

Building Cisco Multilayer Switched Networks

Volume 2

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 2

<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

<i>Configuring PortFast</i>	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
<i>Guarding Against Rogue STP Root Bridges</i>	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
<i>Configuring UplinkFast</i>	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
<i>Configuring BackboneFast</i>	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
<i>Configuring EtherChannel</i>	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

Implementing Redundancy in the Routing Layer

Overview

A network with high availability provides alternative means by which all infrastructure paths and key servers can be accessed at all times. The Hot Standby Routing Protocol (HSRP) is one of those software features configured to provide Layer 3 redundancy to network hosts. HSRP optimization provides immediate or link-specific failover as well as a recovery mechanism. Virtual Router Redundancy Protocol and Gateway Load Balancing Protocol are derivatives of HSRP, providing additional Layer 3 redundancy features such as load balancing. High availability is also accomplished by implementing a combination of device, link, or hardware component redundancy at strategic points in the network. Optimal fault tolerance is provided by redundant supervisor engines in multilayer switches that also host a variety of failover services in software. Another Cisco high availability feature is Single Router Mode, which allows failover to a second Multilayer Switch Feature Card (MSFC).

Module Objectives

Upon completing this module, you will be able to implement redundancy in the routing layer to improve and ensure end-to-end availability of network services. This ability includes being able to meet these objectives:

- Enable HSRP so that redundant routers can provide default gateway functionality
- Tune HSRP to failover quickly and with a tracked interface
- Configure Layer 3 redundancy with VRRP or GLBP so that load balancing is implemented in addition to router failover at the distribution layer
- Configure redundancy features so that a second supervisor engine or power supply provides failover capabilities within a Catalyst switch
- Identify technologies and best practices required to maintain high availability in the Campus Infrastructure module

Configuring Layer 3 Redundancy with HSRP

Overview

Hot Standby Routing Protocol (HSRP) uses a specific software process as it tracks router interfaces and provides a failover route in the event of primary link failure. Routing issues exist as we examine various means of providing redundancy for the default gateway of each segment. HSRP has very specific attributes that warrant further description, as does a delineation of HSRP operations on the network. HSRP interfaces transition through a series of states as they find their role in the capacity of active or standby HSRP router. Various commands are used to configure and verify the operation of HSRP.

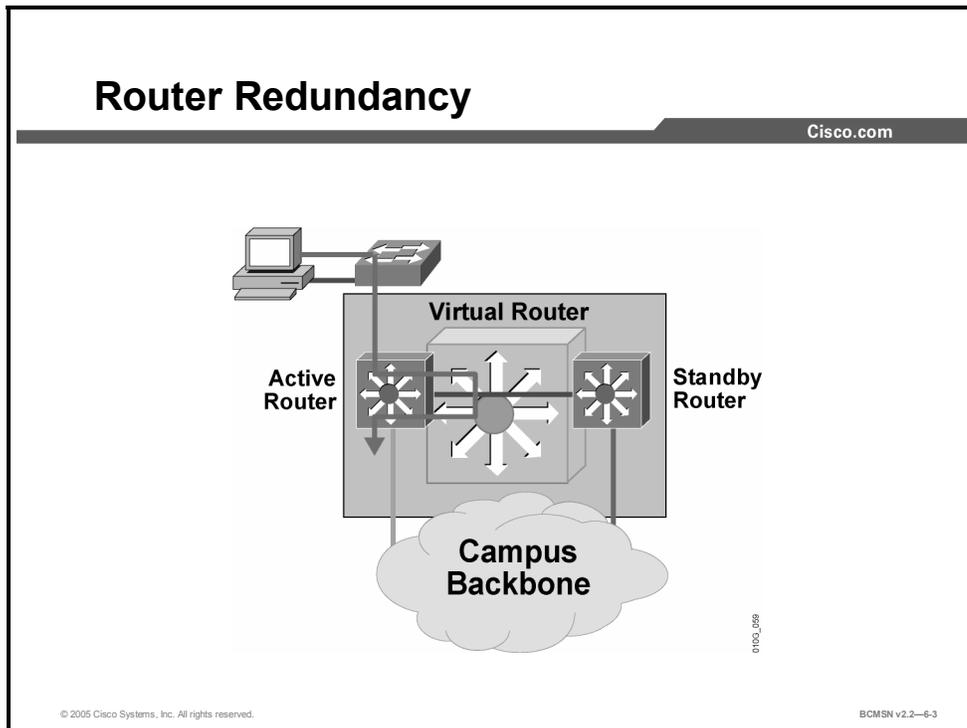
Objectives

Upon completing this lesson, you will be able to enable HSRP so that multiple routers can provide redundant paths off a local segment in the event of default gateway failure. This ability includes being able to meet these objectives:

- Identify how router device redundancy works
- Identify routing issues that occur using default gateways and proxy ARP
- Describe HSRP
- Identify how HSRP operates to provide nonstop path redundancy for IP
- Describe the HSRP states and their functions
- Identify the commands used to configure HSRP
- Enable HSRP

Identifying the Router Redundancy Process

This topic describes how router device redundancy works.

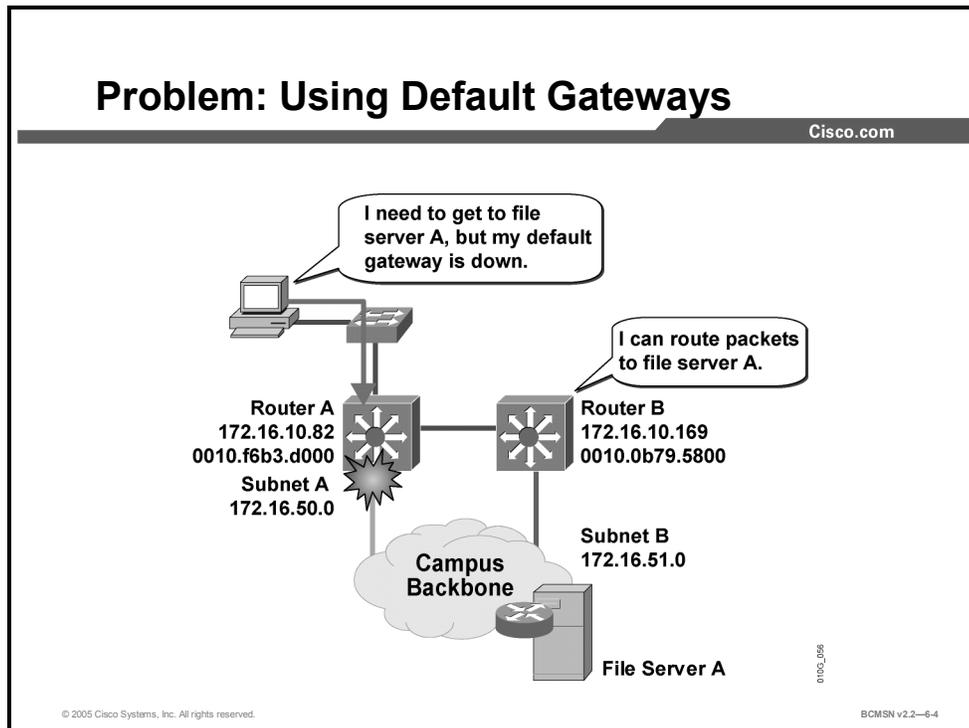


When router redundancy is configured, a virtual router is created. The IP address of the virtual router will be configured as the default gateway for the workstations on a specific IP segment. When frames are to be sent from the workstation to the default gateway, the workstation will use Address Resolution Protocol (ARP) to resolve the MAC address associated with the IP address of the default gateway. The ARP resolution will return the MAC address of the virtual router. Frames sent to the MAC address of the virtual router can then be physically processed by any active or standby router that is part of that virtual router group.

A protocol is used to identify two or more routers as the devices responsible for processing frames sent to the MAC or IP address of a single virtual router. Host devices send traffic to the address of the virtual router. The physical router that forwards this traffic is transparent to the end stations. The redundancy protocol provides the mechanism for determining which router should take the active role in forwarding traffic and determining when that role must be taken over by a standby router. The transition from one forwarding router to another is transparent to the end devices.

Routing Issues

This topic describes routing issues that occur when using default gateways and proxy ARP.



Using Default Gateways

When a default gateway is configured on most devices, there is no means by which to configure a secondary gateway, even if a second route exists to carry packets off the local segment.

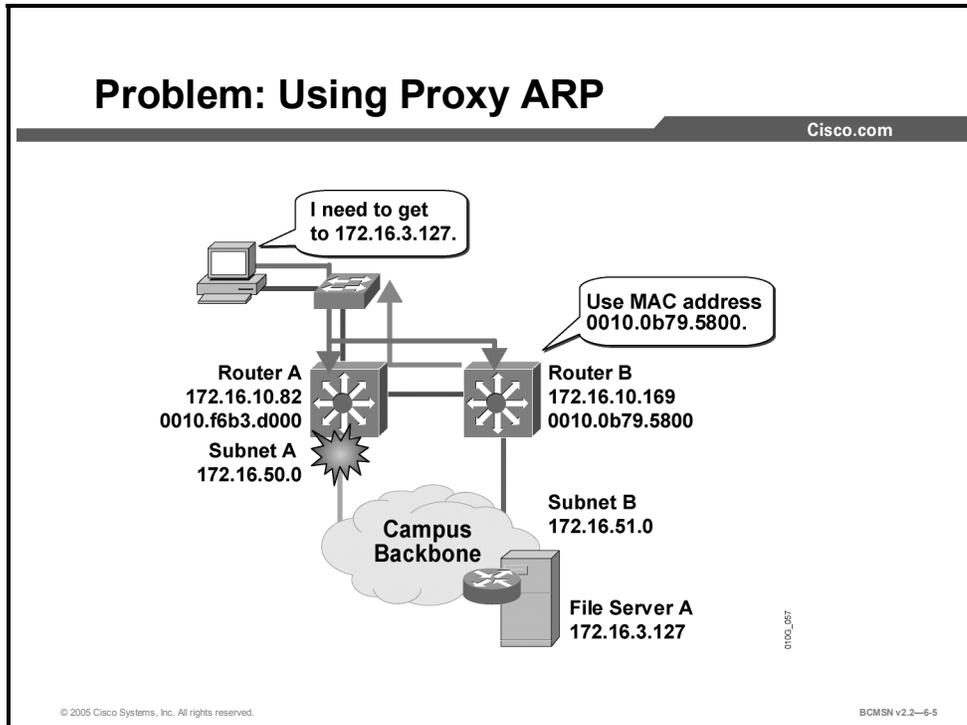
For example, primary and secondary paths between the Building Access submodule and the Building Distribution submodule provide continuous access in the event of a link failure at the Building Access layer. Primary and secondary paths between the Building Distribution layer and the Building Core layer provide continuous operations should a link fail at the Building Distribution layer.

In this example, router A is responsible for routing packets for subnet A, and router B is responsible for handling packets for subnet B. If router A becomes unavailable, routing protocols can quickly and dynamically converge and determine that router B will now transfer packets that would otherwise have gone through router A. Most workstations, servers, and printers, however, do not receive this dynamic routing information.

End devices are typically configured with a single default gateway IP address that does not change when network topology changes occur. If the router whose IP address is configured as the default gateway fails, the local device will be unable to send packets off the local network segment, effectively disconnecting it from the rest of the network. Even if a redundant router exists that could serve as a default gateway for that segment, there is no dynamic method by which these devices can determine the address of a new default gateway.

Using Proxy ARP

This subtopic describes proxy ARP.



Cisco IOS software runs proxy ARP to enable hosts that have no knowledge of routing options to obtain the MAC address of a gateway that is able to forward packets off the local subnet. For example, if the proxy ARP router receives an ARP request for an IP address that it knows is not on the same interface as the request sender, it will generate an ARP reply packet giving its own local MAC address as the destination MAC address of the IP address being resolved. The host that sent the ARP request sends all packets destined for the resolved IP address to the MAC address of the router. The router then forwards the packets toward the intended host, perhaps repeating this process along the way. Proxy ARP is enabled by default.

