports and on point-to-point links. In RSTP, this condition corresponds to a port with a designated role that is in a blocking state. The figure illustrates, step by step, how rapid transition is achieved.

1. A new link is created between the root and switch A, and both ports are in designated blocking state until they receive a BPDU from their counterpart. When a designated port is in a discarding or learning state (and only in this case), it sets the proposal bit on the BPDUs it sends out. This is what happens for port P0 of the root bridge.
2. Switch A sees the proposal BPDU with a superior root ID. It blocks all non-edge designated ports other than the one over which the proposal and agreement process are occurring. This operation is called “sync” and prevents switches below A from causing a loop during the proposal-agreement process. Edge ports need not be blocked and remain unchanged during sync.
3. Bridge A explicitly sends an agreement that allows the root bridge to put the root port P0 in forwarding state. Port P1 becomes the root port for A.

Downstream RSTP Proposal Process

This subtopic discusses the steps in RSTP proposal acknowledgement downstream from the root bridge.

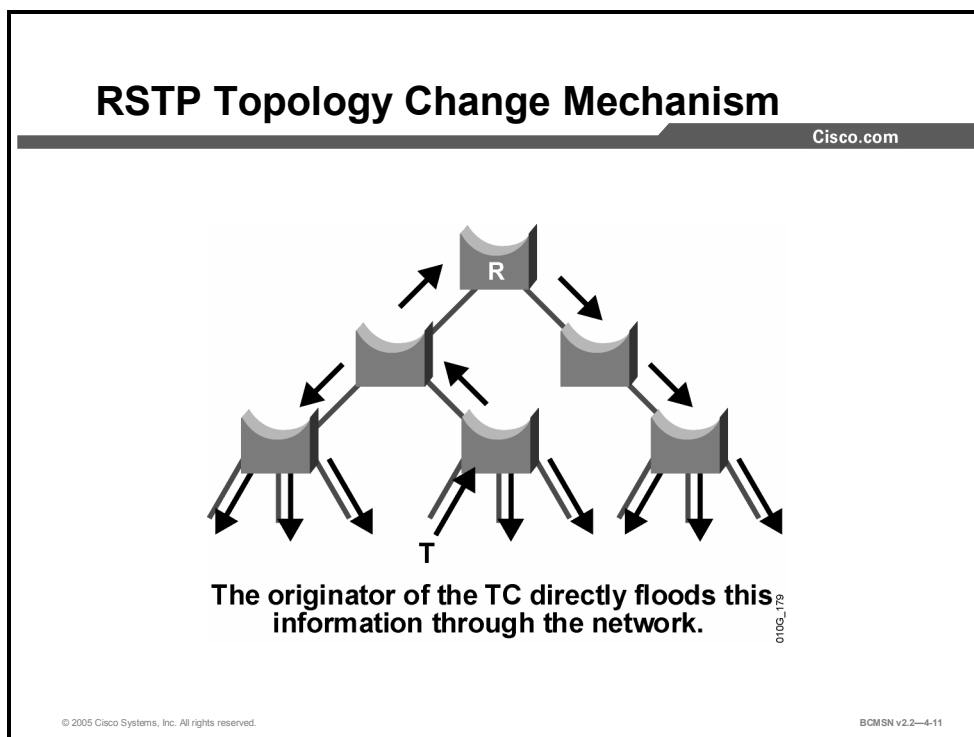


Once switch A and the root bridge are synchronized, the proposal or agreement process continues on switch A out of all of its downstream-designated, non-edge ports, as shown in the figure.

1. Switch B on P5 will see that switch A is discarding and will also transition to the designated discarding state. Switch A then sends its proposal BPDUs down to B with the root ID of the root bridge.
2. Switch B sees a proposal with the superior BPDUs from A and blocks all non-edge-designated ports other than the one over which the proposal and agreement process is occurring.
3. Switch B sends a BPDUs with the agreement bit set, and switch A P3 transitions to forwarding state. The synchronization process continues with switches downstream from B.

Identifying the RSTP Topology Change Notification Process

This topic identifies the actions that apply to notifying topology changes with RSTP in the correct order.



In 802.1D, any port state change generates a TCN. When an 802.1D bridge detects a topology change (TC), it sends TCNs toward the root bridge. The root bridge sets the TC flag on the outbound BPDUs that are relayed to switches down from the root. When a bridge receives a BPDU with the TC flag bit set, the bridge reduces its bridge-table aging time to forward delay seconds. This ensures a relatively quick flushing of the MAC address table.

In RSTP, only non-edge 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 BPDU.

When an RSTP bridge detects a TC, it performs these actions.

Step	Action	Notes
1.	The RSTP bridge starts the TC-While timer.	RSTP sets the TC-While timer with a value equal to twice the hello time for all its non-edge designated ports and the root port, if necessary.
2.	The RSTP bridge flushes the MAC addresses associated with all these ports.	
3.	The TC flag bit is set on all outbound BPDUs.	BPDUs are sent on the root port as long as the While timer is active.

Step	Action	Notes
4.	The bridge receives a BPDU with the TC bit set from a neighbor and clears the MAC addresses on all ports.	The port that received the TC BPDU retains learned MAC addresses.
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 TCN is flooded across the entire network, one switch at a time, from the switch that is the source of the change rather than from the root bridge. 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 the current operating mode for two periods of hello time, the port switches to the mode indicated by the BPDUs.

RSTP Implementation Commands

This topic describes commands for configuring and verifying Rapid Spanning Tree Protocol.

RSTP Implementation Commands

Cisco.com

Configuring

- `spanning-tree mode rapid-pvst`

Verifying

- `show spanning-tree vlan 101`

Debugging

- `debug spanning-tree`

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—4-12

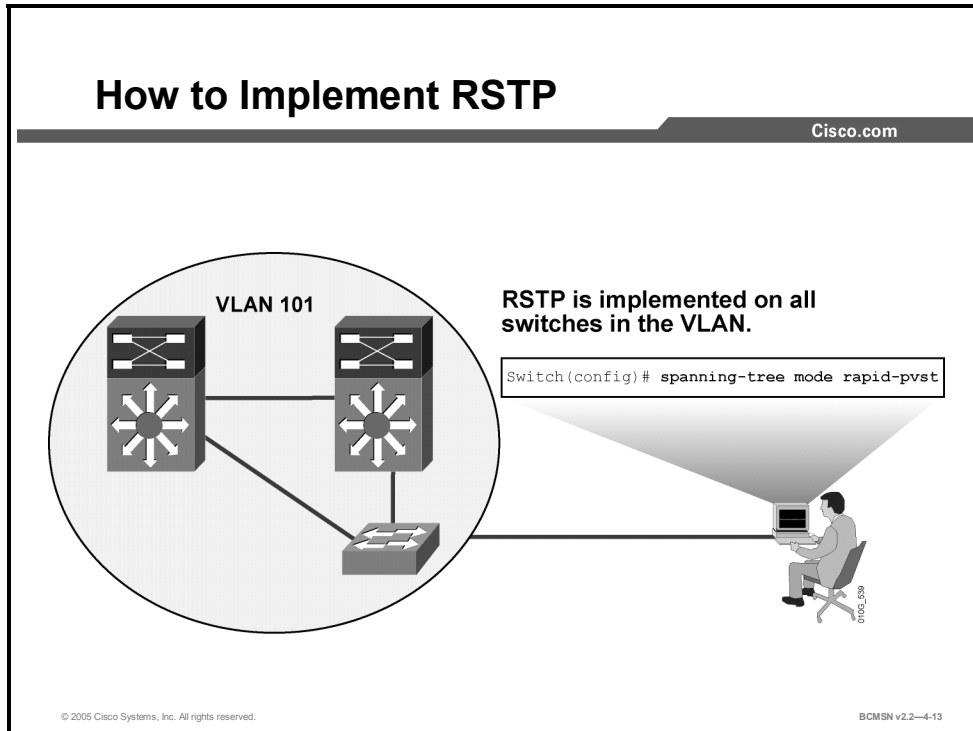
The following commands enable spanning tree per VLAN, and enable rapid PVST+:

RSTP Commands

Command	Description
Switch(config)# <code>spanning-tree mode rapid-pvst</code>	Sets spanning tree mode to Rapid PVST
Switch# <code>show spanning-tree vlan vlan-number [detail]</code>	Shows commands are VLAN based rather than instance based
Switch# <code>debug spanning-tree pvst+</code>	Debugs PVST events
Switch# <code>debug spanning-tree switch state</code>	Debugs port state changes

How to Implement RSTP

This topic shows how to implement and verify RSTP operations.



Configuring RSTP

Step	Description
1.	If spanning tree is disabled, enable it for a VLAN. <code>Switch(config)# spanning-tree vlan <i>vlan-range</i></code>
2.	Set spanning tree mode to Rapid-PVST+. Default is 802.1D (shows as "ieee"). <code>Switch(config)# spanning-tree mode rapid-pvst</code>

Explanation: Enabling PVST

Spanning tree is enabled 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.

Verifying the Rapid PVST Configuration


This subtopic identifies the commands that verify a rapid spanning tree configuration for a VLAN.

Verifying PVST

Cisco.com

```
Switch# show spanning-tree vlan 30
VLAN0030
Spanning tree enabled protocol rstp
Root ID Priority 24606
Address 00d0.047b.2800
This bridge is the root
Hello Time 2 sec Max Age 20 sec Forward Delay 15
sec
Bridge ID Priority 24606 (priority 24576 sys-id-ext
30)
Address 00d0.047b.2800
Hello Time 2 sec Max Age 20 sec Forward Delay 15
sec
Aging Time 300
Interface Role Sts Cost Prio.Nbr Type
-----
Gi1/1 Desg FWD 4 128.1 P2p
Gi1/2 Desg FWD 4 128.2 P2p
Gi5/1 Desg FWD 4 128.257 P2p
```

```
Switch# show spanning-tree vlan 30
```



- Display spanning tree mode is set to rapid PVST.

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-4-14

A variety of **show** commands can be used 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.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Rapid STP provides faster convergence than 802.1D STP when topology changes occur.**
- **RSTP quickly adapts to topology transitions using specific link types, port states, and roles.**
- **Edge ports forward while topology changes occur.**
- **802.1w uses the BPDU differently from 802.1D.**
- **Convergence results from the proposal agreement process conducted switch by switch.**
- **The RSTP topology change notification process differs greatly from 802.1D.**
- **Various commands are used to configure and verify RSTP.**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—4-15

Implementing MST

Overview

Per VLAN Spanning Tree (PVST) creates a single instance of spanning tree for each VLAN in the network. This may impose a processing load on a switch when many VLANs are present. Multiple Spanning Tree (MST) reduces this loading by allowing a single instance of spanning tree to run for multiple VLANs. Specific configuration and verification steps must be followed to properly implement MST.

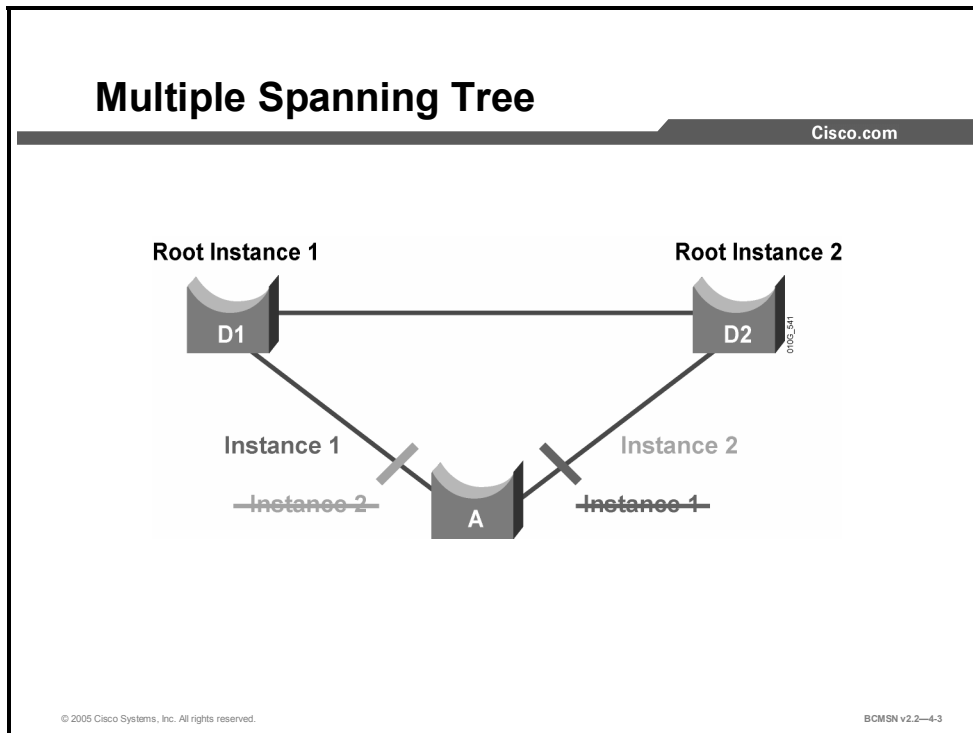
Objectives

Upon completing this lesson, you will be able to implement MST to reduce the number of Spanning Tree Protocol (STP) instances that could occur when running PVST. This ability includes being able to meet these objectives:

- Describe MST
- Identify the characteristics of an MST region
- Identify changes to the Bridge Priority field to accommodate MST instance number
- Identify how MST operates with CST
- Identify the commands used to implement MST
- Implement MST in a switched network

What Is MST?

This topic identifies the features of Multiple Spanning Tree.



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 loading of a switch. The instances of spanning tree are reduced to the number of links (that is, active paths) that are available. If the example in the diagram were implemented via PVST plus (PVST+), there could potentially be 4094 instances of spanning tree, each with its own bridge protocol data unit (BPDU) conversations, root bridge election, and path selections.

In this example, the goal is to achieve load distribution with VLANs 1-500 using one path and with VLANs 501-1000 using the other path, with only two instances of spanning tree. The two ranges of VLANs are mapped to two MST instances (MSTIs), respectively. Rather than maintaining 1000 spanning trees, each switch needs to maintain only two. Implemented in this fashion, MST converges faster than PVST+ and is backward compatible with 802.1D STP, 802.1w (RSTP), and the Cisco PVST+ architecture. Implementation of MST is not required if the Enterprise Composite Network Model is being employed, as the number of active VLAN instances, and hence the STP instances, would be small and very stable due to the design.