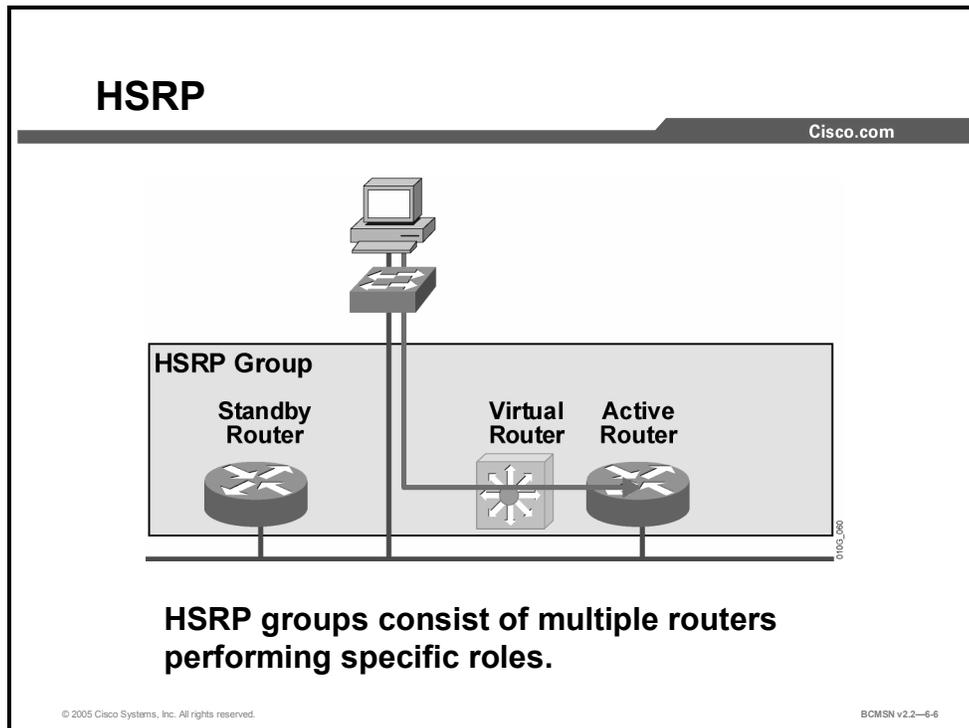
With proxy ARP, the end-user station behaves as if the destination device were connected to its own network segment. If the responsible router fails, the source end station continues to send packets for that IP destination to the MAC address of the failed router and the packets are therefore discarded.

Eventually, the proxy ARP MAC address will age out of the workstation's ARP cache. The workstation may eventually acquire the address of another proxy ARP failover router, but the workstation cannot send packets off the local segment during this failover time.

For further information on proxy ARP, refer to RFC 1027.

Hot Standby Router Protocol

This topic discusses HSRP operations.



HSRP defines a standby group of routers, with one router as the active one. HSRP provides gateway redundancy by sharing IP and MAC addresses between redundant gateways. The protocol consists of virtual MAC and IP addresses that are shared between two routers that belong to the same HSRP group. HSRP can optionally monitor both LAN and serial interfaces via a multicast protocol.

HSRP Terminology

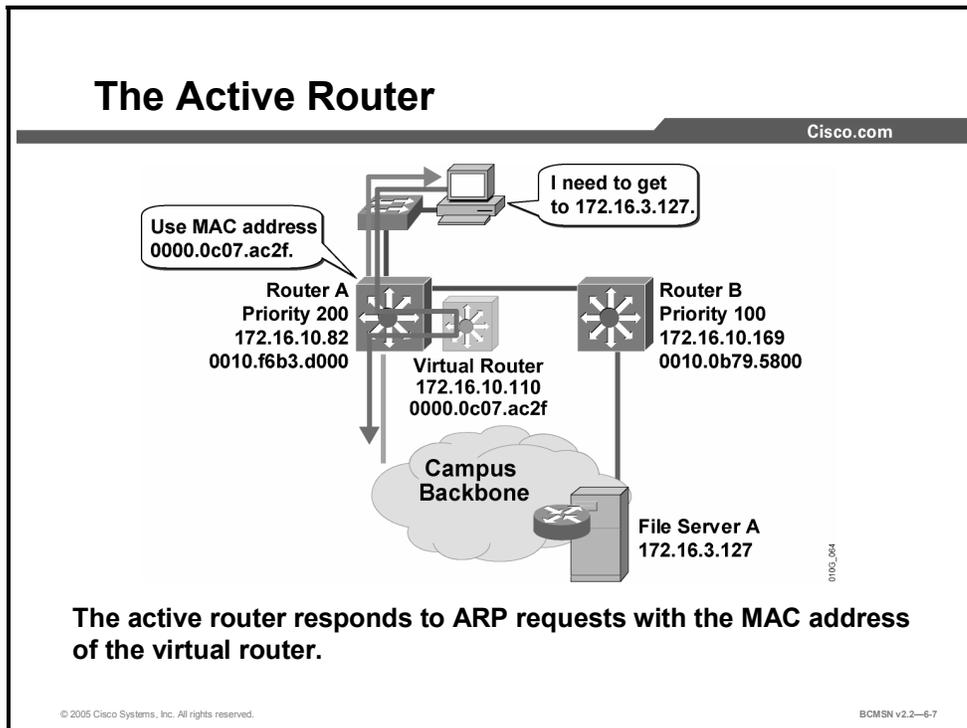
Term	Definition
Active router	The router that is currently forwarding packets for the virtual router
Standby router	The primary backup router
Standby group	The set of routers participating in HSRP that jointly emulate a virtual router
Hello interval time	The interval between successive HSRP hello messages from a given router
Hold interval time	The interval between the receipt of a hello message and the presumption that the sending router has failed

An HSRP group comprises these entities:

- One active router
- One standby router
- One virtual router
- Other routers

Identifying HSRP Operations

This topic describes how HSRP provides redundancy for the IP default gateway.



All the routers in an HSRP group have specific roles and interact in specific manners.

Virtual HSRP Router

The virtual router is simply an IP and MAC address pair that end devices have configured as their default gateway. The active router will process all packets and frames sent to the virtual router address. The virtual router processes no physical frames.

Active HSRP Router

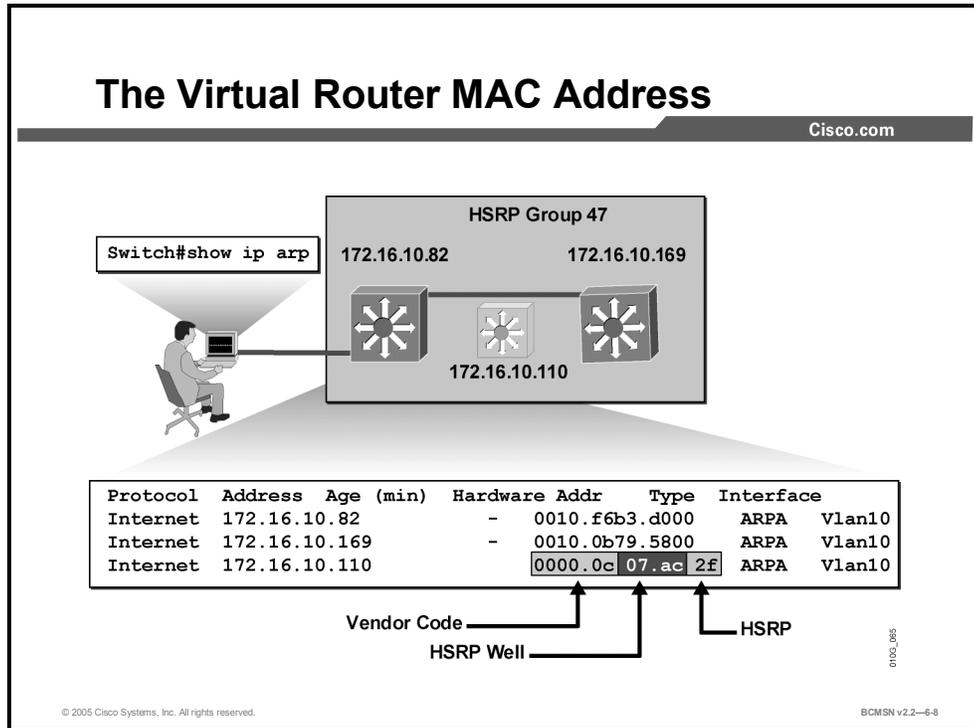
Within an HSRP group, one router is elected to be the active router. The active router physically forwards packets sent to the MAC address of the virtual router.

The active router responds to traffic for the virtual router. If an end station sends a packet to the virtual router MAC address, the active router receives and processes that packet. If an end station sends an ARP request with the virtual router IP address, the active router replies with the virtual router MAC address.

In this example, router A assumes the active role and forwards all frames addressed to the well-known MAC address of 0000.0c07.acxx, where xx is the HSRP group identifier.

ARP Resolution with HSRP

This subtopic describes ARP resolution with HSRP.



ARP establishes correspondences between network addresses, such as an IP address and a hardware Ethernet address. All devices sending packets over IP maintain a table of resolved addresses, including routers.

The IP address and corresponding MAC address of the virtual router are maintained in the ARP table of each router in an HSRP group. As shown in the figure, the command **show ip arp** displays the ARP cache on a router.

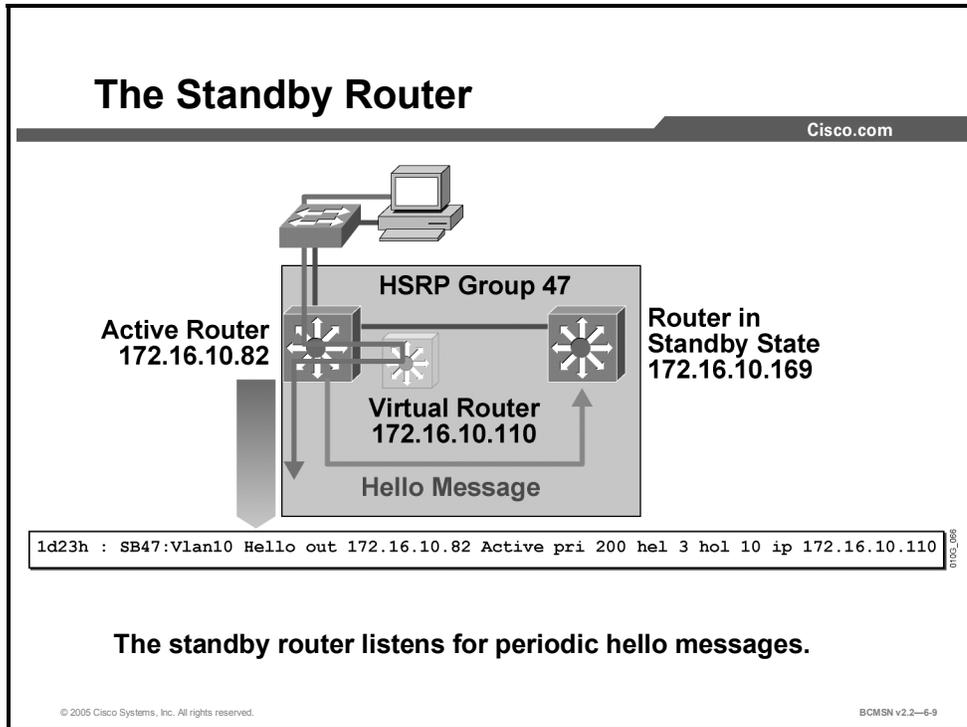
Show IP ARP Output Interpretation

Field	Definition
Protocol	Protocol for network address in the Address field
Address	The network address that corresponds to hardware address
Age (min)	Age, in minutes, of the cache entry
Hardware Addr	The MAC address that corresponds to network address
Type	Type of encapsulation:
Interface	Interface to which this address mapping has been assigned

In the example, the output displays an ARP entry for a router that is a member of HSRP group 47 in VLAN10. The virtual router for VLAN10 is identified as 172.16.10.110. The well-known MAC address that corresponds to this IP address is 0000.0c07.ac2f, where 2f is the HSRP group identifier for group 47. The HSRP group number is the standby group number (47) converted to hexadecimal (2f).

Standby and Other HSRP Routers in the Group

This subtopic describes the HSRP standby and other router roles in an HSRP group.