MST allows you to build multiple spanning trees over trunks by grouping VLANs and associating them with spanning tree instances. Each instance can have a topology independent of other spanning tree instances. This architecture provides multiple active forwarding paths for data traffic and enables load balancing. Network fault tolerance is improved over Common Spanning Tree (CST) because a failure in one instance (forwarding path) does not necessarily affect other instances. This VLAN-to-MST grouping must be consistent across all bridges within an MST region.

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. Bridges with different MST configurations or legacy bridges running 802.1D are considered separate MST regions.

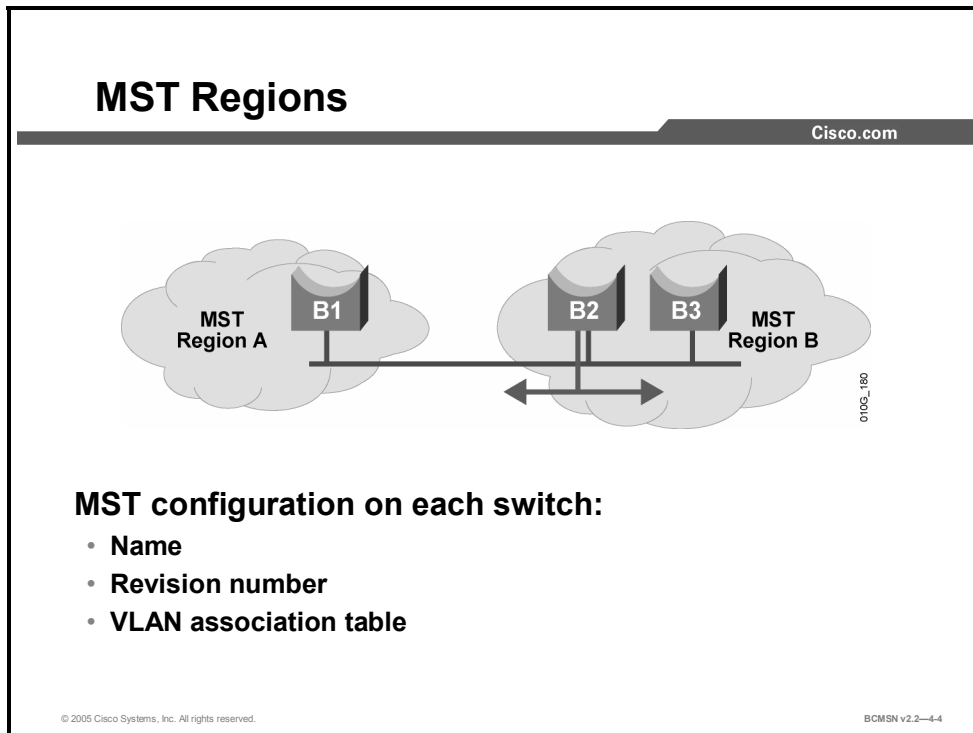
In a Cisco PVST+ environment, the spanning tree parameters are tuned so that half of the VLANs are forwarding on each uplink trunk. This is easily achieved by electing bridge D1 to be the root for VLAN501–1000, and bridge D2 to be the root for VLAN1–500. In this configuration, the following is true:

- Optimum load balancing is achieved.
- 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 (IEEE 802.1s) combines the best aspects of both PVST+ and 802.1D. The idea is that several VLANs can be mapped to a reduced number of spanning tree instances because most networks do not need more than a few logical topologies.

MST Regions

This topic identifies the characteristics of the MST region.



MST differs from the other spanning tree implementations in that it combines some, but not necessarily all, VLANs into logical spanning tree instances. This raises the problem of determining what VLAN is to be associated with what instance. More precisely, this means tagging BPDUs so that receiving devices can identify the instances and the VLANs to which they apply.

The issue is irrelevant in the case of the 802.1D standard, in which all instances are mapped to a unique and common instance CST. In the PVST+ implementation, different VLANs carry the BPDUs for their respective instances (one BPDU per VLAN), based on the VLAN tagging information.

To provide this logical assignment of VLANs to spanning trees, each switch running MST in the network has a single MST configuration that consists of three attributes:

- An alphanumeric configuration name (32 bytes)
- A configuration revision number (two bytes)
- A 4096-element table that associates each of the potential 4096 VLANs supported on the chassis with a given instance

To be part of a common MST region, a group of switches must share the same configuration attributes. It is up to the network administrator to properly propagate the configuration throughout the region.

Note If two switches differ on one or more configuration attributes, they are part of different regions.

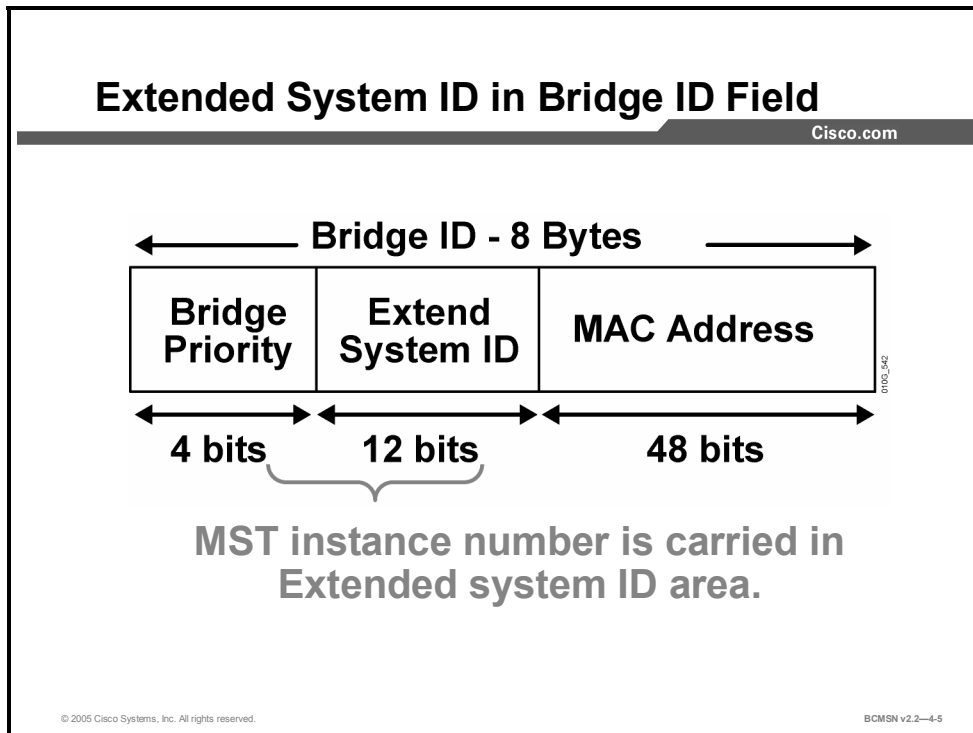
To ensure a consistent VLAN-to-instance mapping, it is necessary for the protocol to be able to exactly identify the boundaries of the regions. For that purpose, the characteristics of the region are included in BPDUs. The exact VLANs-to-instance mapping is not propagated in the BPDU, because the switches only need to know whether they are in the same region as a neighbor.

Therefore, only a digest of the VLANs-to-instance mapping table is sent, along with the revision number and the name. Once a switch receives a BPDU, it extracts the digest (a numerical value derived from the VLAN-to-instance mapping table through a mathematical function) and compares it with its own computed digest. If the digests differ, the mapping must be different, so the port on which the BPDU was received is at the boundary of a region.

In generic terms, a port is at the boundary of a region if the designated bridge on its segment is in a different region or if it receives legacy 802.1D BPDUs. In the figure, the port on B1 is at the boundary of region A, whereas the ports on B2 and B3 are internal to region B.

Extended System ID

This topic discusses extended system ID in the bridge ID (BID).



As with PVST, the 12-bit Extended System ID field is used in MST. In MST, this field carries the MST instance number.

References

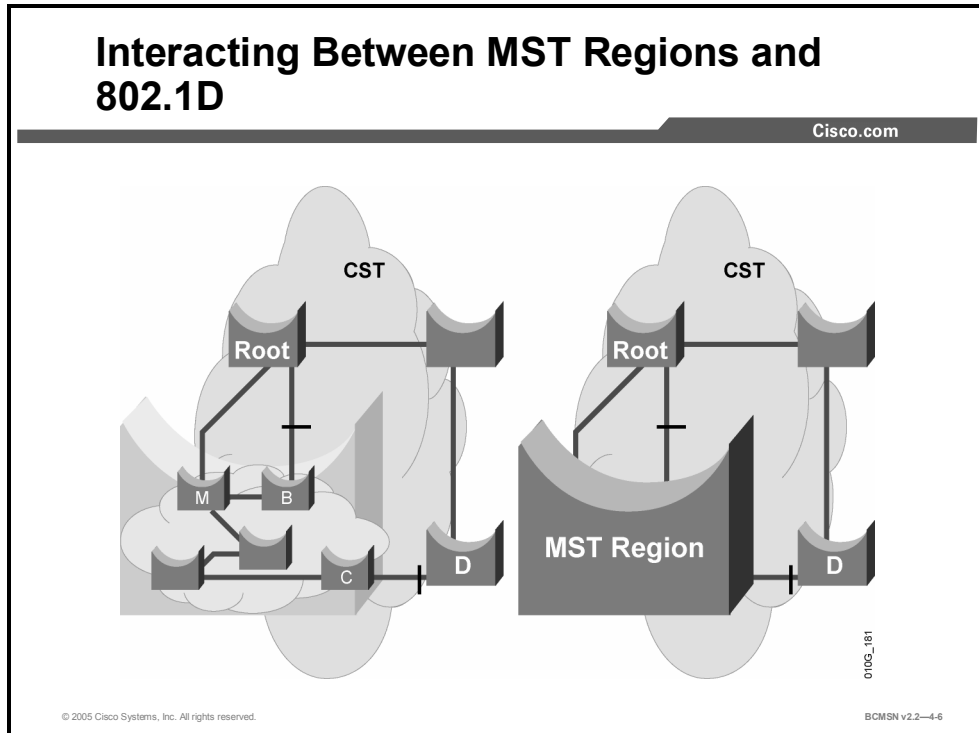
For additional information, refer to this resource:

Cisco Systems, Inc., *Configuring Spanning Tree*:

http://www.cisco.com/en/US/partner/products/hw/switches/ps663/products_configuration_guide_chapter09186a00801162ff.html#1157094

Interacting Between MST Regions and 802.1D Networks

This topic identifies the characteristics of the MST-CST interaction.



One issue that arises from MST design is interoperability with the CST implementation in 802.1D. According to the IEEE 802.1s specification, an MST switch must be able to handle at least one Internal Spanning Tree (IST).

The graphic shows how an MST region comprising multiple switches is viewed as a single entity externally by a network running a single CST. The IST (instance 0) runs on all bridges within an MST region. An important characteristic of the IST instance is that it provides interaction at the boundary of the MST region with other MST regions, and, more important, it is responsible for providing compatibility between the MST regions and the spanning tree of 802.1D (CST) and PVST+ networks connected to the region.

The IST instance receives and sends BPDUs to the CST for compatibility with 802.1D. 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 domains.
- 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, up 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.

- The M-Record is a subfield, within the BPDU of MSTIs, that contains enough information (root bridge and sender bridge priority parameters) for the corresponding instance to calculate the final topology. It does not contain any timer-related parameters (such as hello time, forward delay, and max age) that are typically found in a regular IEEE 802.1D BPDU, as these timers are derived from the IST BPDU timers. It is important to note that within an MST region, all spanning tree instances use the same parameters as the IST.
- 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.
 - For PVLANS, you must map a secondary VLAN to the same instance as the primary.

Each switch in the network that runs MST has a single MST configuration consisting of the following attributes:

- An alphanumeric configuration name (32 bytes)
- A configuration revision number (two bytes)
- A 4096-element table that associates each of the potential 4096 VLANs supported on the chassis to a given instance

To be part of a common MST region, a group of switches must share the same configuration attributes. It is up to the network administrator to properly propagate the configuration throughout the region.

MST Implementation Commands

This topic identifies the commands used to implement MST.

Configuring Multiple Spanning Tree

Cisco.com

```
Switch(config)#spanning-tree mst configuration
```

- Enters MST configuration submenu

```
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

```
Switch(config-mst)#spanning-tree mst instance_number root
primary|secondary
```

- Establishes primary and secondary roots for MST instance

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-4-7

All switches would be configured with the spanning tree MST and extend system-id syntax, and only the distribution switches that terminate the VLANs would have their priority changed.

Step	Description	Notes and Comments
1.	Enter MST configuration submenu. 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	

Step	Description	Notes and Comments
	configuration submode. Switch(config-mst) # end	
8.	Assign root bridge for MST instance. Switch(config-mst) # spanning-tree mst instance_number root primary secondary	This syntax makes switch root primary or secondary (only active if primary fails). It sets primary priority to 24576 and secondary to 28672.
9.	Switch(config) # spanning-tree extend system-id	This enables MAC address reduction, also known as extended system ID in Cisco IOS software.

How to Configure and Verify MST

This topic identifies the commands used to verify MST.

Verifying MST

Cisco.com

```
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
-----
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-4-8

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. This includes MST region name, revision number, and VLAN-to-MST instances mapping.

```
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. Notice that the output is grouped by MSTIs, starting with the IST.

```
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)

Example: Displaying MST Information for a Specific Instance

This example displays spanning tree information for a specific MSTI—particularly port status, costs, and forwarding role.

Verifying MST (Cont.)

Cisco.com

```
Switch#show spanning-tree mst instance_number
```

- **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)
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-4-9

```
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

```

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Multiple Spanning Tree reduces the encumbrance of PVST by allowing a single instance of spanning tree to run for multiple VLANs.**
- **An MST region is a group of MST switches that appears as a single virtual bridge to adjacent CST and MST regions.**
- **Extended system ID ensures that VLAN ID or MST instance can be carried in the Bridge ID field of a BPDU.**
- **An MST region requires an IST and an arbitrary number of MSTIs as it connects to an 802.1Q network at the MST region border.**
- **MST is configured with a unique set of commands.**
- **MST implementation requires configuration and verification using specific configuration and show commands.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—4-10

References

For additional information, refer to this resource:

Cisco Systems, Inc., Catalyst 4500 Series Software Configuration Guide, 7.4: *Configuring Spanning Tree*.

http://www.cisco.com/en/US/partner/products/hw/switches/ps663/products_configuration_guide_chapter09186a00801162ff.html#1157094

Module Summary

This topic summarizes the key points discussed in this module.