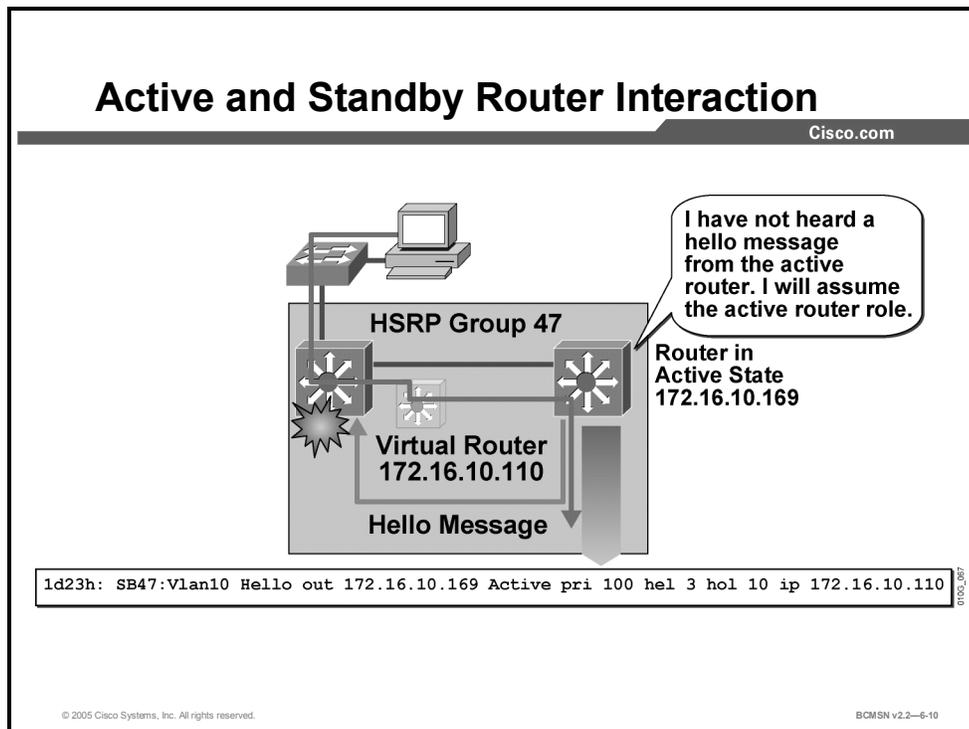


The function of the HSRP standby router is to monitor the operational status of the HSRP group and quickly assume packet-forwarding responsibility if the active router becomes inoperable. Both the active and standby routers transmit hello messages to inform all other routers in the group of their role and status.

An HSRP group may contain other routers that are group members but are not in an active or standby state. These routers monitor the hello messages sent by the active and standby routers to ensure that an active and standby router exist for the HSRP group of which they are a member. These routers do forward any packets addressed to their own specific IP addresses, but they do not forward packets addressed to the virtual router. These routers issue speak messages at every hello interval time.

HSRP Active and Standby Router Interaction

This subtopic describes the interaction between the active and standby routers.



When the active router fails, the other HSRP routers stop seeing hello messages from the active router. The standby router will then assume the role of the active router. If other routers are participating in the group, they then contend to be the new standby router.

In the event that both the active and standby routers fail, all routers in the group contend for the active and standby router roles.

Because the new active router assumes both the IP and MAC addresses of the virtual router, the end stations see no disruption in service. The end-user stations continue to send packets to the virtual router MAC address, and the new active router delivers the packets to the destination.

HSRP States

This topic discusses the different HSRP states.

HSRP States

Cisco.com

An HSRP router can be in one of five different states:

- **Initial**
- **Listen**
- **Speak**
- **Standby**
- **Active**

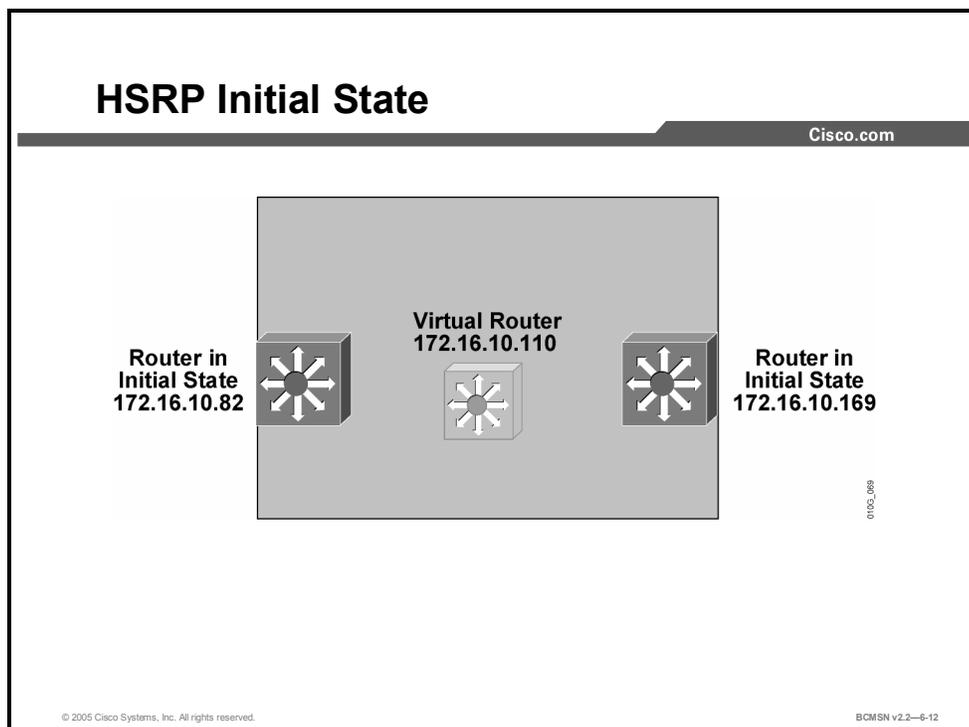
© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-11

A router in an HSRP group can be in one of the following states: initial, listen, speak, standby, or active.

When a router exists in one of these states, it performs the actions required for that state. Not all HSRP routers will transition through all states. For example, a router that is not the standby or active router will not enter the standby or active states.

HSRP Initial State

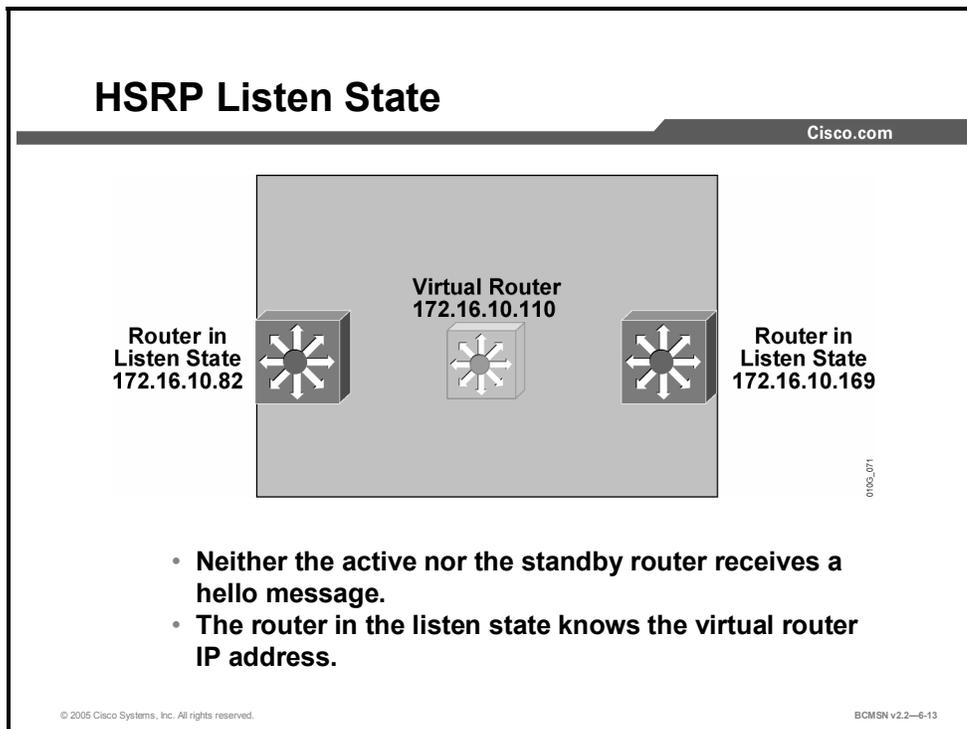
This subtopic describes the HSRP initial state.



All routers begin in the initial state. This is the starting state and indicates that HSRP is not running. This state is entered via a configuration change or when an interface is initiated.

HSRP Listen State

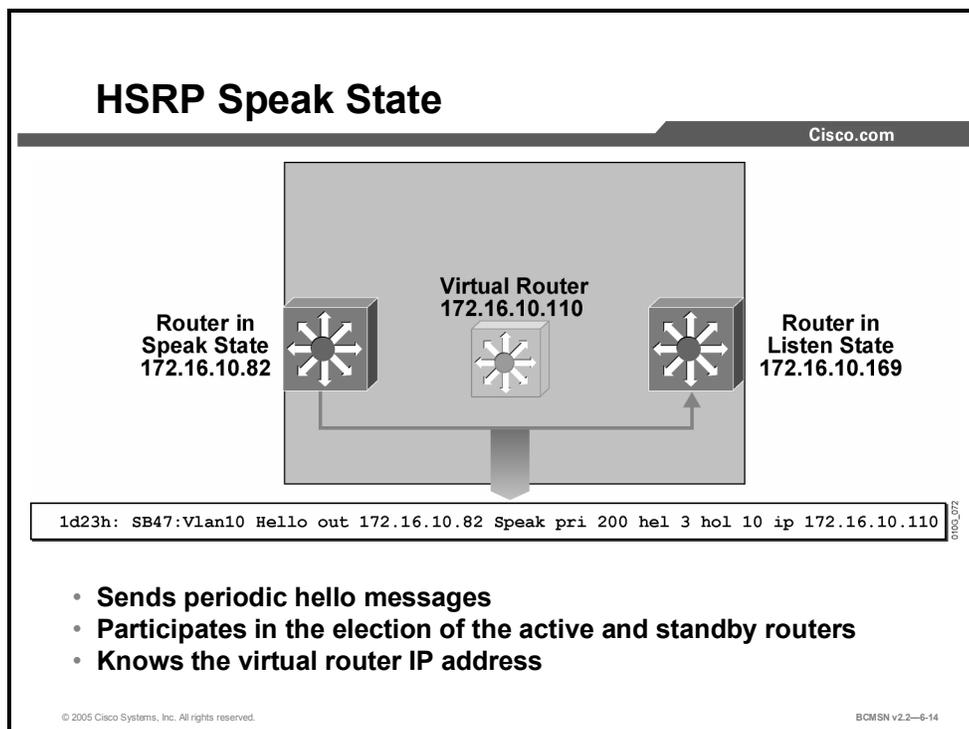
This subtopic describes the HSRP listen state.



In the listen state, the router knows the IP address of the virtual router but is neither the active router nor the standby router. The router listens for hello messages from those routers for a duration called the hold time, which can be configured. The purpose of this listening interval is to determine if there are active or standby routers for the group. Then this router will join the HSRP group, based on its configuration.

HSRP Speak State

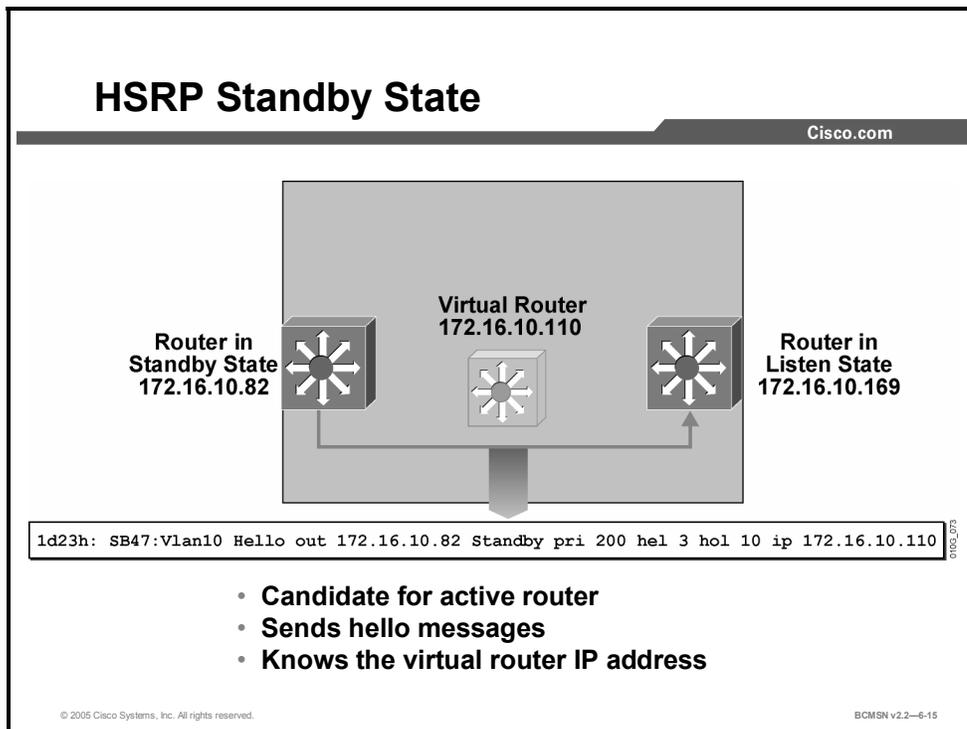
This subtopic describes the HSRP speak state.