Module Summary

Cisco.com

- **Properly identifying STP problems and taking appropriate action is critical to maintaining reliable spanning tree operation.**
- **Unidirectional link detection and loop guard protect the network from anomalous STP conditions that result from unidirectional links.**
- **Through the use of specific port states, port roles, and link types, RSTP quickly adapts to network topology transitions.**
- **Multiple Spanning Tree reduces the burden of excessive STP traffic and CPU processing by allowing a single instance of spanning tree to represent multiple VLANs.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-4-1

Specific steps should be taken to properly troubleshoot STP. STP forwarding loops can be detected and prevented through specific **verification** and **configuration** commands. Many enhancements have been made to the original 802.1D STP protocol. A switched network can quickly adapt to topology changes by implementing Rapid Spanning Tree Protocol. Multiple Spanning Tree implements a minimal number of STP instances in a switched environment.

References

For additional information, refer to these resources:

- Cisco Systems, Inc., *Spanning Tree Protocol Problems and Related Design Considerations*:
http://www.cisco.com/en/US/partner/tech/tk389/tk621/technologies_tech_note09186a00800951ac.shtml
- Cisco Systems, Inc., *Spanning Tree Protocol Enhancements using Loop Guard and BPDU Skew Detection Features*:
http://www.cisco.com/en/US/partner/tech/tk389/tk621/technologies_tech_note09186a0080094640.shtml
- Cisco Systems, Inc., *Catalyst 4500 Series Software Configuration Guide, 7.4: Configuring Spanning Tree*:
http://www.cisco.com/en/US/partner/products/hw/switches/ps663/products_configuration_guide_chapter09186a00801162ff.html#1157094

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which two features apply to loop guard? (Choose two.) (Source: Preventing STP Forwarding Loops)
- A) provides additional protection against Layer 2 STP loops
 - B) allows a blocked 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 state if BPDUs are not received on a nondesignated port
- Q2) Which two features apply to MST? (Choose two.) (Source: Implementing MST)
- A) It can group a set of VLANs to a single spanning tree instance.
 - B) It groups a set of instances to a single VLAN.
 - C) A failure in one instance can cause a failure in another instance.
 - D) The total number of spanning tree instances should match the number of redundant switch paths.
- Q3) What is the purpose of MST? (Source: “Implementing MST”)
- A) to increase the number of STP instances
 - B) to decrease the number of STP instances
 - C) to provide for load balancing
 - D) to be able to maintain two times the typical number of STP instances
- Q4) Which command will enable spanning tree debugging messages for event detail? (Source: “Troubleshooting Spanning Tree”)
- A) **debug spanning-tree all**
 - B) **debug spanning-tree issues**
 - C) **debug spanning-tree events**
 - D) **debug stp events**
- Q5) Which two of the following are valid port states for RSTP? (Choose two.) (“Implementing RSTP”)
- A) alternative port
 - B) backup port
 - C) listening port
 - D) learning port

Module Self-Check Answer Key

Q1) A, D

Q2) A, D

Q3) B

Q4) C

Q5) A, B

Implementing Multilayer Switching

Overview

A switch with multiple VLANs requires a means of passing Layer 3 traffic between those VLANs. This module describes both the process and various methods of routing traffic from VLAN to VLAN. A router that is external to the Layer 2 switch hosting the VLANs can provide the inter-VLAN routing. When routing occurs within a Catalyst multilayer switch, Cisco Express Forwarding (CEF) is deployed to facilitate Layer 3 switching through hardware-based tables, providing an optimal packet forwarding process. When CEF is implemented, routing is enabled between VLANs through the configuration of Switch Virtual Interfaces associated with the various VLANs on the multilayer switch.

Module Objectives

Upon completing this module, you will be able to implement multilayer switching, thereby achieving higher data throughput speeds when routing between VLANs. This ability includes being able to meet these objectives:

- Configure routing between VLANs on a router that is external to a switch
- Deploy CEF-based multilayer switching so that all packets are CEF switched in hardware
- Enable routing between VLANs on a multilayer switch

Describing Routing Between VLANs

Overview

Layer 2 switching involves processing frames with respect to their data link layer headers. Information from those headers is stored within the content addressable memory (CAM) table in the switch, which in turn provides the information required to make the forwarding decisions as frames traverse the switch. When multiple Layer 2 VLANs are configured on a switch, a Layer 3 process is required for inter-VLAN communication. VLAN-to-VLAN packet transfer can occur on a Layer 3 device external to the switch.

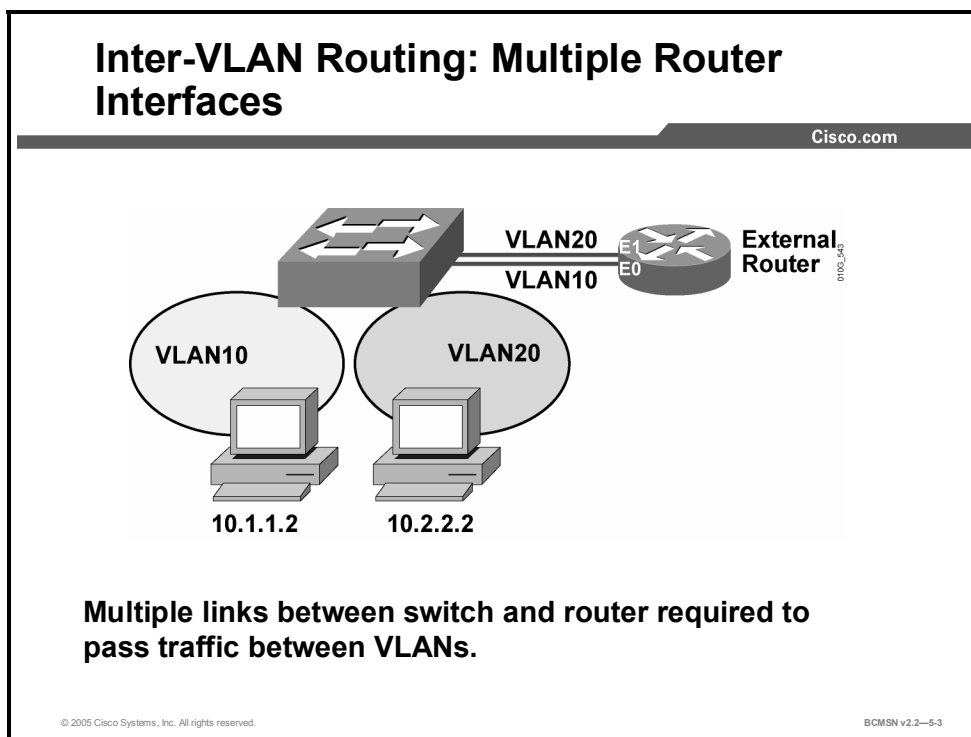
Objectives

Upon completing this lesson, you will be able to how to configure routing between VLANs on a router that is external to a single switch hosting multiple VLANs. This ability includes being able to meet these objectives:

- Identify the process of routing between VLANs on an external router using multiple interfaces
- Identify the process of routing between VLANs on an external router using a single trunk
- Configure and verify inter-VLAN routing using an external router using ISL or 802.1Q

Inter-VLAN Routing Using Multiple Interfaces on an External Router

This topic describes how to use an external router with multiple interfaces to route traffic between VLANs.



On a single VLAN on a single switch, all frames are processed and passed between ports based upon inspection of Layer 2 headers. If packets must be passed between VLANs, this requires packet forwarding based upon inspection of the Layer 3 header, specifically the destination IP address.

In a multilayer switch, Layer 2 and Layer 3 processing occurs on the same device. However, if a switch supports multiple VLANs but has no Layer 3 capability to route packets between those VLANs, the switch must be connected to a router external to the switch. This can be accomplished by setting up one router interface for each VLAN on the switch. The router then provides all Layer 3 routing functionality between VLANs on the switch.

In the figure, the clients on VLAN10 need to establish sessions with a server that is in VLAN20. This will require that traffic be routed between the VLANs, as described in this table.

Step	Action
1.	The router accepts packets from VLAN10 on E0, which hosts the IP address for the default gateway for all hosts on VLAN10.
2.	The router performs Layer 3 processing, inspecting the destination IP address, and the routing table shows that subnet is directly connected on E1.
3.	Packets destined for all IP addresses on VLAN20 are sent out E1.

External Router with Multiple Interface: Advantages and Disadvantages

This subtopic describes advantages and disadvantages of inter-VLAN routing on an external router using multiple interfaces.

Inter-VLAN Routing: Multiple Router Interfaces (Cont.)

Cisco.com

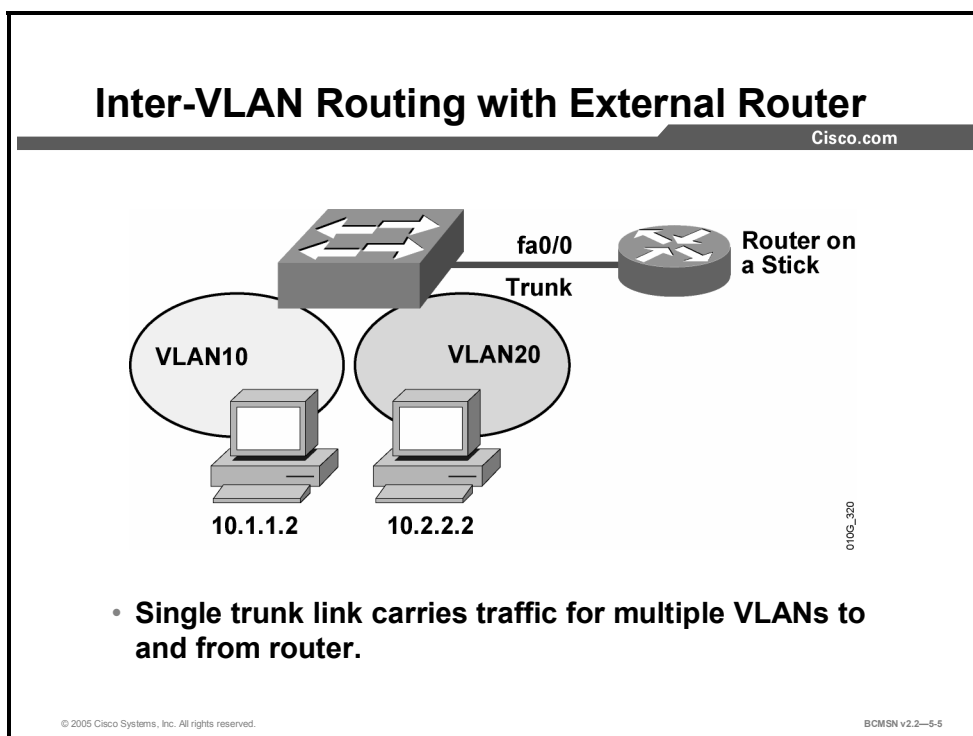
- **Advantages**
 - Layer 3 services not required on switch
 - Use existing router equipment
 - Trunking protocols not required
- **Disadvantages**
 - Route processing required, speed suboptimal
 - Requires multiple media between devices
 - 802.1Q tagging not available

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-5-4

The figure describes the advantages and disadvantages of using an external router with multiple interfaces for the purpose of inter-VLAN routing.

Inter-VLAN Routing Using an External Router and a Single Trunk

This topic describes how a router is used to route traffic between the VLANs on a switch.



If a switch supports multiple VLANs but has no Layer 3 capability to route packets between those VLANs, the switch must be connected to a router external to the switch. This will be most efficiently accomplished by providing a single trunk link between the switch and the router that can carry the traffic of multiple VLANs, which can in turn be routed by the router. This single physical link must be Fast Ethernet or greater to support Inter-Switch Link (ISL) encapsulation, but 802.1Q is supported on 10 Mb Ethernet router interfaces.

In the figure, the clients on VLAN10 need to establish sessions with a server that is in VLAN 20. This will require that traffic be routed between the VLANs as described in this table.

Step	Action
1.	The router accepts the packets from VLAN10 on its subinterface in that VLAN.
2.	The router performs Layer 3 processing based on the destination network address.
3.	Because the destination network is associated with a VLAN accessed over the trunk link, the router applies the appropriate VLAN identification to the packet.
4.	The router then routes the packet out the appropriate subinterface on VLAN20.

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 (subinterface) for each VLAN running between the switch and the router and ISL, or 802.1Q trunking must be enabled on the single physical connection between the router and switch. The routing table will

show as directly connected to all the subnets associated with the VLANs configured on the router subinterfaces. The router must learn routes to networks not configured on directly connected interfaces through dynamic routing protocols.

External Router with Single Interface: Advantages and Disadvantages

This subtopic describes the advantages and disadvantages of inter-VLAN routing on an external router.

Inter-VLAN Routing on an External Router

Cisco.com

- **Advantages**
 - Implementation is simple
 - Layer 3 services are not required on switch.
 - Router provides communications between VLANs.
- **Disadvantages**
 - Router is single point of failure.
 - Single traffic path may become congested.
 - Latency may be introduced as frames leave switch chassis.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—5-6

The figure describes the advantages and disadvantages of using an external router for inter-VLAN routing.

Inter-VLAN Routing Using External Router Configuration Commands

This topic describes the commands used to configure inter-VLAN routing on an external router.

Inter-VLAN Routing External Router Configuration Commands

Cisco.com

Configure on subinterface

- **encapsulation isl (or dot1Q) 10**
- **ip address 10.1.10.1 255.255.255.0**

Verify

- **show vlan 10**
- **show ip route**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—5-7

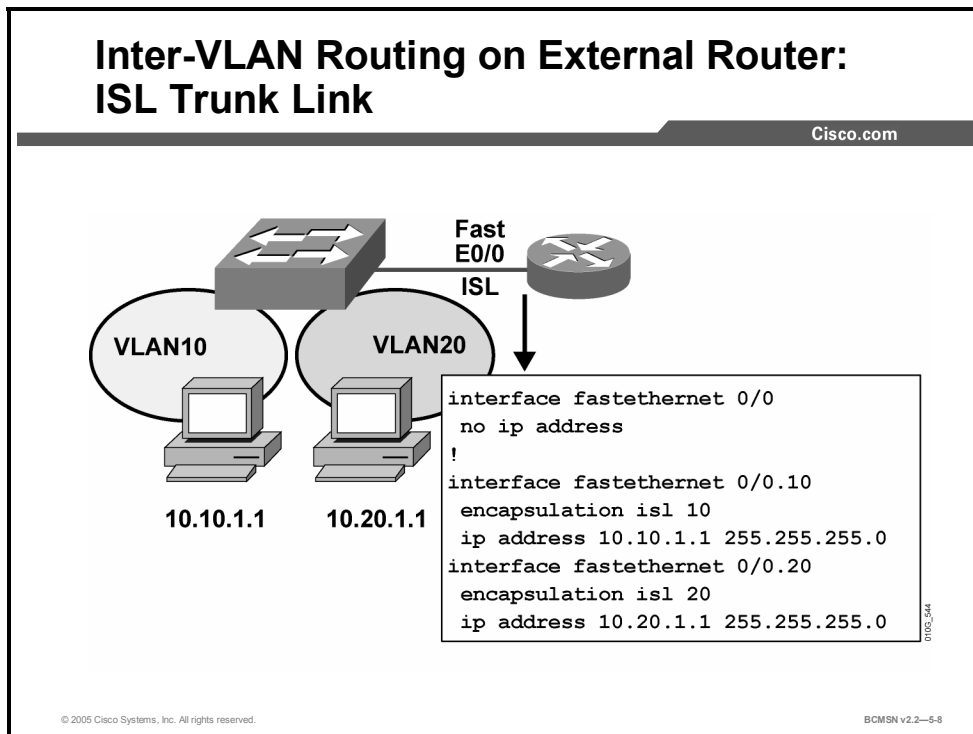
Inter-VLAN routing can be configured using an external router over either ISL or 802.1Q trunks. The commands for configuring the trunk interface on the router are shown in the table.

Inter-VLAN Routing on an External Router Commands

Command	Description
Router(config-subif) # encapsulation isl vlan_ID	This command sets the encapsulation to ISL on the subinterface and associates that subinterface with a particular VLAN. This command is repeated for each VLAN instance on the trunk.
Router(config-subif) # encapsulation dot1Q vlan_ID	This command sets the encapsulation to 802.1Q on the subinterface and associates that subinterface with a particular VLAN. This command is repeated for each VLAN instance on the trunk.
Router# show vlan [vlan_ID]	This command displays the subinterface to VLAN mappings, the trunk protocol configured, and the number of packets received and transmitted on each VLAN through the configured subinterfaces.
Router# show ip protocols	This command displays the configured routing protocols on the router along with the interfaces and networks configured for routing.
Router# show ip route	This command displays the routing information that will be used to move data between VLANs on this device.

How to Configure Inter-VLAN Routing Using an External Router

This topic illustrates how to configure routing between VLANs using an external router.



A router interface providing inter-VLAN routing on a trunk link must be configured with a subinterface for each VLAN that will be serviced across the link. Each subinterface on the physical link must then be configured with the same trunk encapsulation protocol. That protocol, either ISL or 802.1Q, is typically determined by what was configured on the switch side of the link.

Configuring an External Router using ISL Encapsulation

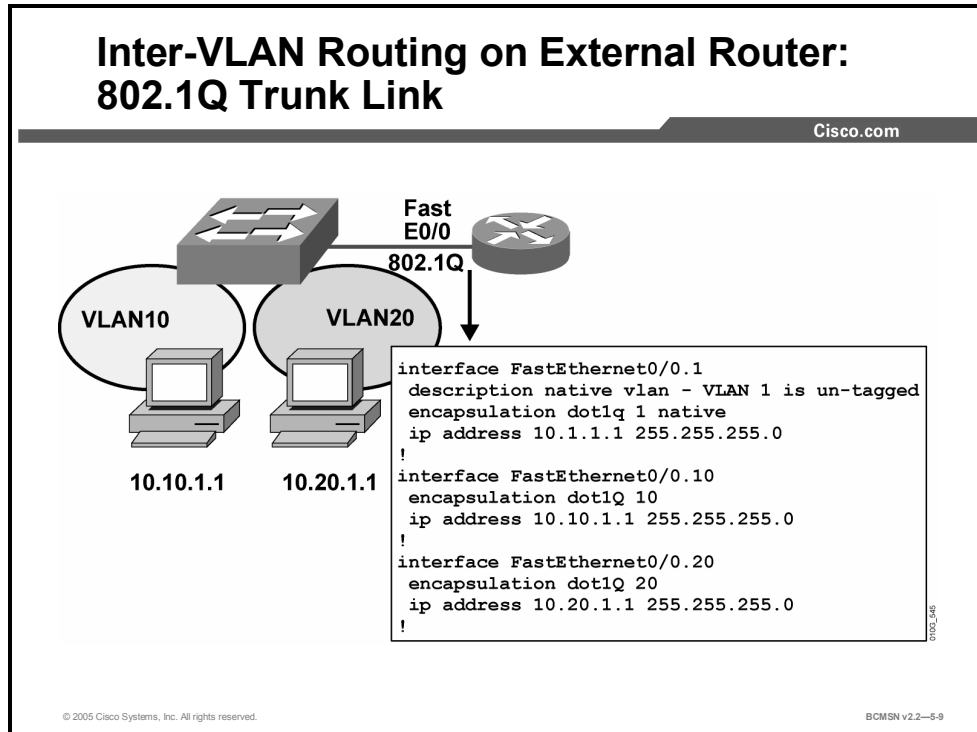
Use the **encapsulation isl *vlan_id*** subinterface configuration command to enable ISL trunking on a router subinterface.

Step	Action
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 The subnets of the VLANs are directly connected to the router. Routing between these subnets does not require a dynamic routing protocol. Routes to the subnets associated with each VLAN will appear in the routing table on directly connected interfaces.

Configuring an External Router using 802.1Q

This subtopic illustrates how to enable routing between VLANs over 802.1Q trunks.



Use the **encapsulation dot1q** subinterface configuration command to enable 802.1Q encapsulation on a router subinterface.

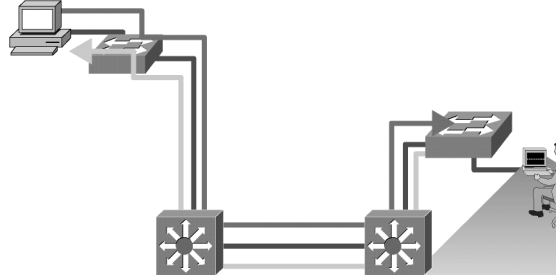
The use of the **native** keyword indicates that the untagged frames are associated with a particular subinterface. The other individual VLAN subinterfaces are configured as they are with ISL. The subinterface number need not match the dot-1Q VLAN number.

Verifying the Inter-VLAN Routing Configuration using Ping

This subtopic discusses how to verify the inter-VLAN routing configuration using ping.

Verifying Inter-VLAN Routing

Cisco.com



```
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.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-5-10

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, use the **ping** command from privileged mode:

```
Router# ping destination-ip-address
```

- Step 1** From the router, attempt to ping a host address on each VLAN to verify router connectivity.
- Step 2** From a host on a particular VLAN, attempt to ping a host on another VLAN to verify routing across the external router.

The **ping** command will return one of these responses:

- **Success rate is 100 percent or *ip-address* is alive:** This response occurs in 1 to 10 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 Time to Live (TTL) times out. The TTL default is 2 seconds.

Verifying the Inter-VLAN Routing Configuration

This subtopic describes commands used to verify inter-VLAN routing configuration.

Verifying the Inter-VLAN Routing Configuration

Cisco.com

```
Switch#show vlan
```

- Displays the current IP configuration per VLAN

```
Switch#show ip route
```

- Displays IP route table information

```
Switch#show ip interface brief
```

- Displays IP address on interfaces and current state of interface

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-5-11

Use the **show** commands to display the current (running) configuration, IP routing information, and IP protocol information. This will verify if the routing table represents the subnets of all VLANs.

Example: Displaying Inter-VLAN Configuration Information

```
Router#show vlans
```

```
Virtual LAN ID: 10 (Inter Switch Link Encapsulation)
```

```
  vLAN Trunk Interface: FastEthernet1/0.1
```

```
  Protocols Configured: Address:          Received:
Transmitted:
```

```
          IP          10.0.10.1          0          20
```

```
Virtual LAN ID: 20 (Inter Switch Link Encapsulation)
```

```
  vLAN Trunk Interface: FastEthernet1/0.2
```

```
  Protocols Configured: Address:          Received:
Transmitted:
```

```
          IP          10.0.20.1          0          20
```

Example: Displaying Routing Table Information

```
Router#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 - per-user static route, o - ODR
```

```
        P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
10.0.0.0/24 is subnetted, 2 subnets
```