In the speak state, the router sends periodic hello messages and is actively participating in the election of the active router or standby router or both. A router cannot enter the speak state unless the router has the IP address of the virtual router. The router will remain in the speak state unless it becomes an active or standby router.

Standby State

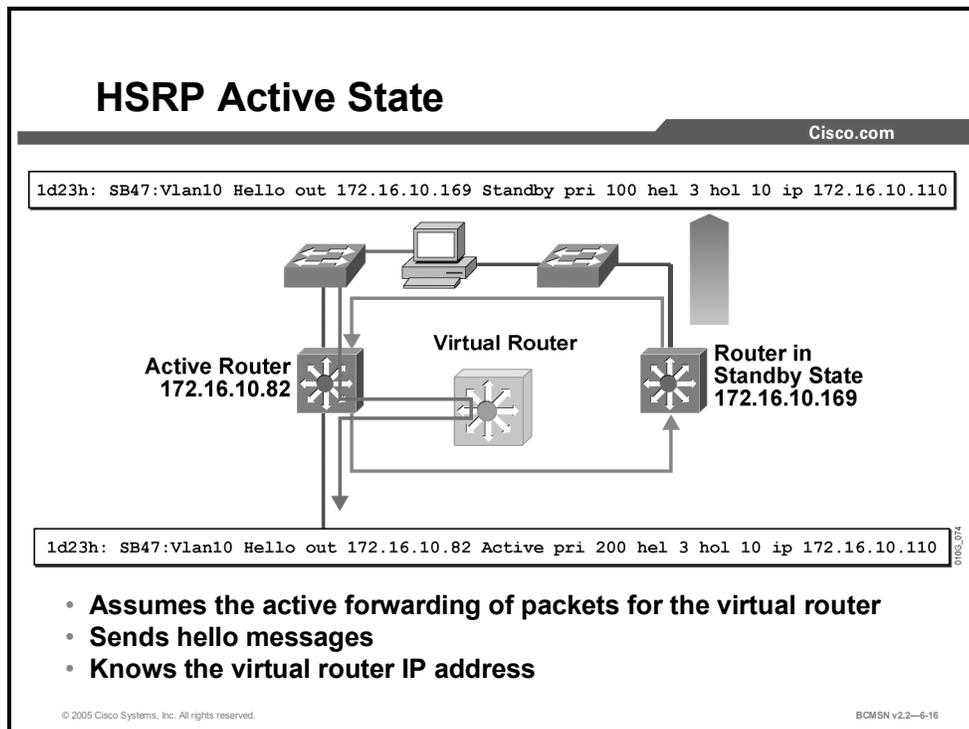
This subtopic describes the HSRP standby state.



In the standby state, because the router is a candidate to become the next active router, it sends periodic hello messages. It will also listen for hello messages from the active router. There will be only one standby router in the HSRP group.

Active State

This subtopic describes the HSRP active state.



In the active state, the router is currently forwarding packets that are sent to the virtual MAC address of the group. It also replies to ARP requests directed to the virtual router's IP address. The active router sends periodic hello messages. There must be one active router in each HSRP group.

HSRP Configuration Commands

This topic lists the commands used to configure and verify HSRP.

About HSRP Configuration Commands

Cisco.com

Configure

- **standby 47 ip 10.1.1.1**
- **standby 47 priority 150**

Verify

- **show running-config**
- **show standby**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—6-17

Commands Used to Configure and Verify HSRP

Command	Description
Switch(config-if)# standby group-number ip ip-address	Configures HSRP on this interface for this group number. IP address is that of the virtual gateway. Default group number is 0.
Switch(config-if)# no standby group-number ip ip-address	Disables HSRP on the interface
Switch(config-if)# standby group-number priority priority-value	Sets the priority of this router to <i>value</i> for this HSRP group. Range is 0-255. Default priority value is 100.
Switch(config-if)# no standby group-number priority	Removes any set priority value. Returns priority to default of 100.
Switch# show running-config	Displays HSRP parameters configured on each interface
Switch# show standby [interface] [group] [brief]	Show standby is all that is required. Use other commands to minimize output.

How to Enable HSRP

This topic describes how to enable and verify HSRP operations.

Configuring an HSRP Standby Interface

Cisco.com

```
Switch#show standby vlan 10
interface Vlan10
ip address 172.16.10.82 255.255.255.0
no ip redirects
standby 47 ip 172.16.10.110
```

Virtual Router IP Address
Standby Group Number

↓ ↓

```
Switch(config-if)#standby 47 ip 172.16.10.110
```

Enabling HSRP on a Cisco router interface automatically disables Internet Control Message Protocol redirects.

© 2005 Cisco Systems, Inc. All rights reserved.
BCMSN v2.2-6-18

Steps for Configuring HSRP

Step	Action
1.	Configure HSRP group on an interface.
2.	Verify the configuration.
3.	Establish priorities.
4.	Verify the HSRP standby priority.
5.	Verify all HSRP operations.

Configure HSRP Group on an Interface

```
Switch(config-if)#standby group-number ip ip-address
```

Variable	Definition
<i>group-number</i>	(Optional) Indicates the HSRP group to which this interface belongs. Specifying a unique group number in the standby commands enables the creation of multiple HSRP groups. The default group is 0.
<i>ip-address</i>	Indicates the IP address of the virtual HSRP router.

While running HSRP, the end-user stations must not discover the actual MAC addresses of the routers in the standby group. Any protocol that informs a host of a router actual address must be disabled. To ensure that the actual addresses of the participating HSRP routers are not discovered, enabling HSRP on a Cisco router interface automatically disables Internet Control Message Protocol (ICMP) redirects on that interface.

After the **standby ip** command is issued, the interface changes to the appropriate state. When the router successfully executes the command, the router issues an HSRP message.

To remove an interface from an HSRP group, enter the **no standby group ip** command.

Verifying HSRP Configuration

This subtopic shows the command used to verify the HSRP configuration.

Displaying the Standby Brief Status

Cisco.com

```
Switch#show standby brief
          P indicates configured to preempt.

```

Interface	Grp	Prio	P	State	Active addr	Standby addr	Group addr
Vl11	11	110		Active	local	172.16.11.114	172.16.11.115

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-6-19

The following example states that interface VLAN10 is a member of HSRP group 47, that the virtual router IP address for the group is 172.16.10.110, and that ICMP redirects are disabled.

```
Switch#show running-config
Building configuration...
```

Current configuration:

```
!
(text deleted)
interface Vlan10
ip address 172.16.10.82 255.255.255.0
no ip redirects
standby 47 ip 172.16.10.110
!
```

Another means of verifying the HSRP configuration is with this command:

```
Switch# show standby brief
```

It displays abbreviated information about the current state of all HSRP operations on this device.

Establish HSRP Priorities

In this step the HSRP priority is set and verified.

Configuring HSRP Standby Priority

Cisco.com

```
Switch#show standby vlan 10
interface Vlan10
ip address 172.16.10.82 255.255.255.0
no ip redirects

standby 47 priority 150
standby 47 ip 172.16.10.110
```

Assigned Priority →
Standby Group Number →

```
Switch(config-if)#standby 47 priority 150
```

- The router in an HSRP group with the highest priority becomes the forwarding router.
- The default priority is 100.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-20

Each standby group has its own active and standby routers. The network administrator can assign a priority value to each router in a standby group, allowing the administrator to control the order in which active routers for that group are selected.

To set the priority value of a router, enter this command in interface configuration mode:

```
Switch(config-if)#standby group-number priority priority-value
```

Variable	Definition
<i>group-number</i>	Indicates the HSRP group. This number can be in the range of 0 to 255.
<i>priority-value</i>	Indicates the number that prioritizes a potential hot standby router. The range is 0 to 255; the default is 100.

During the election process, the router in an HSRP group with the highest priority becomes the forwarding router. In the case of a tie, the router with the highest configured IP address will become active.

To reinstate the default standby priority value, enter the **no standby priority** command.

Verify the HSRP Standby Priority

The following example states that interface VLAN10 has a priority value of 150 in HSRP group 47. If this priority value is the highest number in that HSRP group, the routing device on which this interface resides is the active router for that group.

```
Switch#show running-config
Building configuration...

Current configuration:
!
(text deleted)
interface Vlan10
ip address 172.16.10.32 255.255.255.0
no ip redirects
standby 47 priority 150
standby 47 ip 172.16.10.110
```

To display the status of the HSRP router, enter one of these commands:

```
Switch#show standby [interface [group]] [active | init | listen |
standby] [brief]
```

```
Switch#show standby delay [type-number]
```

If the optional interface parameters are not indicated, the **show standby** command displays HSRP information for all interfaces.

Verify All HSRP Operations

This example shows the output of the **show standby** command:

```
Switch#show standby Vlan10 47
Vlan11 - Group 47
  Local state is Active, priority 150, may preempt
  Hellotime 3 holdtime 10
  Next hello sent in 00:00:02.944
  Hot standby IP address is 172.16.10.110 configured
  Active router is local
  Standby router is 172.16.10.82 expires in 00:00:08
  Standby virtual mac address is 0000.0c07.ac2f
  Tracking interface states for 1 interface, 1 up:
    Up Vlan51 Priority decrement: 40
```

This is an example of the output resulting when you specify the **brief** parameter:

```
Switch#show standby brief
Interface  Grp   Prio P State   Active addr   Standby addr
Group addr
Vl10      47   150  P Active   local         172.16.10.82
172.16.11.10
Vl12      12   100   Standby   172.16.102.82 local
172.16.12.10
```

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- Router redundancy is needed in the campus network.
- Routing issues can occur when using default gateways and proxy ARP.
- HSRP provides router redundancy to end devices.
- HSRP operates to provide nonstop path redundancy for IP.
- An HSRP-enabled router will exist in a specific state or transition through a series of states.
- Specific commands are used to configure and verify HSRP.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-21

Optimizing HSRP

Overview

Hot Standby Router Protocol (HSRP) has options that allow it to be configured for expedited failover, recovery from failover or to specify which interface is to be monitored for HSRP failover. Specific commands are used to optimize and tune HSRP operations for greatest failover resiliency. There is also a set of commands for verifying and debugging HSRP general and optimized operations.

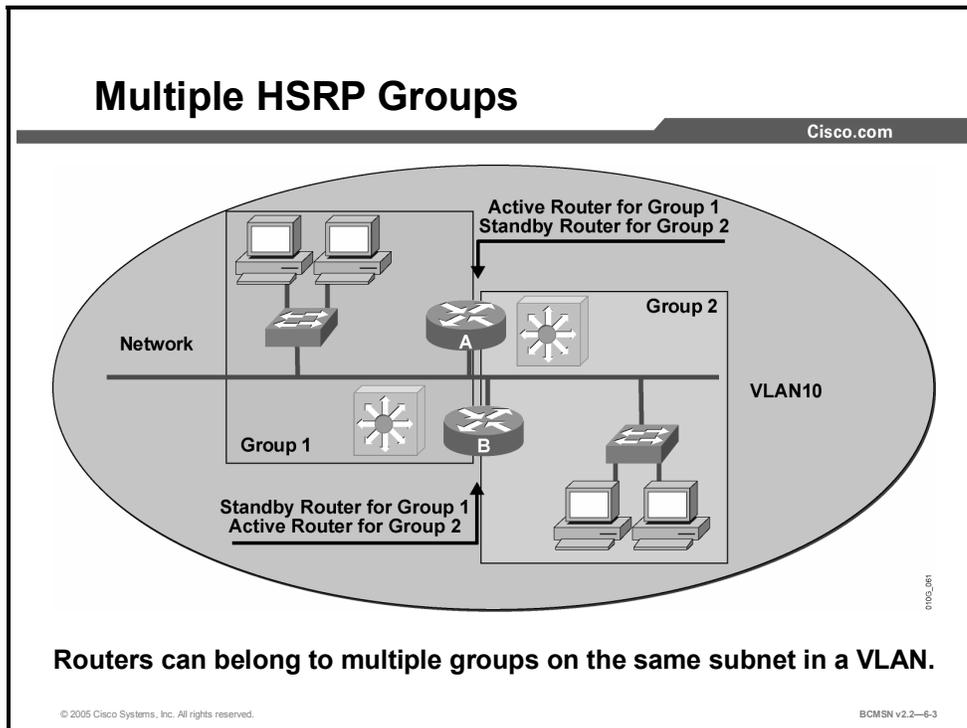
Objectives

Upon completing this lesson, you will be able to tune HSRP so that failover and recovery are immediate. This ability includes being able to meet these objectives:

- Describe how a single router can be a member of multiple HSRP groups
- Identify the options that can be configured to optimize HSRP
- Identify the commands used to optimize HSRP
- Determine which HSRP operations require tuning
- Identify and execute the commands required to debug HSRP operations

Load Sharing

This topic describes how and why a single router might be a member of multiple HSRP groups.



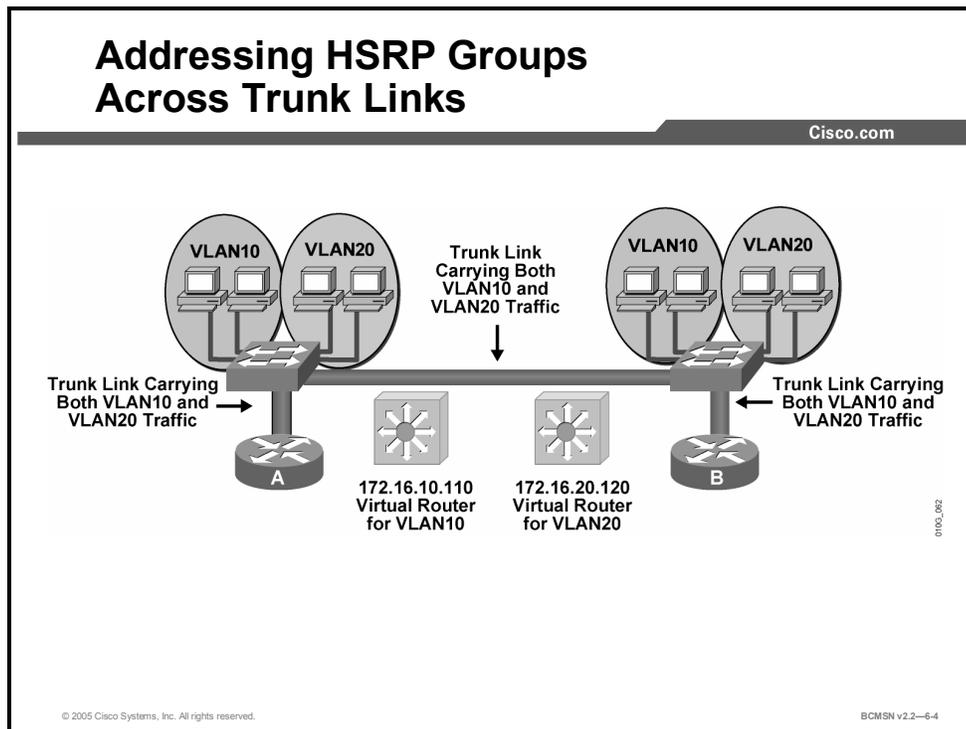
To facilitate load sharing, a single router may be a member of multiple HSRP groups on a single segment. Multiple standby groups further enable redundancy and load sharing within networks. While a router is actively forwarding traffic for one HSRP group, the router can be in standby or listen state for another group. Each standby group emulates a single virtual router. There can be up to 255 standby groups on any LAN.

Caution Increasing the number of groups in which a router participates increases the load on the router. This can have an impact on the performance of the router.

In the figure, both router A and router B are members of groups 1 and 2. However, router A is the active forwarding router for group 1 and the standby router for group 2. Router B is the active forwarding router for group 2 and the standby router for group 1.

Addressing HSRP Groups Across Trunk Links

This subtopic describes how HSRP devices can provide primary paths for some data and backup paths for other data when VLANs and trunks are deployed in the network.



Routers can simultaneously provide redundant backup and perform load sharing across different IP subnets.

For each standby group, an IP address and a single well-known MAC address with a unique group identifier is allocated to the group.

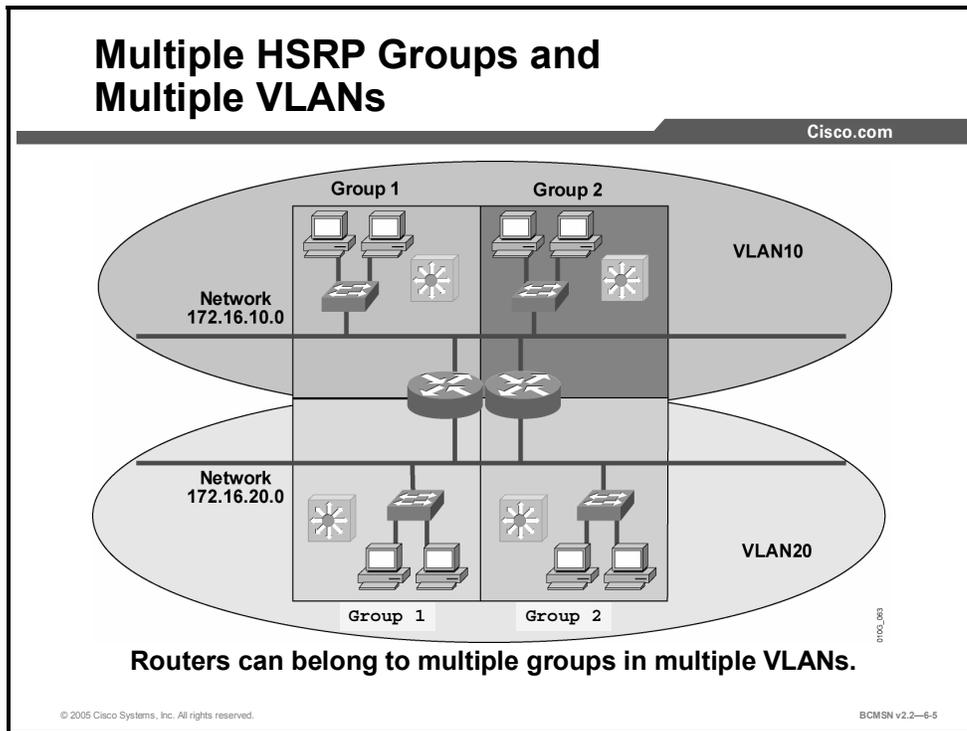
The IP address of a group is in the range of addresses belonging to the subnet that is in use on the LAN. However, the IP address of the group must differ from the addresses allocated as interface addresses on all routers and hosts on the LAN, including virtual IP addresses assigned to other HSRP groups.

In the figure, two HSRP-enabled routers participate in two separate VLANs, using Inter-Switch Link (ISL) or 802.1Q. Running HSRP over trunking allows users to configure redundancy among multiple routers that are configured as front ends for VLAN IP subnets. By configuring HSRP over ISLs, users can eliminate situations in which a single point of failure causes traffic interruptions. This feature inherently provides some improvement in overall networking resilience by providing load balancing and redundancy capabilities between subnets and VLANs.

Note A route processor theoretically can support up to 32,650 subinterfaces; however, the actual number of supported interfaces is limited by the capacity of the route processor and the number of VLANs.

Supporting Multiple Subnets with Multiple HSRP Groups

This subtopic describes how to optimize the use of resources by allowing multiple routers to provide load balancing and redundancy.



Routers can belong to multiple groups within multiple VLANs. As members of multiple hot standby groups, routers can simultaneously provide redundant backup and perform load sharing across different IP subnets.

Although multiple routers can exist in an HSRP group, only the active router forwards the packets sent to the virtual router.

A single group can also be configured to host the virtual router of multiple IP subnets. For example, the default gateway virtual addresses of 172.16.10.1 and 172.16.20.1 can both be associated with a single HSRP group number.

HSRP Optimization Options

This topic describes the options that can be used to tune HSRP.

HSRP Optimization Options

Cisco.com

These items can be configured to optimize HSRP:

- **HSRP standby preempt**
- **Hello message timers**
- **HSRP interface tracking**

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

Three options are available to optimize the operation of HSRP in the campus network.

HSRP Standby Preempt

In the event of an active router failure, a standby router will assume the role of active router. When the original active router comes back online, this new active router will retain the role of active router even when the former active router has a higher priority. The preempt option allows a router to regain its role of active router even if there is an existing active router on the segment.

HSRP Hello Message Timer Adjustment

An HSRP-enabled router sends hello messages to indicate that the router is running and is capable of becoming either the active or the standby router. The hello message contains the priority of the router as well as hellotime and holdtime parameter values. The hellotime parameter value indicates the interval between the hello messages that the router sends. The holdtime parameter value indicates the amount of time that the current hello message is considered valid. The standby timer includes an “msec” parameter for faster failovers.

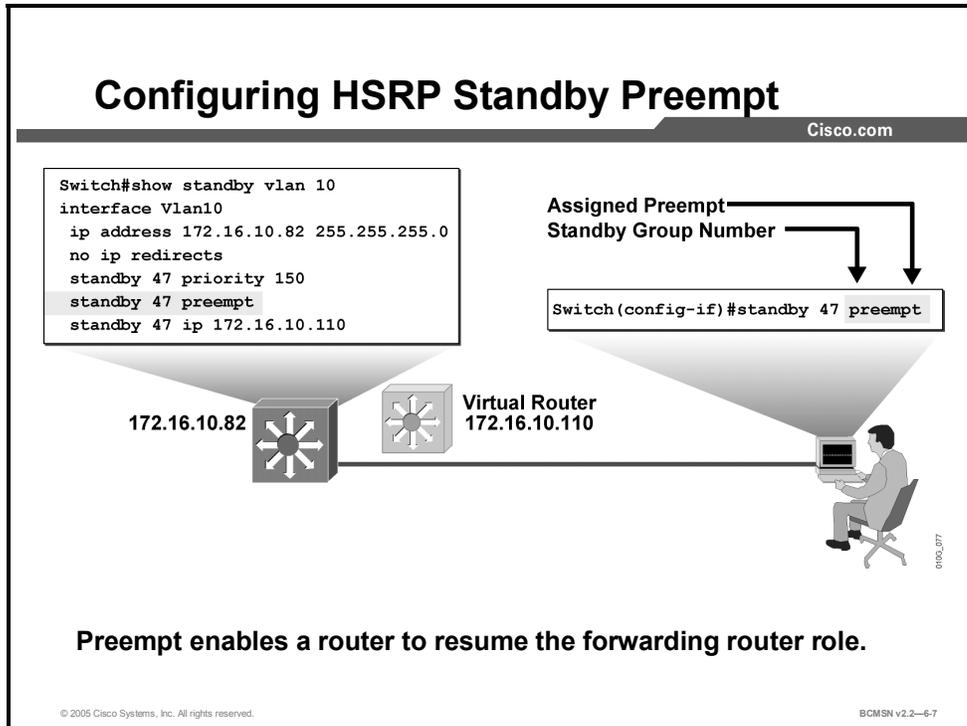
HSRP Interface Tracking

In some situations, it is not the status of the HSRP interface itself that could cause an HSRP failover but the failure of another route used by the active HSRP router. For example, if the active router has a WAN link back to corporate and that WAN link goes down, it would be desirable for the LAN interface of that router to give up its role as active HSRP router so it could be assumed by a standby router with a functioning link back to corporate. In this case,

HSRP would be configured to track the status of the WAN interface and the LAN interface on the same router would relinquish its role as active router if the WAN interface went down.

HSRP Standby Preempt

The standby router automatically assumes the active router role when the active router fails or is removed from service. This new active router remains the forwarding router even when the former active router with the higher priority regains service in the network.