```
C        10.0.10.0 is directly connected, FastEthernet1/0.1
```

```
C        10.0.20.0 is directly connected, FastEthernet1/0.2
```


Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Packets can be passed between VLANs using an external router with at least two interfaces.**
- **A single trunk link and router interface with one subinterface per VLAN can be configured to route between VLANs.**
- **Either ISL or 802.1Q can be configured as the trunking protocol between the switch and the router interface routing traffic between the VLANs on the trunk.**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—5-12

Deploying CEF-Based Multilayer Switching

Overview

Layer 3 switching provides a wire-speed mechanism by which to route packets between VLANs using tables that store Layer 2 and Layer 3 forwarding information in hardware. Cisco Express Forwarding (CEF) is the most efficient means of providing Layer 3 switching on a multilayer switch. CEF uses a very specific process to build forwarding tables in hardware and then uses that table information to forward packets at line speed.

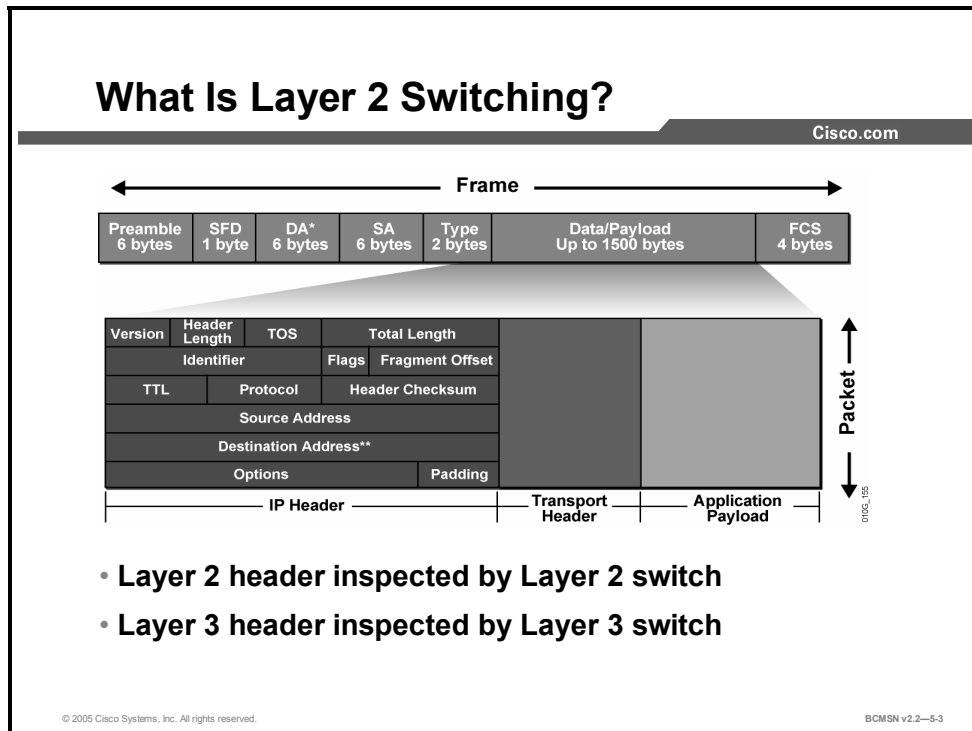
Objectives

Upon completing this lesson, you will be able to deploy CEF-based Multilayer Switching (MLS) so that all packets are CEF switched in hardware. This ability includes being able to meet these objectives:

- Define the Layer 2 switching process
- Define multilayer switching
- Define a CEF-based multilayer switch
- Identify the process that a multilayer switch uses to forward packets
- Configure and verify CEF operation on a Cisco Catalyst multilayer switch
- Identify common CEF problems and solutions
- List steps used to troubleshoot problems with a CEF-based multilayer switch

What Is Layer 2 Switching?

This topic identifies the features of a Layer 2 switch.



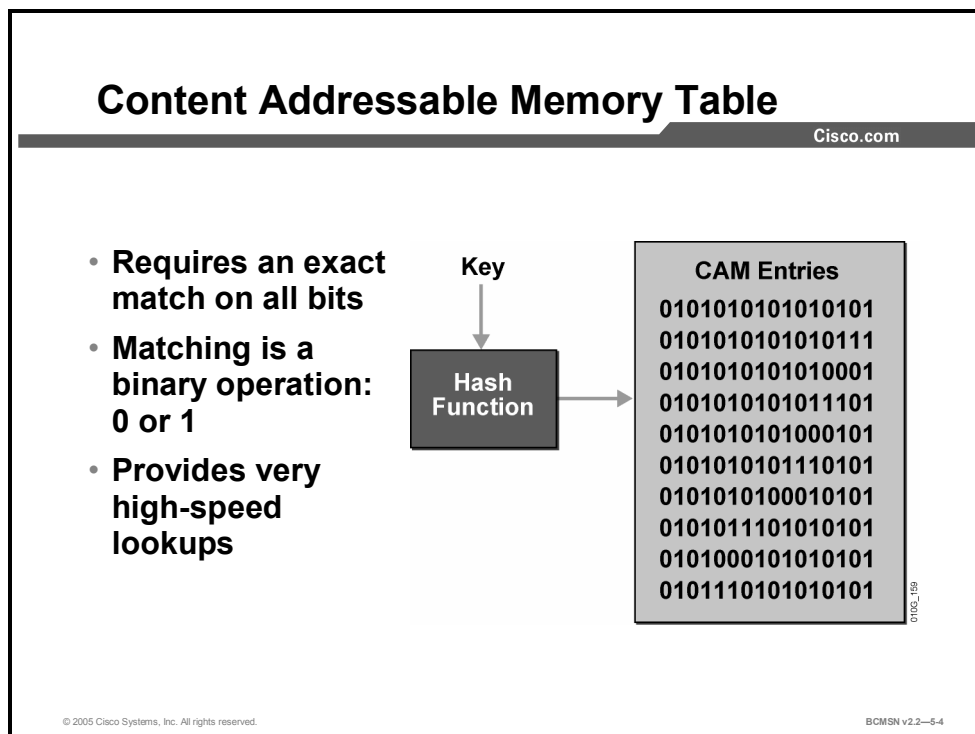
Layer 2 switching forwards frames based on information in the Layer 2 frame header as shown in the figure. Layer 2 switching occurs in hardware, thereby decreasing latency introduced by the software switching typically found in original bridge platforms. Switch hardware utilizes specialized chips, called ASICs, to handle frame manipulation and forwarding. Because the majority of frame manipulation and forwarding decisions occur in hardware, Layer 2 switching can provide wire-speed performance in ideal circumstances.

A Layer 2 switch builds a forwarding table as it records the source MAC address and the inbound port number of received frames. Because the switch simply moves frames from one port to another, based on the information in the forwarding table, operation is said to be transparent; the sending end station is unaware of the switch path traversed by the frame.

Additionally, to facilitate wire-speed lookups, the frame can be checked against an access control list (ACL) and quality of service (QoS) criteria that originate in Layer 3 software but are stored in tables in switch hardware. This process provides frame forwarding at wire speed while still qualifying the forwarding based on upper-layer criteria.

What Are Layer 2 Switching Tables?

This subtopic describes the multilayer switching table architectures.



Routing, switching, ACL, and QoS tables are stored in a high-speed table memory so that forwarding decisions and restrictions can be made in high-speed hardware. Cisco Catalysts have two primary table architectures:

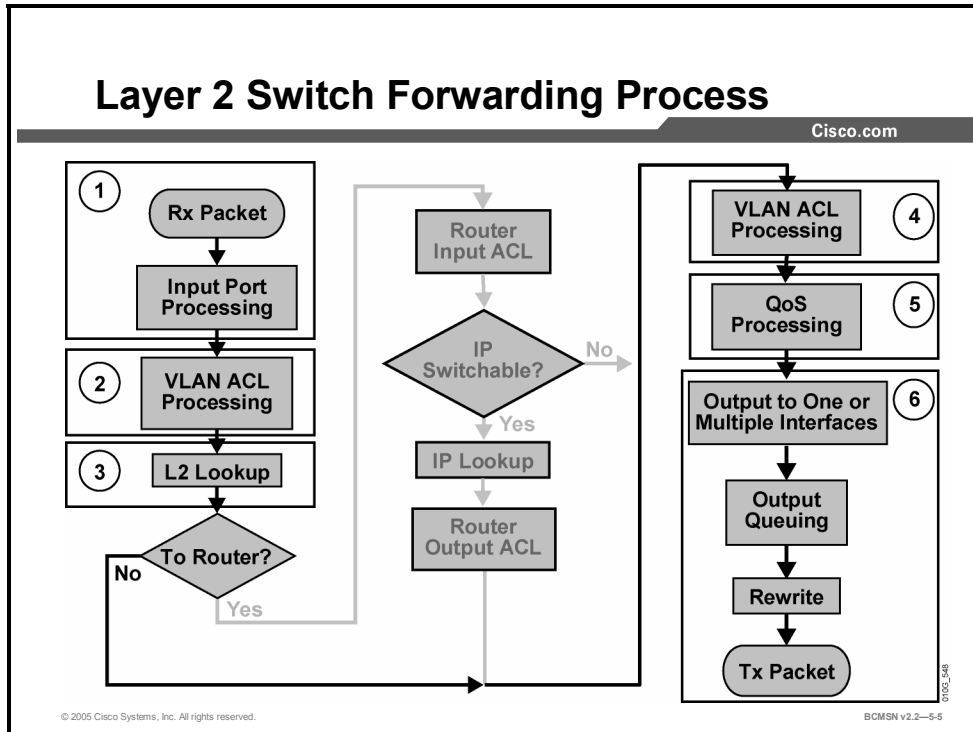
- **Content addressable memory (CAM) table:** This is the primary table used to make Layer 2 forwarding decisions. The table is built by recording the source address and inbound port of all frames. When a frame arrives at the switch with a destination MAC address of an entry in the CAM table, the frame is forwarded out only the port associated with that specific MAC address.
- **Ternary CAM (TCAM) table:** This table stores ACL, QoS, and other information generally associated with upper-layer processing. The TCAM details are discussed as a part of the Layer 3 switching process.

Table lookups are done with efficient search algorithms. A “key” is created to compare the frame to the table content. For example, the destination MAC address and VLAN ID (VID) of a frame would constitute the key for Layer 2 table lookup. This key is fed into a hashing algorithm, which produces a pointer into the table. The system uses the pointer to access a smaller specific area of the table without requiring searching the entire table.

In a Layer 2 table, all bits of all information are significant for frame forwarding (for example, VLANs, destination MAC addresses, and destination protocol types). However, in more complicated tables associated with upper-layer forwarding criteria, some bits of information may be too inconsequential to analyze. For example, an ACL may require a match on the first 24 bits of an IP address, but the last 8 bits may be insignificant information.

Identifying the Layer 2 Switch Forwarding Process

This subtopic describes how a switch processes a Layer 2 frame.



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.

How a Layer 2 Switch Forwards Packets

Step	Action
1.	The Layer 2 engine receives a frame.
2.	The Layer 2 engine performs the input ACL lookup.
3.	The Layer 2 lookup engine looks up the destination MAC address.
4.	The Layer 2 forwarding engine performs the outbound security ACL lookup.
5.	The Layer 2 forwarding engine performs the outbound QoS lookup.
6.	The Layer 2 forwarding engine forwards the packet.

What Is Multilayer Switching?

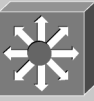
This topic identifies the operation of the key components required to implement Layer 3 switching.

What Is Multilayer Switching?

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

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-5-6

MLS includes the capability to switch data based on information at multiple layers. MLS also refers to a class of high-performance routers that provide Layer 3 services and simultaneously forward packets at wire speed through switching hardware. A Layer 3 switch performs packet switching, route processing, and intelligent network services.

Layer 3 switch processing forwards packets at wire speed by using ASIC hardware instead of microprocessor-based engines such as might be found on a traditional router. Specific Layer 3 components such as routing tables or ACLs are cached into hardware. The Layer 3 packet headers of data traffic will be analyzed and packets forwarded at line speeds based upon that cached information.

Layer 3 switching can occur at two different locations on the switch.

- **Centralized switching:** Switching decisions are made on the route processor by a central forwarding table, typically controlled by an ASIC.
- **Distributed switching:** Switching decisions can be made on a port or line-card level rather than on a central route processor. Cached tables are distributed and synchronized to various hardware components so processing can be distributed throughout the switch chassis.

Layer 3 switching takes place using one of these two methods, which are platform dependent.

- **Route caching:** Also known as flow-based or demand-based switching, a Layer 3 route cache is built in hardware as the switch sees traffic flow into the switch.

- **Topology-based switching:** Information from the routing table is used to populate the route cache regardless of traffic flow. The populated route cache is called the Forwarding Information Base (FIB). CEF is the facility that builds the FIB.

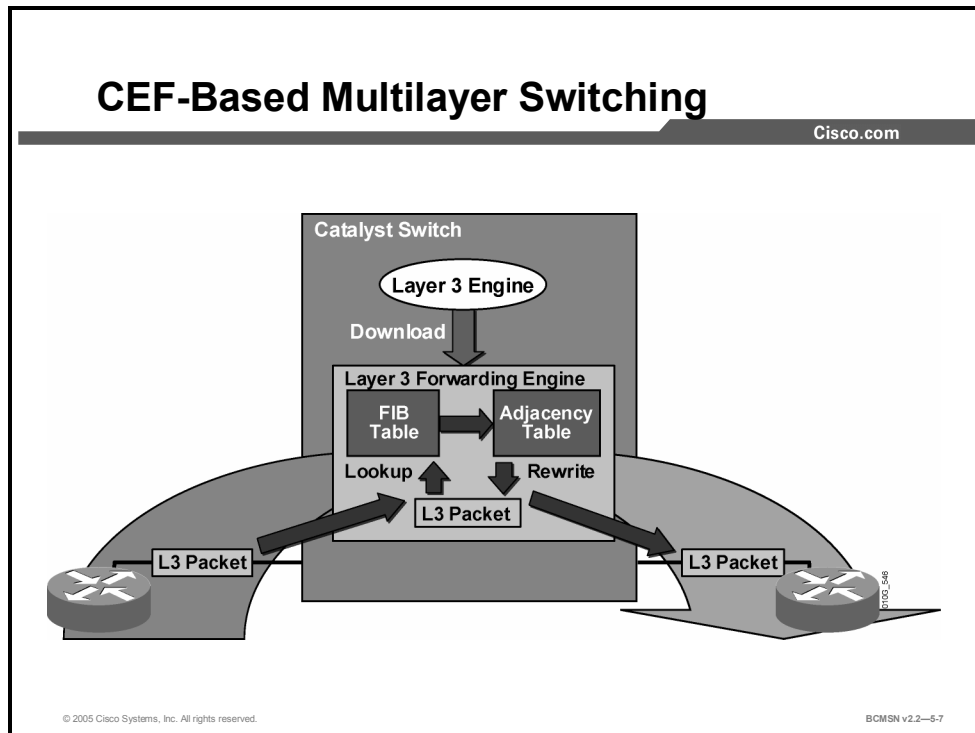
References

For additional information, refer to this resource:

http://www.cisco.com/en/US/partner/products/hw/switches/ps5304/products_configuration_guide_chapter09186a00800f0a7b.html#22246

What Is a CEF-Based Multilayer Switch?

This topic defines Cisco Express Forwarding.



Cisco Layer 3 devices can use a variety of methods to switch packets from one port to another. The most basic method of switching packets between interfaces is called process switching. Process switching moves packets between interfaces, based on information in the routing table and the Address Resolution Protocol (ARP) cache, on a scheduled basis. As packets arrive, they will be moved into a queue to wait for further processing. When the scheduler runs, the outbound interface will be determined and the packet will be switched. Waiting for the scheduler introduces latency.

To speed the switching process, strategies exist to switch packets on demand as they arrive on an interface and to cache information necessary to make packet-forwarding decisions.

CEF uses these strategies to expediently switch data packets to their destination. It caches information generated by the Layer 3 routing engine. CEF caches routing information in one table (the FIB) and caches Layer 2 next-hop addresses for all FIB entries in an adjacency table. Because CEF maintains multiple tables for forwarding information, parallel paths can exist and enable CEF to load balance per packet.

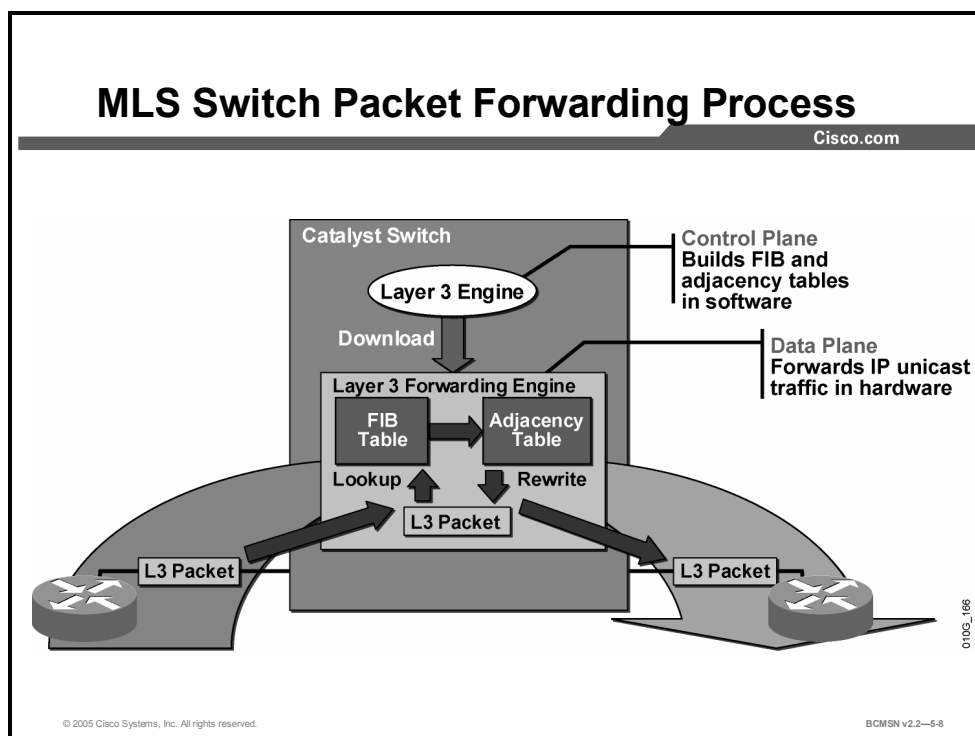
CEF operates in one of two modes.

- **Central CEF mode:** The CEF FIB and adjacency tables reside on the route processor, and the route processor performs the express forwarding. Use this CEF mode when line cards are not available for CEF switching, or when features are not compatible with distributed CEF.
- **Distributed CEF (dCEF) mode:** When dCEF is enabled, line cards maintain identical copies of the FIB and adjacency tables. The line cards can perform the express forwarding by themselves, relieving the main processor of involvement in the switching operation.

Distributed CEF uses an Inter-Process Communication (IPC) mechanism to ensure synchronization of FIBs and adjacency tables on the route processor and line cards.

Identifying the Multilayer Switch Packet Forwarding Process

This topic illustrates the process used by a multilayer switch to forward packets.



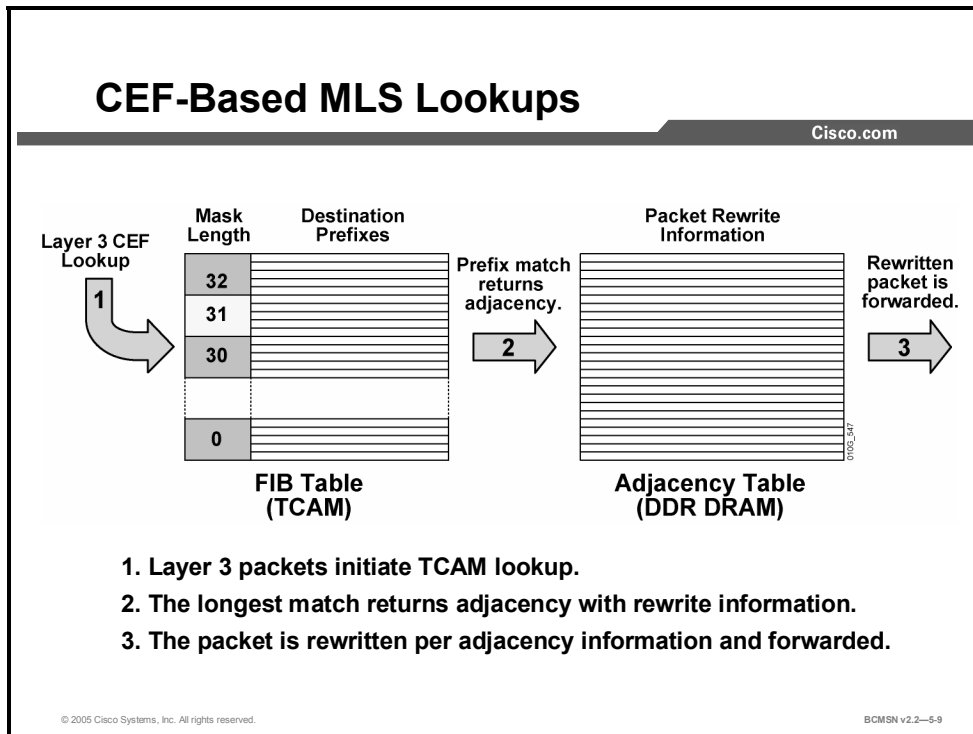
CEF separates the control plane hardware from the data plane hardware and switching. ASICs in switches are used to separate the control plane and data plane, thereby achieving higher data throughput. The control plane is responsible for building the FIB table and adjacency tables in software. The data plane is responsible for forwarding IP unicast traffic using hardware.

When traffic cannot be processed in hardware, it must receive processing in software by the Layer 3 engine, thereby not receiving the benefit of expedited hardware-based forwarding. A number of different packet types may force the Layer 3 engine to process them. Some examples of IP exception packets are the following:

- IP 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 Time to Live (TTL) counter
- Packets that are forwarded to a tunnel interface
- Packets that arrive with unsupported encapsulation types
- Packets that are routed to an interface with unsupported encapsulation types
- Packets that exceed the maximum transmission unit (MTU) of an output interface and must be fragmented

CEF-Based Tables and MLS Lookups

This subtopic describes the CEF process.



CEF-based tables are initially populated and used as follows:

- The FIB is derived from the IP routing table and is arranged for maximum lookup throughput.
- The adjacency table is derived from the 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.
- When the CEF TCAM table is full, a wildcard entry redirects frames to the Layer 3 engine.
- When the adjacency table is full, a CEF TCAM table entry points to the Layer 3 engine to redirect the adjacency.
- The FIB lookup is based on the Layer 3 destination address prefix (longest match).

FIB Table Updates

The FIB table is updated when the following occur:

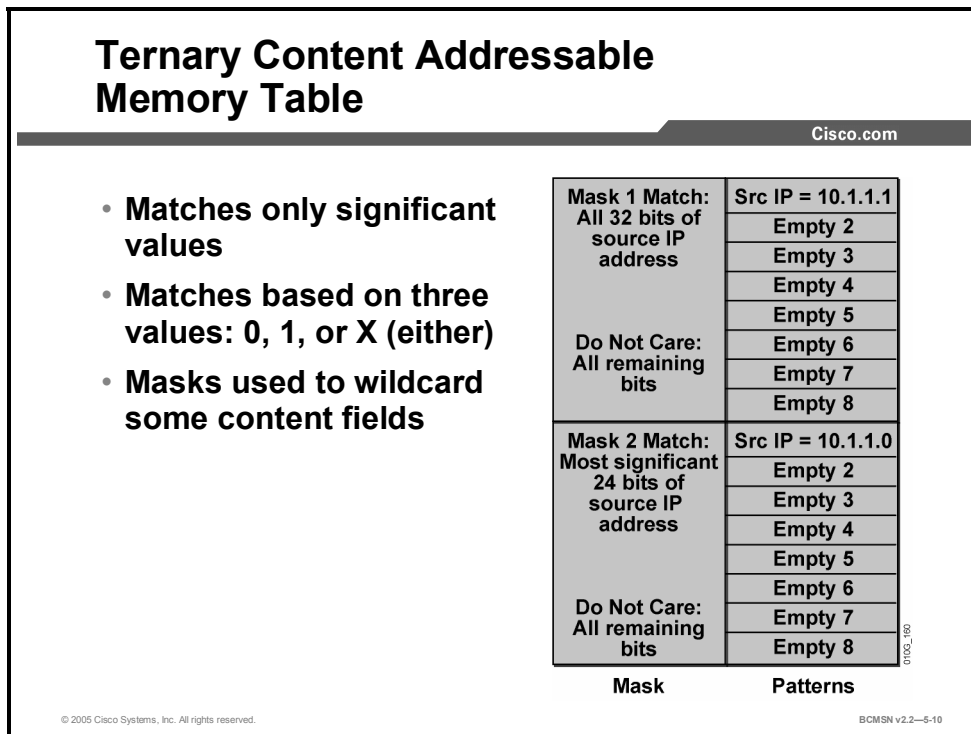
- An ARP entry for the destination next hop changes, ages out, or is removed.
- The routing table entry for a prefix changes.
- The routing table entry for the next hop changes.

These are the basic steps for initially populating the adjacency table.

- 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 a 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.

Ternary Content Addressable Memory Table

This subtopic describes how the TCAM is used.



In specific high-end switch platforms, the TCAM is a portion of memory designed for rapid, hardware-based table lookups of Layer 3 and Layer 4 information. In the TCAM, a single lookup provides all Layer 2 and Layer 3 forwarding information for frames, including CAM and ACL information.

The figure displays the ACL information stored in the TCAM table that would result in a packet being permitted or denied.

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 as wildcards in some content fields.

These two access control list entries are referenced in the figure, as it shows how their values would be stored in the TCAM:

```
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
```

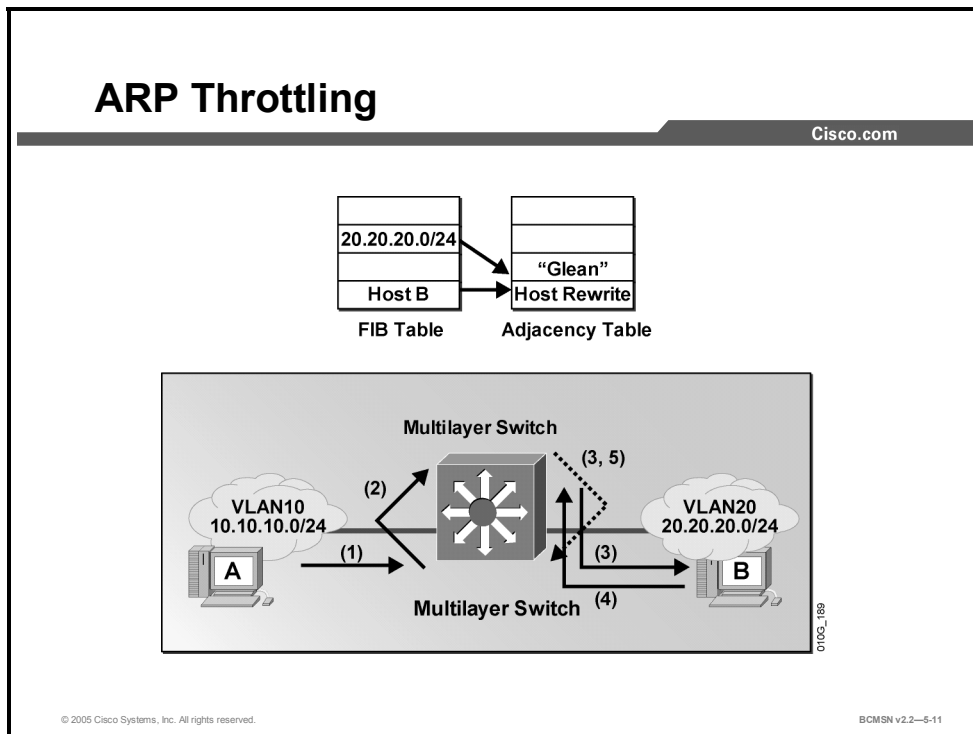
The TCAM table entries in the figure consist of types of regions:

- **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.

ARP Throttling

This subtopic describes ARP throttling.



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. A glean adjacency entry indicates that a particular next hop should be directly connected, but there is no MAC header rewrite information available.
- 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.

The Layer 3 engine installs adjacency for host B and removes the drop adjacency. 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, 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.

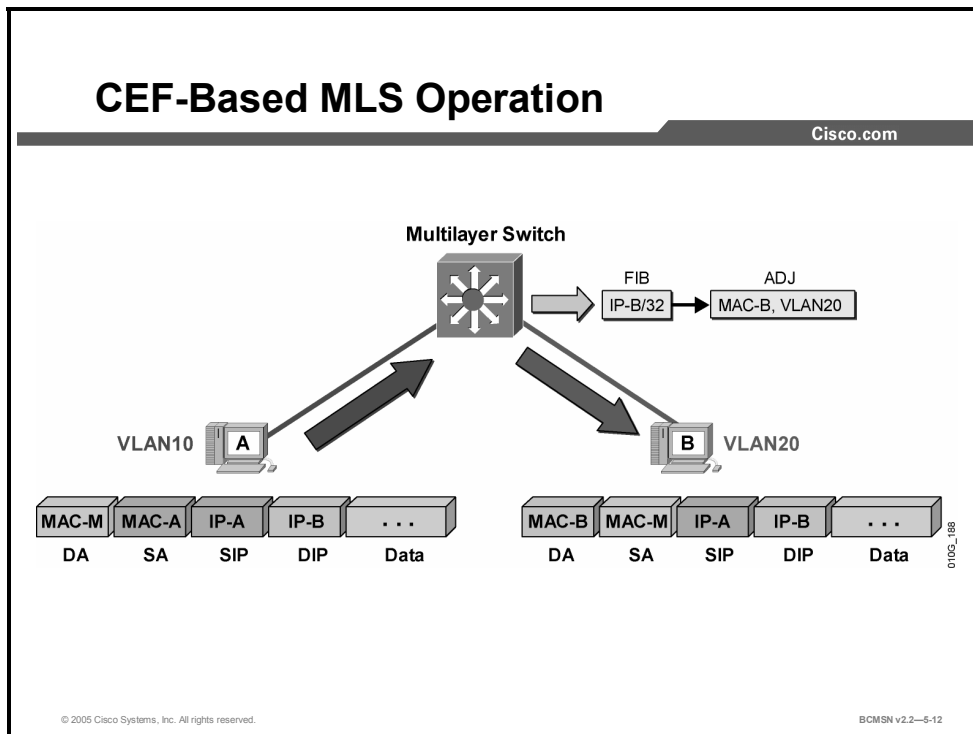
In addition to adjacencies associated with next-hop interfaces (host-route adjacencies), other types of adjacencies are used to expedite switching when certain exception conditions exist. When the prefix is defined, prefixes requiring exception processing are cached with one of the following special 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 appended to packets, FIB requires the appended header 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.

CEF-Based MLS Operation

This subtopic describes CEF-based MLS operation.



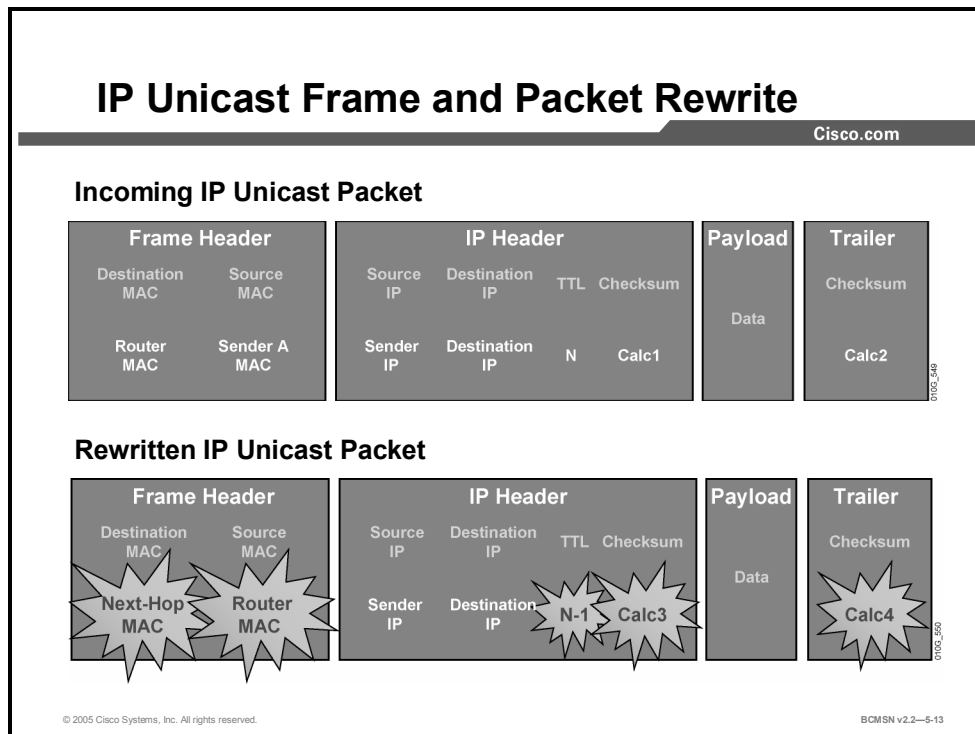
These are the steps that would occur when you use CEF to forward frames between host A and host B on different VLANs.

- Step 1** 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 2** 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 3** The Layer 3 engine installs an ARP throttling adjacency in the switch for the host B IP address.
- Step 4** The Layer 3 engine sends ARP requests for host B on VLAN20.
- Step 5** Host B sends an ARP response to the Layer 3 engine.
- Step 6** The Layer 3 engine installs the resolved adjacency in the switch (removing the ARP throttling adjacency).
- Step 7** The switch forwards the packet to host B.
- Step 8** The switch receives a subsequent packet for host B (IP-B).
- Step 9** 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.

The switch rewrites packets per the adjacency information and forwards the packet to host B on VLAN20.

Frame Rewrite Using CEF

This subtopic describes the frame rewrite process used by CEF.



The figure shows how the frame and packet header would be altered when CEF is used to forward frames. 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 TTL is decremented by one and, as a result, the IP header checksum is recalculated.
- The frame checksum must be recalculated.

Configuring and Verifying CEF

This topic will review how to change the status of CEF and verify its operation.

Configuring and Verifying CEF

Cisco.com

Configuring CEF

- **ip cef** (enabled by default)
- **ip route-cache cef** (only on VLAN interface)

Verifying CEF

- **show ip cef fa 0/1 detail**
- **show adjacency fa 0/1 detail**

Catalyst 3550, 3750, 4000, and 6500

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—5-14

Use these commands to configure CEF when possible and verify its operation.

CEF Configuration Commands

Command	Description
Switch(config-if)# ip cef	On a Catalyst 4000, enables CEF if it has been previously disabled. CEF is on by default.
Switch(config-if)# no ip cef	Disables CEF on a Catalyst 4000.
Switch(config-if)# ip route-cache cef	On a Catalyst 3550, enables CEF if it has been previously disabled on an interface. CEF is on by default.
Switch(config-if)# no ip route-cache cef	Disables CEF on a Catalyst 3550.
Switch# show ip cef [type mod/port] [detail]	Verifies CEF operation.
Switch# show interface type mod/port begin L3	Displays information about Layer 3 switched traffic for the interface.
Switch# show interface type mod/port include switched	Displays counts on packets switched at Layer 2 and 3.

Hardware Layer 3 switching is permanently enabled on Catalyst 6500 series supervisor engine 720 with Policy Feature Card 2 (PFC3), Multilayer Switch Feature Card 3 (MSFC3), and Distributed Forwarding Card (DFC). No configuration is required, and CEF cannot be disabled.

The **no ip cef** command can be used to disable CEF on the Catalyst 4000 or the **no ip route-cache cef** command on a Catalyst 3550 interface.

If CEF is enabled globally, it is automatically enabled on all interfaces as long as IP routing is enabled on the device. It can then be enabled or disabled on an interface basis. Cisco recommends that CEF be enabled on all Layer 3 interfaces. If CEF is disabled 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.

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. This ensures that packets for a given host pair arrive in order. 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 increases.

Verifying CEF

This subtopic lists commands used to verify the operation of CEF.

Verifying CEF

Cisco.com