The former active router can be configured to resume the forwarding router role from a router with a lower priority. To enable a router to resume the forwarding router role, enter the following command in interface configuration mode:

```
Switch(config-if)#standby [group-number] preempt [{delay} [minimum  
delay]  
[sync delay]]
```

When the **standby preempt** command is issued, the interface changes to the appropriate state.

To remove the interface from preemptive status, enter the **no standby group preempt** command.

Example: Displaying HSRP Preempt

The following example states that interface VLAN10 is configured to resume its role as the active router in HSRP group 47, assuming that interface VLAN10 on this router has the highest priority in that standby group.

```
Switch#show running-config
Building configuration...
```

```
Current configuration:
!
```

```
(text deleted)
interface Vlan10
ip address 172.16.10.82 255.255.255.0
no ip redirects
standby 47 priority 150
standby 47 preempt
standby 47 ip 172.16.10.110
```

Hello Message Timers

An HSRP-enabled router sends hello messages to indicate that the router is running and is capable of becoming either the active or the standby router.

Configuring the Hello Message Timers

Cisco.com

```
Building configuration...

Current configuration:
(text deleted)
!
interface Vlan10
 ip address 172.16.10.82 255.255.255.0
 no ip redirects
 standby 47 timers 5 15
 standby 47 ip 172.16.10.10
```

Holdtime
Hellotime

Switch(config-if)#standby 47 timers 5 15

The holdtime parameter value should be at least three times the value of the hellotime parameter.

© 2005 Cisco Systems, Inc. All rights reserved.
BCMSN v2.2-6-8

The hello message contains the priority of the router and also hellotime and holdtime parameter values. The hellotime parameter value indicates the interval between the hello messages that the router sends. The holdtime parameter value indicates the amount of time that the current hello message is considered valid.

If an active router sends a hello message, receiving routers consider that hello message to be valid for one holdtime. The holdtime value should be at least three times the value of the hellotime. The holdtime value must be greater than the value of the hellotime.

Both the hellotime and the holdtime parameters are configurable. To configure the time between hello messages and the time before other group routers declare the active or standby router to be nonfunctioning, enter this command in interface configuration mode:

```
Switch(config-if)#standby group-number timers hellotime holdtime
```

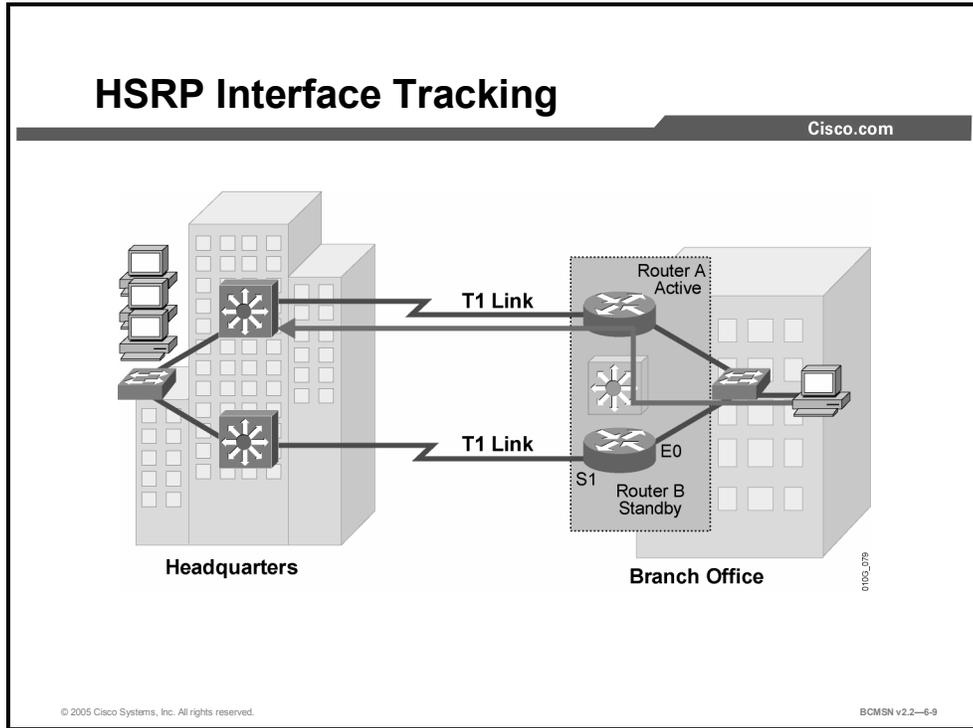
Standby Message Timer Configuration Options

Variable	Description
<i>group-number</i>	(Optional) Group number on the interface to which the timers apply. The default is 0.
<i>hellotime</i>	Hello interval in seconds. This is an integer from 1 through 255. The default is 3 seconds.
<i>holdtime</i>	Time, in seconds, before the active or standby router is declared to be down. This is an integer from 1 through 255. The default is 10 seconds.

To reinstate the default standby timer values, enter the **no standby group timers** command.

HSRP Interface Tracking

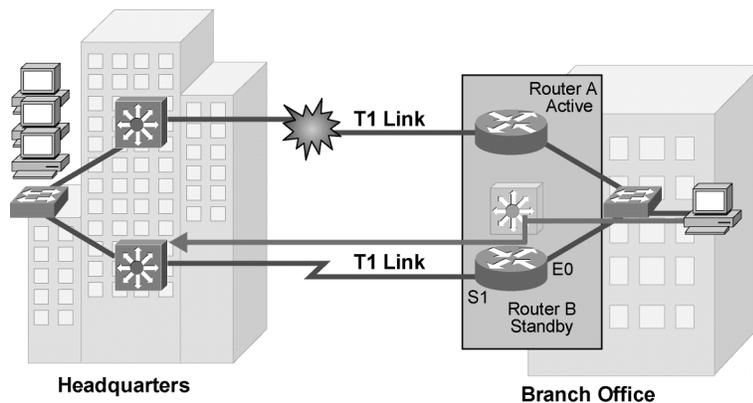
In some situations, the status of an interface directly affects which router needs to become the active router. This is particularly true when each of the routers in an HSRP group has a different path to resources within the campus network.



In this example, router A and router B reside in a branch office. These two routers each support a T1 link to headquarters. Router A has the higher priority and is the active forwarding router for standby group 47. Router B is the standby router for that group. Routers A and B are exchanging hello messages through their E0 interfaces.

HSRP Interface Tracking (Cont.)

Cisco.com



The T1 link between the active forwarding router for the standby group and headquarters experiences a failure. Without HSRP enabled, router A would detect the failed link and send an Internet Control Message Protocol (ICMP) redirect to router B. However, when HSRP is enabled, ICMP redirects are disabled. Therefore, neither router A nor the virtual router sends an ICMP redirect. In addition, although the S1 interface on router A is no longer functional, router A still communicates hello messages out interface E0, indicating that router A is still the active router. Packets sent to the virtual router for forwarding to headquarters cannot be routed.

Interface tracking enables the priority of a standby group router to be automatically adjusted, based on availability of the interfaces of that router. When a tracked interface becomes unavailable, the HSRP priority of the router is decreased. The HSRP tracking feature reduces the likelihood that a router with an unavailable key interface will remain the active router.

In this example, the E0 interface on router A tracks the S1 interface. If the link between the S1 interface and headquarters fails, the router automatically decrements the priority on that interface and stops transmitting hello messages out interface E0. Router B assumes the active router role when no hello messages are detected for the specific holdtime period.

Configuring HSRP Tracking

To configure HSRP tracking, enter the command in the figure in interface configuration mode.

Configuring HSRP Tracking

Cisco.com

```
Switch(config-if)#standby [group-number] track type number  
[interface-priority]
```

- **Configures HSRP tracking**

```
Switch(config)#interface vlan 10  
Switch(config-if)#standby 1 track GigabitEthernet 0/7 50  
Switch(config-if)#standby 1 track GigabitEthernet 0/8 60
```

- **Example of HSRP tracking**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-6-11

HSRP Tracking Configuration Arguments

Variable	Description
<i>group-number</i>	(Optional) Indicates the group number on the interface to which the tracking applies. The default number is 0.
<i>type</i>	Indicates the interface type (combined with the interface number) that will be tracked.
<i>number</i>	Indicates the interface number (combined with the interface type) that will be tracked.
<i>interface-priority</i>	(Optional) Indicates the amount by which the hot standby priority for the router is decremented when the interface becomes disabled. The priority of the router is incremented by this amount when the interface becomes available. The default value is 10.

To disable interface tracking, enter the **no standby group track** command.

The command to configure HSRP tracking on a multilayer switch is the same as on the external router, except that the interface type can be identified as a switch virtual interface (**vlan** followed by the *vlan number* assigned to that interface) or by a physical interface.

The internal routing device uses the same command as the external routing device to disable interface tracking.

Multiple tracking statements may be applied to an interface. For example, this may be useful if the currently active HSRP interface will only relinquish its status upon the failure of two (or more) tracked interfaces.

Tuning HSRP Operations

This topic describes which HSRP operations require tuning in a network.

Tuning HSRP

Cisco.com

- **Configure preempt delay timer.**
- **Preempt occurs only after full router boot.**
- **Preempt decreases time for network convergence.**

The chart, titled 'Test tool timeout—30 seconds', shows 'Time to converge in seconds' on the y-axis (0 to 35) and two categories on the x-axis: 'No Preempt Delay' and 'Preempt Delay Tuned'. For 'No Preempt Delay', five bars represent different IOS versions: 3550 IOS (~28s), 2950 IOS (~28s), 4006 (CatOS) (~28s), 4507 (IOS) (~28s), and 6500 (CatOS) (~28s). For 'Preempt Delay Tuned', the bars are significantly shorter, with a callout box indicating 'More than 30 seconds of delay or loss tuned away' pointing to the bars. A legend on the right lists the IOS versions: 3550 IOS, 2950 IOS, 4006 (CatOS), 4507 (IOS), 6500 (CatOS), and 6500 (IOS).

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-12

HSRP timers can be adjusted to tune the performance of HSRP on distribution devices, thereby increasing their resilience and reliability in routing packets off the local VLAN.

Subsecond Failover

The HSRP hello and holdtime can be set to millisecond values so that HSRP failover occurs in less than 1 second. Here is an example:

```
Switch(config-if)#standby 1 timers msec 200 msec 750
```

Preempt Time Aligned with Router Boot Time

Preempt is an important feature of HSRP that allows the primary router to resume the active role when it comes back online after a failure or maintenance event. Preemption is a desired behavior as it forces a predictable routing path for the VLAN during normal operations and ensures that the Layer 3 forwarding path for a VLAN parallels the Layer 2 Spanning Tree Protocol (STP) forwarding path whenever possible.

When a preempting device is rebooted, HSRP preempt communication should not begin until the distribution switch has established full connectivity to the rest of the network. This allows the routing protocol convergence to occur more quickly once the preferred router is in an active state. To accomplish this, measure the system boot time and set the HSRP preempt delay to a value 50 percent greater than the boot time. This ensures that the primary distribution switch establishes full connectivity to the network before HSRP communication occurs.

For example, if the boot time for the distribution device is 120 seconds, the preempt configuration would appear as follows:

```
standby 1 preempt  
standby 1 preempt delay minimum 180
```

HSRP debug Commands

This topic describes the commands used to debug HSRP operations.

About the HSRP Debug Command

Cisco.com

- **debug standby events**
- **debug standby terse**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-13

The following commands are used to debug HSRP operation:

Command	Description
Switch# debug standby [errors] [events] [packets]	Displays all state changes to HSRP, including all hello packets. Arguments minimize output.
Switch# debug standby terse	Displays all HSRP errors, events, and packets, <i>except</i> hello and advertisement packets.

References

For additional information, refer to this resource:

http://cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a0080094a91.shtml#hsrpdebug

How to Debug HSRP Operations

This topic describes how to use the debug command to view and verify HSRP operations.

Debugging HSRP

Cisco.com

```
Switch#debug standby

*Mar 1 00:22:30.443: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:22:32.019: SB11: V111 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:22:33.331: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:22:34.927: SB11: V111 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:22:36.231: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:22:37.823: SB11: V111 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:22:39.163: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:22:40.735: SB11: V111 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:22:42.119: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:22:43.663: SB11: V111 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
*Mar 1 00:22:45.067: SB11: V111 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar 1 00:22:46.567: SB11: V111 Hello in 172.16.11.112 Standby pri 50 ip 172.16.11.115
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-6-14

The Cisco IOS implementation of HSRP supports the **debug** command. Enabling debug displays HSRP state changes and debug output regarding the transmission and receipt of HSRP packets. To enable HSRP debugging, enter this command in privileged EXEC mode:

```
Switch#debug standby
```

Caution Because debugging output is assigned high priority in the CPU process, this command can render the system unusable.

Example: HSRP Debugging on Negotiation for Role of Active Router

This example displays the **debug standby** command output as the DSW111 router with the IP address 172.16.1.111 initializes and negotiates for the role of active router:

```
*Mar 8 20:34:10.221: SB11: V111 Init: a/HSRP enabled
*Mar 8 20:34:10.221: SB11: V111 Init -> Listen
*Mar 8 20:34:20.221: SB11: V111 Listen: c/Active timer expired
(unknown)
*Mar 8 20:34:20.221: SB11: V111 Listen -> Speak
*Mar 8 20:34:20.221: SB11: V111 Hello out 172.16.11.111 Speak pri
100 ip 172.16.11.115
*Mar 8 20:34:23.101: SB11: V111 Hello out 172.16.11.111 Speak pri
100 ip 172.16.11.115
*Mar 8 20:34:25.961: SB11: V111 Hello out 172.16.11.111 Speak pri
100 ip 172.16.11.115
*Mar 8 20:34:28.905: SB11: V111 Hello out 172.16.11.111 Speak pri
```

```

100 ip 172.16.11.115
*Mar  8 20:34:30.221: SB11: V111 Speak: d/Standby timer expired
(unknown)
*Mar  8 20:34:30.221: SB11: V111 Standby router is local
*Mar  8 20:34:30.221: SB11: V111 Speak -> Standby
*Mar  8 20:34:30.221: SB11: V111 Hello  out 172.16.11.111 Standby pri
100 ip 172.16.11.115
*Mar  8 20:34:30.221: SB11: V111 Standby: c/Active timer expired
(unknown)
*Mar  8 20:34:30.221: SB11: V111 Active router is local
*Mar  8 20:34:30.221: SB11: V111 Standby router is unknown, was local
*Mar  8 20:34:30.221: SB11: V111 Standby -> Active
*Mar  8 20:34:30.221: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state
Standby -> Active
*Mar  8 20:34:30.221: SB11: V111 Hello  out 172.16.11.111 Active  pri
100 ip 172.16.11.115

```

To disable the debugging feature, enter either the **no debug standby** command or the **no debug all** command.

Example: HSRP Debugging on First and Only Router on Subnet

In this example, because DSW111 (172.16.11.111) is the only router on the subnet and because it is not configured for preempt, this router will go through all the HSRP states before becoming the active router. Notice at time stamp Mar 8 20:34:10.221 that the interface comes up and DSW111 enters the listen state. The router stays in the listen state for the holdtime of 10 seconds. DSW111 then goes into the speak state at time stamp Mar 8 20:34:20.221 for 10 seconds. When the router is speaking, it sends its state out every 3 seconds, according to its hello interval. After 10 seconds in speak state, the router has determined that there is no standby router at time stamp Mar 8 20:34:30.221 and enters the standby state. The router has also determined that there is not an active router; therefore, the router immediately enters the active state at time stamp Mar 8 20:34:30.221. From then on, the active router will send its active state hello message every 3 seconds. Because there are no other routers on this broadcast domain, no hellos are being received.

```

DSW111(config)#interface vlan 11
DSW111(config-if)#no shut

*Mar  8 20:34:08.925: %SYS-5-CONFIG_I: Configured from console by
console
*Mar  8 20:34:10.213: %LINK-3-UPDOWN: Interface Vlan11, changed state
to up
*Mar  8 20:34:10.221: SB:  V111 Interface up
*Mar  8 20:34:10.221: SB11: V111 Init: a/HSRP enabled
*Mar  8 20:34:10.221: SB11: V111 Init -> Listen
*Mar  8 20:34:11.213: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Vlan11, changed state to up
*Mar  8 20:34:20.221: SB11: V111 Listen: c/Active timer expired
(unknown)
*Mar  8 20:34:20.221: SB11: V111 Listen -> Speak
*Mar  8 20:34:20.221: SB11: V111 Hello  out 172.16.11.111 Speak  pri
100 ip 172.16.11.115
*Mar  8 20:34:23.101: SB11: V111 Hello  out 172.16.11.111 Speak  pri
100 ip 172.16.11.115
*Mar  8 20:34:25.961: SB11: V111 Hello  out 172.16.11.111 Speak  pri
100 ip 172.16.11.115
*Mar  8 20:34:28.905: SB11: V111 Hello  out 172.16.11.111 Speak  pri
100 ip 172.16.11.115
*Mar  8 20:34:30.221: SB11: V111 Speak: d/Standby timer expired

```

```
(unknown)
*Mar 8 20:34:30.221: SB11: V111 Standby router is local
*Mar 8 20:34:30.221: SB11: V111 Speak -> Standby
*Mar 8 20:34:30.221: SB11: V111 Hello out 172.16.11.111 Standby pri
100 ip 172.16.11.115
*Mar 8 20:34:30.221: SB11: V111 Standby: c/Active timer expired
(unknown)
*Mar 8 20:34:30.221: SB11: V111 Active router is local
*Mar 8 20:34:30.221: SB11: V111 Standby router is unknown, was local
*Mar 8 20:34:30.221: SB11: V111 Standby -> Active
*Mar 8 20:34:30.221: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state
Standby -> Active
*Mar 8 20:34:30.221: SB11: V111 Hello out 172.16.11.111 Active pri
100 ip 172.16.11.115
*Mar 8 20:34:33.085: SB11: V111 Hello out 172.16.11.111 Active pri
100 ip 172.16.11.115
*Mar 8 20:34:36.025: SB11: V111 Hello out 172.16.11.111 Active pri
100 ip 172.16.11.115
*Mar 8 20:34:38.925: SB11: V111 Hello out 172.16.11.111 Active pri
100 ip 172.16.11.115
```

Example: HSRP on NonPreempt Configured Router Coming Up

Router DSW111 (172.16.11.111) is configured with a priority of 100. This priority is higher than the priority of the current active router, DSW112 (172.16.11.112), which has a priority of 50. Note that router DSW111 is *not* configured with preempt. Only when it is configured with preempt will a router with a higher priority immediately become the active router. After router DSW111 goes through the HSRP initialization states, it will come up as the standby router.

```
DSW111(config)#interface vlan 11
DSW111(config-if)#no shut

*Mar 1 00:12:16.871: SB11: Vl11 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:16.871: SB11: Vl11 Active router is 172.16.11.112
*Mar 1 00:12:16.891: %SYS-5-CONFIG_I: Configured from console by
console
*Mar 1 00:12:18.619: %LINK-3-UPDOWN: Interface Vlan11, changed state
to up
*Mar 1 00:12:18.623: SB: Vl11 Interface up
*Mar 1 00:12:18.623: SB11: Vl11 Init: a/HSRP enabled
*Mar 1 00:12:18.623: SB11: Vl11 Init -> Listen
*Mar 1 00:12:19.619: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Vlan11, changed state to up
*Mar 1 00:12:19.819: SB11: Vl11 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:19.819: SB11: Vl11 Listen: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
*Mar 1 00:12:22.815: SB11: Vl11 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:22.815: SB11: Vl11 Listen: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
*Mar 1 00:12:25.683: SB11: Vl11 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:25.683: SB11: Vl11 Listen: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
*Mar 1 00:12:28.623: SB11: Vl11 Listen: d/Standby timer expired
(unknown)
*Mar 1 00:12:28.623: SB11: Vl11 Listen -> Speak
*Mar 1 00:12:28.623: SB11: Vl11 Hello out 172.16.11.111 Speak pri
100 ip 172.16.11.115
*Mar 1 00:12:28.659: SB11: Vl11 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:28.659: SB11: Vl11 Speak: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
*Mar 1 00:12:31.539: SB11: Vl11 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:31.539: SB11: Vl11 Speak: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
*Mar 1 00:12:31.575: SB11: Vl11 Hello out 172.16.11.111 Speak pri
100 ip 172.16.11.115
*Mar 1 00:12:34.491: SB11: Vl11 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:34.491: SB11: Vl11 Speak: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
*Mar 1 00:12:34.547: SB11: Vl11 Hello out 172.16.11.111 Speak pri
100 ip 172.16.11.115
*Mar 1 00:12:37.363: SB11: Vl11 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:37.363: SB11: Vl11 Speak: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
```

```
*Mar 1 00:12:37.495: SB11: V111 Hello out 172.16.11.111 Speak pri
100 ip 172.16.11.115
*Mar 1 00:12:38.623: SB11: V111 Speak: d/Standby timer expired
(unknown)
*Mar 1 00:12:38.623: SB11: V111 Standby router is local
*Mar 1 00:12:38.623: SB11: V111 Speak -> Standby
*Mar 1 00:12:38.623: SB11: V111 Hello out 172.16.11.111 Standby pri
100 ip 172.16.11.115
*Mar 1 00:12:40.279: SB11: V111 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:40.279: SB11: V111 Standby: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
*Mar 1 00:12:41.551: SB11: V111 Hello out 172.16.11.111 Standby pri
100 ip 172.16.11.115
*Mar 1 00:12:43.191: SB11: V111 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:43.191: SB11: V111 Standby: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
*Mar 1 00:12:44.539: SB11: V111 Hello out 172.16.11.111 Standby pri
100 ip 172.16.11.115
*Mar 1 00:12:46.167: SB11: V111 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:46.167: SB11: V111 Standby: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
*Mar 1 00:12:47.415: SB11: V111 Hello out 172.16.11.111 Standby pri
100 ip 172.16.11.115
*Mar 1 00:12:49.119: SB11: V111 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:12:49.119: SB11: V111 Standby: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
*Mar 1 00:12:50.267: SB11: V111 Hello out 172.16.11.111 Standby pri
100 ip 172.16.11.115
```

Example: HSRP on Preempt-Configured Router Coming Up

DSW111 (172.16.11.11) is configured with a priority of 100. This priority is higher than the priority of the active router, DSW112 (172.16.11.112). DSW111 is also configured with preempt. Only when a router is configured with preempt will that router with a higher priority transition into the active state. At time stamp Mar 1 00:16:43.099, the interface VLAN11 on DSW111 comes up and transitions into the listen state. At time stamp Mar 1 00:16:43.295, DSW111 receives a hello message from the active router (DSW112). DSW111 determines that the active router has a lower priority. At time stamp Mar 1 00:16:43.295, DSW111 immediately sends out a coup message, indicating that DSW111 is transitioning into the active router. DSW112 enters the speak state and eventually becomes the standby router.