```
Switch#show ip cef [type mod/port/ vlan_interface] [detail]
```

```
Switch# show ip cef vlan 11 detail

IP CEF with switching (Table Version 11), flags=0x0
 10 routes, 0 reresolve, 0 unresolved (0 old, 0 new), peak 0
 13 leaves, 12 nodes, 14248 bytes, 14 inserts, 1 invalidations
 0 load sharing elements, 0 bytes, 0 references
 universal per-destination load sharing algorithm, id 4B936A24
 2(0) CEF resets, 0 revisions of existing leaves
 Resolution Timer: Exponential (currently 1s, peak 1s)
 0 in-place/0 aborted modifications
 refcounts: 1061 leaf, 1052 node

Table epoch: 0 (13 entries at this epoch)

172.16.11.0/24, version 6, epoch 0, attached, connected
0 packets, 0 bytes
 via Vlan11, 0 dependencies
  valid glean adjacency
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—5-15

Verifying Layer 3 Switching

This subtopic discusses how the **show ip cef detail** command indicates if CEF is running globally. Specify an interface to verify CEF operation on that interface.

Verify Layer 3 Switching

Cisco.com

```
Switch#show interface {{type mod/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  
.....  
Switch#
```

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—5-16

Display CEF Statistics

This subtopic discusses the **show interface** command with the `| begin L3` argument to verify that Layer 3 traffic is being switched thereby utilizing CEF.

Displaying Hardware Layer 3 Switching Statistics

Cisco.com

```
Switch#show interfaces {{type mod/port}} | {port-channel number}} include switched
```

```
Switch#show interfaces gigabitethernet 9/5 | include switched
L2 Switched: ucast: 8199 pkt, 1362060 bytes - mcast: 6980 pkt, 371952 bytes
L3 in Switched: ucast: 3045 pkt, 742761 bytes - mcast: 0 pkt, 0 bytes mcast
L3 out Switched: ucast: 2975 pkt, 693411 bytes - mcast: 0 pkt, 0 bytes
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—5-17

Use the **show interfaces** command with the `| include switch` argument to show switching statistics at each layer for the interface. Verify that Layer 3 packets are being switched.

Displaying Detailed Adjacency Information

Here is the command used to display detailed information about the adjacency table.

Adjacency Information

Cisco.com

```
Switch#show adjacency [{"type mod/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
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2--5-18

Each time an adjacency entry is created, a Layer 2 data link–layer header for that adjacent node is precomputed and stored in the adjacency table. This information is subsequently used for encapsulation during CEF switching of packets.

Output from the command **show adjacency detail** displays the content of the information to be used during this Layer 2 encapsulation. Verify that the header information displays as would be expected during Layer 2 operations, not using precomputed encapsulation from the adjacency table. Adjacency statistics are updated approximately every 60 seconds.

Also, the **show cef drops** command will display if packets are being dropped due to adjacencies that are either incomplete or nonexistent. There are two known reasons for incomplete or nonexistent adjacencies.

- The router cannot use ARP successfully for the next-hop interface.
- After a **clear ip arp** or a **clear adjacency** command, the router marks the adjacency as incomplete, and then it fails to clear the entry.

The symptoms of an incomplete adjacency include random packet drops during a ping test. Use the **debug ip cef** command to view CEF drops due to an incomplete adjacency.

Debugging CEF Operations

The debug facility can be used to display detailed information on CEF operations.

Debugging CEF Operations

Cisco.com

```
Switch#debug ip cef {drops | access-list | 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

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—5-19

Use the **debug ip cef** arguments to limit the debug output, thereby reducing the overhead of the debug command and providing focus on a specific CEF operation.

```
debug ip cef {drops [access-list] | receive [access-list] | events
[access-list] | prefix-ipc [access-list] | table [access-list]}
```

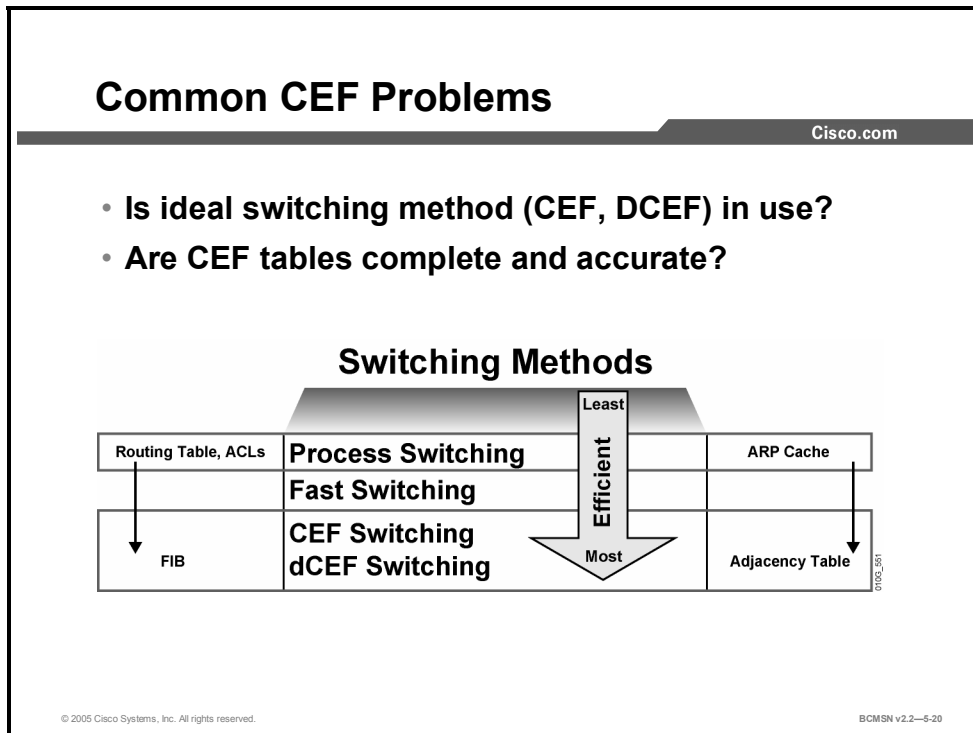
Adding an argument to the **debug** command limits the debug output as follows.

- *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
 - Notification that 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

Common CEF Problems and Solutions

This topic describes common problems and associated solutions when verifying CEF operation.



CEF is the fastest means of switching Layer 3 packets in hardware. The CEF tables stored in hardware are populated from information gathered by the route processor. Troubleshooting CEF operations therefore has two primary steps.

- Ensure that the normal Layer 3 operations on the route processor are functioning properly so that the switch tables will be populated with accurate and complete information.
- Verify that information from the route processor has properly populated the FIB and adjacency table, and is being used by CEF to switch Layer 3 packets in hardware.

Troubleshooting CEF is, in essence, verifying that packets are indeed receiving the full benefit of CEF switching and not being “punted” to a slower packet switching or processing method. The Cisco term "punt" describes the action of sending a packet "down" to the next-fastest switching level. The following list defines the order of preferred Cisco IOS switching methods, from fastest to slowest.

- Distributed CEF
- CEF
- Fast switching
- Process switching

A punt occurs when the preferred switching method did not produce a valid path or, in CEF, a valid adjacency. If the CEF lookup process fails to find a valid entry in the FIB, CEF will install a punt adjacency to the less-preferred system. CEF will punt all packets with that adjacency to the next-best switching mode in order to forward all the packets by some means, even if that means is less efficient.

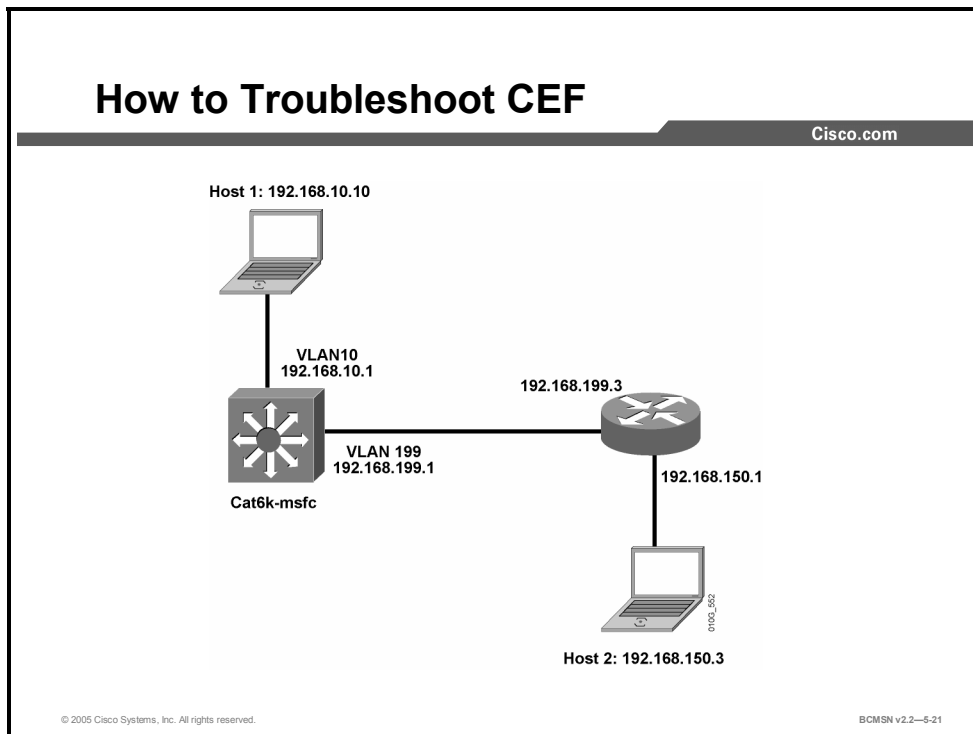
Some basic CEF problems and associated solutions are described in the “CEF Problems and Solutions” table.

CEF Problems and Solutions

	Problem	Solution
1.	CEF or dCEF has been disabled on an interface or line card.	<p>Verify that IP routing is enabled; disabling IP routing will disable CEF globally.</p> <p>Determine if lack of memory resource is disabling CEF or dCEF. This is possible if an inadequate amount of memory is available to store a large number of routing table entries in the FIB.</p> <p>Review the switch configuration for features that may not be compatible with CEF. Specifically look for commands and features related to switching services.</p> <p>MALLOCFAIL and FIBDISABLE messages will indicate if CEF is not functional.</p>
2.	Routes are not appearing in routing table as designed.	Troubleshoot routing protocols to ensure that routing table is being properly populated.
3.	ARP is not resolving MAC addresses.	<p>Verify that the IP host exists on a media accessible to the host trying to ARP resolve. Use a packet analyzer to determine if ARP is not functioning or if the sending host is not caching the ARP resolutions.</p> <p>Check that subsystems of ARP are functioning correctly: IP addresses may be invalid, local addresses may resolve, but remote addresses may not resolve to a default gateway, and so forth.</p>
4.	Existing routing table entries do not appear in the CEF FIB table, or next hop is not appropriate.	This is a Cisco IOS software CEF issue.
5.	Resolved ARP addresses do not appear in the adjacency table.	This is a Cisco IOS software CEF issue.
6.	Part of an adjacency is an “incomplete” or “drop” adjacency.	Check that unsupported software features have not been enabled.
7.	CEF operations are suspect or their existence needs to be verified.	Debug CEF to observe all messages regarding CEF operations. Messages are indicative of successful or failed CEF operations.

How to Troubleshoot Layer 3 Connectivity in a CEF-based Multilayer Switch

This topic describes a process by which to verify and troubleshoot CEF operation.



The following items will be covered in this topic.

- Troubleshoot host connectivity using CEF
- Displaying detailed adjacency information
- Debugging CEF operations

Troubleshoot Host Connectivity Using CEF

The CEF tables stored in hardware are populated from information gathered by the route processor. To properly troubleshoot CEF operations, first ensure that the normal Layer 3 operations on the route processor are functioning properly so that the CEF tables will be populated with accurate and complete information. Next, verify that information from the route processor has properly populated the FIB and adjacency table used by CEF to perform Layer 3 switching of packets.

The steps below will verify if packet transfer between these hosts is occurring using CEF:

- Host 1 in VLAN10 with an IP address of 192.168.10.10
- Host 2 in VLAN150 with an IP address of 192.168.150.3

1. Verify CEF Operations

Verify that CEF is operational at global or interface level using these commands:

```
show ip cef summary
show ip cef vlan 10
```

Note CEF cannot be turned off on most Catalyst platforms. If CEF is not operational, it is likely that the Catalyst has disabled the feature. This may be due to a software, feature, or hardware incompatibility or due to inadequate memory to support a large FIB and adjacency table.

2. Verify the Configuration

If CEF is not operational, display the running configuration to determine if any switching functions have been configured that might disable CEF operations.

If CEF is operational, display the running configuration to verify the IP configuration of the Layer 3 interfaces used for the hosts to communicate. The IP addresses should be appropriate for the subnet, and the interfaces should not be shut down. This is a sample of the configuration output expected for the VLANs associated with the host communication. On this router, VLAN 199 is the transit path that will be traversed to arrive at subnet 192.168.150.0:

```
Switch# show running-config
interface VLAN 10
  description Source VLAN
  ip address 192.168.10.1 255.255.255.0
!
interface VLAN 199
  description Transit VLAN
  ip address 192.168.199.1 255.255.255.0
```

3. Verify Population of the Routing Table on the Route Processor

The routing protocols and route processor must populate the routing table accurately before those routing table entries can be of use as they are transferred to the FIB to facilitate Layer 3 switching. Verify the routing table by referring to a network diagram, knowing what routes should appear in the routing table, and then execute the **show ip route** command. In the case of troubleshooting connectivity to the specific network of the destination host (192.168.150.3/24), use this command:

```
Switch# show ip route | include 192.168.150.0
O 192.168.150.0/24 [110/2] via 192.168.199.3, 00:13:00, VLAN 199
```

The network is accessible via next-hop address 192.168.199.3; therefore the ARP entry by which to access 192.168.150.3 should be the MAC address resolved for 192.168.199.3.

4. Verify an ARP Entry on the Route Processor

Verify that there is an ARP entry for the next-hop IP address before checking if that entry is represented in the adjacency table.

```
Switch# show ip arp 192.168.199.3

Protocol    Address          Age    Hardware Addr   Type   Interface
Internet    192.168.199.3   176    0030.7150.6800  ARPA   VLAN 199
```

5. Verify the CEF FIB Table Entry for the Route

Step 3 verified that a route to network 192.168.150.0 existed in the routing table. Now, verify that a CEF FIB entry exists to that same destination to ensure that packets will be CEF switched using the FIB rather than process switched using the routing table.

```
Switch# show ip cef 192.168.150.0

192.168.150.0/24, version 298, cached adjacency 192.168.199.3
0 packets, 0 bytes
via 192.168.199.3, VLAN 199, 0 dependencies
next-hop 192.168.199.3, VLAN 199
valid cached adjacency
```

This output verifies that there is a valid CEF entry for the destination network; packets can be CEF switched to the destination host.

6. Verify an Adjacency Table Entry for the Destination

Now, verify that the FIB entry shown in Step 5 has an associated adjacency table entry by using this command:

```
Switch# show adjacency detail | begin 192.168.199.3

IP VLAN 199 192.168.199.3(7)
0 packets, 0 bytes
003071506800
.....
...
.
```

The above output indicates that there is an adjacency for the next-hop IP address. The destination MAC address (003071506800) is the MAC address in the ARP table, as displayed in Step 4, above.

The counters (0 packets, 0 bytes) are almost always 0, as packets are switched in hardware and, as such, they never reach the route processor, which is required to increment counters.

7. Verify CEF from the Supervisor Engine

The CEF FIB and adjacency table entries shown in the example can also be verified from the supervisor engine on modular switch platforms such as the 6500 series switches. This step is not necessary on fixed configuration switches such as the 3550.

To display an FIB entry for the specific network from the supervisor engine:

```
Console> (enable) show mls entry cef ip 192.168.150.0/24
```

Mod	FIB-Type	Destination-IP	Destination-Mask	NextHop-IP	Weight
15	resolved	192.168.150.0	255.255.255.255	192.168.199.3	1

To display an FIB entry for the specific network from the supervisor engine:

```
Console> (enable) show mls entry cef ip 192.168.150.0/24 adjacency
```

```
Mod:15
```

```
Destination-IP : 192.168.199.3 Destination-Mask : 255.255.255.255
```

```
FIB-Type : resolved
```

AdjType	NextHop-IP	NextHop-Mac	VLAN	Encp	TX-Packets
connect	192.168.199.3	00-30-71-50-68-00	199	ARPA	0

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Layer 2 switching is performed on an MLS.**
- **Layer 3 switching is high-performance packet switching in hardware.**
- **MLS functionality is implemented through CEF.**
- **CEF uses tables in hardware to forward packets.**
- **Specific commands are used to enable and verify CEF operations.**
- **CEF problems can be matched to specific solutions.**
- **Specific commands are used to troubleshoot and solve CEF problems.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—5-22

References

For additional information, refer to these resources:

Cisco Systems, Inc., *How to Choose the Best Router Switching Path for Your Network*.

http://www.cisco.com/en/US/partner/tech/tk827/tk831/technologies_white_paper09186a00800a62d9.shtml

- Cisco Systems, Inc., *Cisco Express Forwarding White Paper*.
http://www.cisco.com/warp/customer/cc/pd/iosw/iore/tech/cef_wp.htm

Enabling Routing Between VLANs on a Multilayer Switch

When Multiple VLANs are configured on a multilayer switch, routing between those VLANs can occur on the switch itself through the configuration of Layer 3 Switch Virtual Interfaces (SVI). SVIs are configured and verified using Layer 3 Cisco IOS commands to facilitate inter-VLAN routing on a multilayer switch. It is also possible to convert layer 2 switch ports to operate as layer 3 interfaces.

Objectives

Upon completing this lesson, you will be able to enable routing between VLANs. This ability includes being able to meet these objectives:

- Describe a Layer 3 Switch Virtual Interface
- Describe a routed interface on a multilayer switch
- Identify the commands used to configure inter-VLAN communication on a multilayer switch
- Configure and verify inter-VLAN routing on a multilayer switch

Layer 3 Switch Virtual Interface

This topic describes a VLAN Switch Virtual Interface.

Layer 3 Switch Virtual Interfaces

Cisco.com

- **A Switch Virtual Interface (SVI) allows communication to switch ports with the same VLAN ID**
- **Created upon entering VLAN interface configuration mode**
- **Configure an SVI for each VLAN to route traffic**
- **Supports routing protocol and bridging configurations**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-5-3

A Switch Virtual Interface (SVI) is a virtual Layer 3 interface that can be configured for any VLAN that exists on a Layer 3 switch. It is “virtual” in that there is no physical interface for the VLAN, and yet it can accept configuration parameters applied to any Layer 3 router interface. The SVI for the VLAN provides Layer 3 processing for packets from all switch ports associated with that VLAN. Only one SVI can be associated with a VLAN. You configure an SVI for a VLAN for the following reasons:

- To provide a default gateway for a VLAN so traffic can be routed between VLANs
- To provide fallback bridging if it is required for nonroutable protocols
- To provide Layer 3 IP connectivity to the switch
- To support routing protocol and bridging configurations

By default, an SVI is created for the default VLAN (VLAN1) to permit remote switch administration. Additional SVIs must be explicitly created.

SVIs are created the first time a VLAN interface configuration mode is entered for a particular VLAN SVI. The VLAN corresponds to the VLAN tag associated with data frames on an ISL or 802.1Q encapsulated trunk or to the VLAN ID (VID) configured for an access port. Configure and assign an IP address to each VLAN SVI that is to route traffic off of and onto the local VLAN.

Routed Interfaces on a Multilayer Switch

This topic describes routed interfaces on a multilayer switch.

Routed Interfaces on a Multilayer Switch

Cisco.com

- **Physical switch port with Layer 3 capability**
- **Not associated with a VLAN**
- **Serves as default gateway for devices out that switch port**
- **Requires removal of Layer 2 port functionality**

```
Switch(config)#interface FastEthernet 0/1
Switch(config-if)#no switchport
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-5-4

A routed switch port is a physical switch port on a multilayer switch that is capable of Layer 3 packet processing. A routed port is not associated with a particular VLAN, as is an access port or SVI. A routed port behaves like a regular router interface, except that it does not support VLAN subinterfaces. Routed switch ports can be configured using most commands applied to a physical router interface, including the assignment of an IP address and the configuration of Layer 3 routing protocols.

A routed switch port is similar to an SVI in that it is a switch port that provides Layer 3 packet processing. SVIs generally provide Layer 3 services for devices connected to the ports of the switch where the SVI is configured. Routed switch ports can provide a Layer 3 path into the switch for a number of devices on a specific subnet, all of which are accessible from a single, physical switch port.

The number of routed ports and SVIs that can be configured on a switch is not limited by software. However, the interrelationship between these interfaces other features configured on the switch may overload the CPU due to hardware limitations.

Routed switch ports are typically configured by removing the Layer 2 switch port capability of the switch port. On most switches, the ports are Layer 2 ports by default. On some switches, the ports are Layer 3 ports by default. The layer at which the port functions determines the commands that can be configured on the port.

Configuration Commands for Inter-VLAN Communication on a Multilayer Switch

This topic describes commands used to configure inter-VLAN routing on a multilayer switch through a Switched Virtual Interface.

Configuration Commands for Inter-VLAN Communication on an SVI

Cisco.com

Configure

- **ip routing**
- **interface vlan 10**
 - **ip address 10.1.1.1 255.255.255.0**
- **router eigrp 50**
 - **network 10.0.0.0**

Verify

- **show ip route**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-5-5

These commands are used to configure inter-VLAN routing on a multilayer switch through an SVI.

Inter-VLAN Routing through SVI Commands

Command	Description
Switch(config)# ip routing	Enables Layer 3 routing on the switch. Enabled by default on some Catalyst platforms.
Switch(config)# interface vlan <i>vlan-id</i>	Creates SVI for the VLAN and moves to interface configuration mode for that SVI.
Switch(config-if)# ip address <i>ip-address mask</i>	Assigns a Layer 3 address to the SVI. This will likely be the default gateway for host on that VLAN.
Switch(config)# router <i>routing-protocol <options></i>	(Optional) Invokes a dynamic routing protocol, Enhanced Interior Gateway Routing Protocol (EIGRP), so the routing services on this switch can exchange routing updates with other devices.
Switch# show ip route	Indicates if the networks associated with the SVIs appear in the routing table.

How to Configure Inter-VLAN Routing on a Multilayer Switch

This topic describes how to configure a Switched Virtual Interface for inter-VLAN routing.

Configuring Inter-VLAN Routing Through an SVI

Cisco.com

```
Switch(config)#ip routing
```

- Enables IP routing on the switch

```
Switch(config)#interface vlan vlan-id
```

- Creates and enters interface configuration mode for SVI

```
Switch(config-if)#ip address ip-address mask
```

- Assigns an IP address to the SVI

```
Switch(config)#router ip_routing_protocol <options>
```

- Specifies the IP routing protocol if needed

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—5-6

To configure inter-VLAN routing on a Catalyst SVI, perform these steps.

Step	Description	Notes and Comments
1.	Enable IP routing on the router. <code>Switch(config)#ip routing</code>	
2.	Create the SVI interface or navigate to configuration mode for the interface <code>Switch(config)#interface VLAN vlan-id</code>	
3.	Assign an IP address to the SVI for the VLAN. <code>Switch(config-if)#ip address n.n.n.n subnet-mask</code>	
4.	(Optional) Specify an IP routing protocol. <code>Switch(config)#router ip_routing_protocol <options></code>	This step is necessary for the switch to exchange dynamic routing updates with other routing devices. The routing protocol specified may require additional options. Refer to the documentation for the routing protocol for further details.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **SVI is a VLAN of switch ports represented by one interface to the routing system.**
- **Routed switch interface has Layer 3 attributes.**
- **Specific commands are used to configure and verify routing on multilayer switch interfaces.**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-5-7

Module Summary

This topic summarizes the key points discussed in this module.

Summary

Cisco.com

- **An external router can be configured to route packets between the VLANs on a Layer 2 switch.**
- **CEF-based multilayer switching facilitates packet switching in hardware.**
- **Multilayer switches allow routing and the configuration of interfaces to pass packets between VLANs.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—5-1

The configuration of multiple Layer 2 VLANs requires that Layer 3 routing occur between those VLANs. This inter-VLAN routing can be provided external to a Layer 2 switch or within a multilayer switch through the configuration of Switch Virtual Interfaces and IP routing. When routing occurs within a Catalyst multilayer switch, Cisco Express Forwarding is deployed to facilitate Layer 3 switching through hardware-based tables, providing an optimal packet forwarding process.

References

For additional information, refer to these resources:

- Cisco Systems, Inc., *How to Choose the Best Router Switching Path for Your Network*.
http://www.cisco.com/en/US/partner/tech/tk827/tk831/technologies_white_paper09186a00800a62d9.shtml
- Cisco Systems, Inc., *Cisco Express Forwarding White Paper*.
http://www.cisco.com/warp/customer/cc/pd/iosw/iore/tech/cef_wp.htm

Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

- Q1) Which statement accurately describes CEF-based MLS? (Source: Deploying CEF-Based Multilayer Switching)
- A) All packets require minor processing.
 - B) All traffic switching is flow-based.
 - C) Packets are not forwarded based on flow.
 - D) Only the first packet in each flow is handled in software.
- Q2) In order to forward packets between VLANs at wire speed, what technology is required? (Source: Implementing Multilayer Switching)
- A) fast switching
 - B) CEF-based multilayer switching
 - C) an external router
 - D) STP
- Q3) If you want to perform inter-VLAN routing on a single connection to an external router, which two of the following must be configured on the router? (Choose two.) (Source: Describing Routing Between VLANs)
- A) subinterface
 - B) unique MAC addresses for each participating interface
 - C) a trunking encapsulation
 - D) STP
- Q4) On a multilayer switch, in order to change a switch port into a Layer 3-capable interface, what command must be issued? (Source: Enabling Routing Between VLANs on an MLS)
- A) **IP address**
 - B) **IP routing**
 - C) **Layer 3 enable**
 - D) **no switchport**

Module Self-Check Answer Key

- Q1) C
- Q2) B
- Q3) A, C
- Q4) D