```
DSW111(config)#interface vlan 11
DSW111(config-if)#no shut
```

```
*Mar 1 00:16:41.295: %SYS-5-CONFIG_I: Configured from console by
console
*Mar 1 00:16:43.095: %LINK-3-UPDOWN: Interface Vlan11, changed state
to up
*Mar 1 00:16:43.099: SB: Vl11 Interface up
*Mar 1 00:16:43.099: SB11: Vl11 Init: a/HSRP enabled
*Mar 1 00:16:43.099: SB11: Vl11 Init -> Listen
*Mar 1 00:16:43.295: SB11: Vl11 Hello in 172.16.11.112 Active pri
50 ip 172.16.11.115
*Mar 1 00:16:43.295: SB11: Vl11 Active router is 172.16.11.112
*Mar 1 00:16:43.295: SB11: Vl11 Listen: h/Hello rcvd from lower pri
Active router (50/172.16.11.112)
*Mar 1 00:16:43.295: SB11: Vl11 Active router is local, was
172.16.11.112
*Mar 1 00:16:43.295: SB11: Vl11 Coup out 172.16.11.111 Listen pri
100 ip 172.16.11.115
Mar 1 00:16:43.295
*Mar 1 00:16:43.299: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state
Listen -> Active
*Mar 1 00:16:43.299: SB11: Vl11 Hello out 172.16.11.111 Active pri
100 ip 172.16.11.115
*Mar 1 00:16:43.303: SB11: Vl11 Hello in 172.16.11.112 Speak pri
50 ip 172.16.11.115
*Mar 1 00:16:44.095: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Vlan11, changed state to up
*Mar 1 00:16:46.187: SB11: Vl11 Hello in 172.16.11.112 Speak pri
50 ip 172.16.11.115
*Mar 1 00:16:46.207: SB11: Vl11 Hello out 172.16.11.111 Active pri
100 ip 172.16.11.115
*Mar 1 00:16:49.095: SB11: Vl11 Hello in 172.16.11.112 Speak pri
50 ip 172.16.11.115
*Mar 1 00:16:49.195: SB11: Vl11 Hello out 172.16.11.111 Active pri
100 ip 172.16.11.115
*Mar 1 00:16:52.079: SB11: Vl11 Hello in 172.16.11.112 Speak pri
50 ip 172.16.11.115
*Mar 1 00:16:52.147: SB11: Vl11 Hello out 172.16.11.111 Active pri
100 ip 172.16.11.115
*Mar 1 00:16:53.303: SB11: Vl11 Hello in 172.16.11.112 Standby pri
50 ip 172.16.11.115
*Mar 1 00:16:53.303: SB11: Vl11 Standby router is 172.16.11.112
*Mar 1 00:16:55.083: SB11: Vl11 Hello out 172.16.11.111 Active pri
100 ip 172.16.11.115
*Mar 1 00:16:56.231: SB11: Vl11 Hello in 172.16.11.112 Standby pri
50 ip 172.16.11.115
```

```
*Mar  1 00:16:58.023: SB11: V111 Hello  out 172.16.11.111 Active  pri
100 ip 172.16.11.115
*Mar  1 00:16:59.223: SB11: V111 Hello  in  172.16.11.112 Standby pri
50 ip 172.16.11.115
*Mar  1 00:17:00.983: SB11: V111 Hello  out 172.16.11.111 Active  pri
100 ip 172.16.11.115
*Mar  1 00:17:02.211: SB11: V111 Hello  in  172.16.11.112 Standby pri
50 ip 172.16.11.115
*Mar  1 00:17:03.847: SB11: V111 Hello  out 172.16.11.111 Active  pri
100 ip 172.16.11.115
```

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **A single router can be a member of multiple HSRP groups.**
- **Preempt, timers, and interface tracking are options that can be configured to optimize HSRP.**
- **HSRP preempt can be tuned by adjusting timers that can thereby reduce failover time.**
- **Specific debug commands are used to view HSRP state changes.**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-15

References

For additional information, refer to this resource:

- Cisco Systems, Inc., *Hot Standby Router Protocol Features and Functionality*, http://cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a0080094a91.shtml#hsrpdebug

Configuring Layer 3 Redundancy with VRRP and GLBP

Overview

As the name would imply, Virtual Route Redundancy Protocol (VRRP) provides router interface failover in a manner similar to Hot Standby Router Protocol (HSRP) but with added features and IEEE compatibility. The process by which VRRP operates is defined in this lesson. The Gateway Load Balancing Protocol (GLBP) and its operations will be defined and differentiated from both HSRP and VRRP. Specific commands are used to implement and to verify VRRP and GLBP.

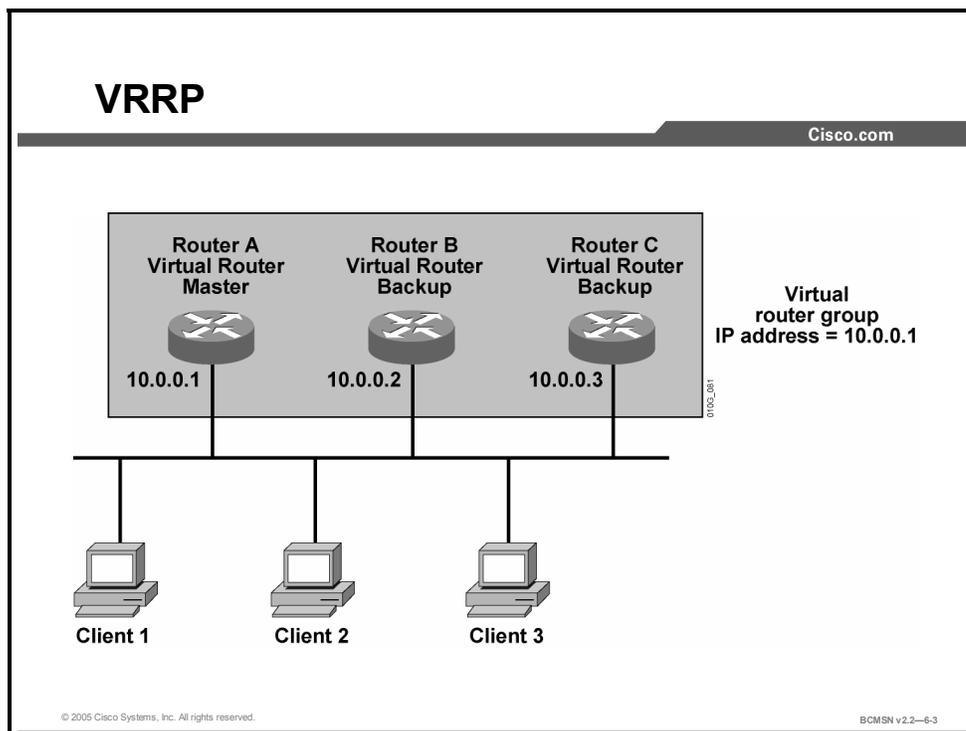
Objectives

Upon completing this lesson, you will be able to configure Layer 3 redundancy with VRRP or GLBP so that load balancing is implemented in addition to router failover at the distribution layer. This ability includes being able to meet these objectives:

- Describe the VRRP
- Identify how VRRP supports transitions from a master to a backup router
- Describe the GLBP
- Identify how GLBP balances traffic on a per-host basis using a round-robin scheme
- Identify the commands used to configure VRRP and GLBP
- Enable VRRP and GLBP to improve availability

Virtual Router Redundancy Protocol

This topic will describe the Virtual Router Redundancy Protocol (VRRP).



Like HSRP, VRRP allows a group of routers to form a single virtual router. The LAN workstations are then configured with the address of the virtual router as their default gateway. VRRP differs from HSRP in the following ways:

- VRRP is an IEEE standard for router redundancy; HSRP is a Cisco proprietary.
- The virtual router, representing a group of routers, is known as a VRRP group.
- The active router is referred to as the master virtual router.
- The master virtual router may have the same IP address as the virtual router group.
- Multiple routers can function as backup routers.

In the example, routers A, B, and C are members of a VRRP group. The IP address of the virtual router is the same as that of the LAN interface of router A (10.0.0.1). Router A is responsible for forwarding packets sent to this IP address.

The clients have a gateway address of 10.0.0.1. Routers B and C are backup routers. If the master router fails, the backup router with the highest priority becomes the master router. When router A recovers, it resumes the role of master router.

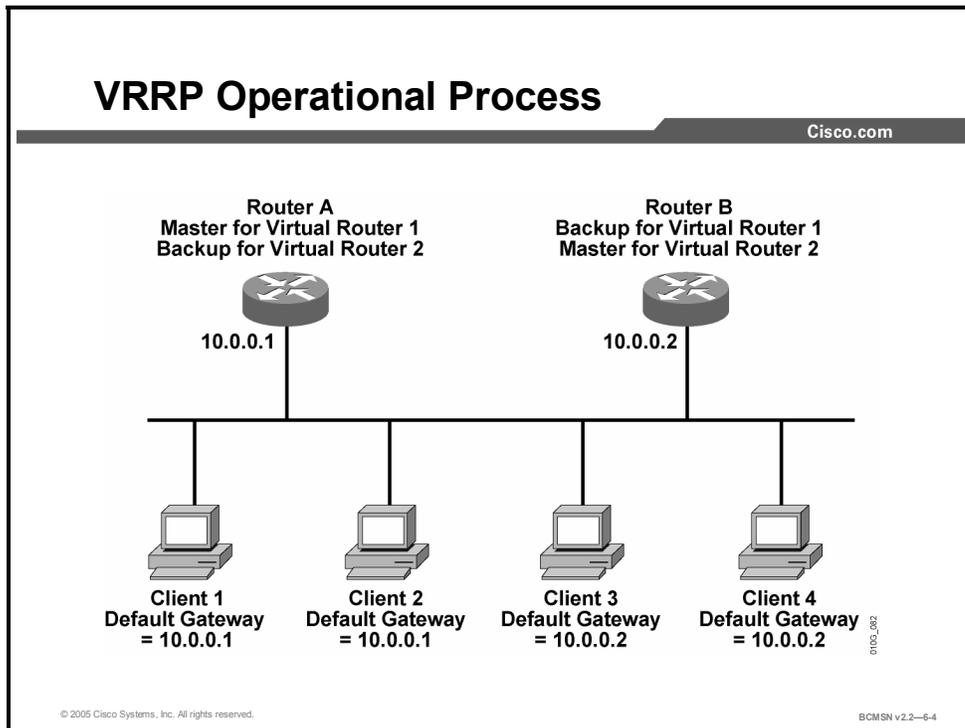
VRRP offers these redundancy features:

- VRRP provides redundancy for the real IP address of a router or for a virtual IP address shared among the VRRP group members.
- If a real IP address is used, the router with that address becomes the master. If a virtual IP address is used, the master is the router with the highest priority.

- A VRRP group has one master router and one or more backup routers. The master router uses VRRP messages to inform group members of the IP addresses of the backup routers.

Identifying the VRRP Operations Process

This topic describes VRRP operations.



This figure shows a LAN topology in which VRRP is configured so that routers A and B share the load of being the default gateway for clients 1 through 4. Routers A and B act as backup virtual routers to one another should either one fail.

In this example, two virtual router groups are configured. For virtual router 1, router A is the owner of IP address 10.0.0.1, and therefore the master virtual router for clients configured with that default gateway address. Router B is the backup virtual router to router A.

For virtual router 2, router B is the owner of IP address 10.0.0.2 and is the master virtual router for clients configured with the default gateway IP address of 10.0.0.2. Router A is the backup virtual router to router B.

Given that the IP address of the VRRP group is that of a physical interface on one of the group members, the router owning that address will be the master in the group. Its priority is set to 255. Backup router priority values can range from 1 to 254; the default value is 100.

The master sends the advertisement on multicast 224.0.0.18 on a default interval of 1 second. A VRRP flow message is similar in concept to an HSRP coup message. A master with a priority of 0 triggers a transition to a backup router.

The dynamic failover, when the active (master) becomes unavailable, uses two timers within VRRP: the advertisement interval and the master down interval. The advertisement interval is the time interval between advertisements (in seconds). The default interval is 1 second. The master down interval is the time interval for backup to declare the master down (in seconds).

Gateway Load Balancing Protocol

This topic describes the GLBP protocol.

Gateway Load Balancing Protocol

Cisco.com

- **Single gateway IP address on all routers**
- **Traffic to single gateway, distributed across routers**
- **Automatic rerouting in the event of any failure**
- **Fully utilizes resources on all routers in group**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—6-5

While HSRP and VRRP provide gateway resiliency, the standby members of the redundancy group their upstream bandwidth are not used while the device is in standby mode. Only the active router for HSRP and VRRP groups forwards traffic for the virtual MAC. Resources associated with the standby router are not fully utilized. Some load balancing can be accomplished with these protocols through the creation of multiple groups and through the assignment of multiple default gateways, but this configuration creates an administrative burden.

Cisco designed GLBP to allow automatic selection and simultaneous use of multiple available gateways as well as automatic failover between those gateways. Multiple routers share the load of frames that, from a client perspective, are sent to a single default gateway address. With GLBP, resources can be fully utilized without the administrative burden of configuring multiple groups and managing multiple default gateway configurations as is required with HSRP and VRRP.

Identifying the GLBP Operations Process

This topic describes how GLBP operates in a campus network.

GLBP Operations

Cisco.com

- **Election of Active Virtual Gateway**
- **Virtual MACs assigned to group members**
- **Client ARP's default gateway**
- **Active gateway replies with varied MACs**
- **Group members service failed virtual MACs**

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

GLBP allows automatic selection and simultaneous use of all available gateways in the group. The members of a GLBP group elect one gateway to be the Active Virtual Gateway (AVG) for that group. Other members of the group provide backup for the AVG should it become unavailable. The AVG assigns a virtual MAC address to each member of the GLBP group. All routers become Active Virtual Forwarders (AVFs) for frames addressed to that virtual MAC address. As clients send Address Resolution Protocol (ARP) requests for the address of the default gateway, the AVF sends these virtual MAC addresses in the ARP replies. A GLBP group can have up to four group members.

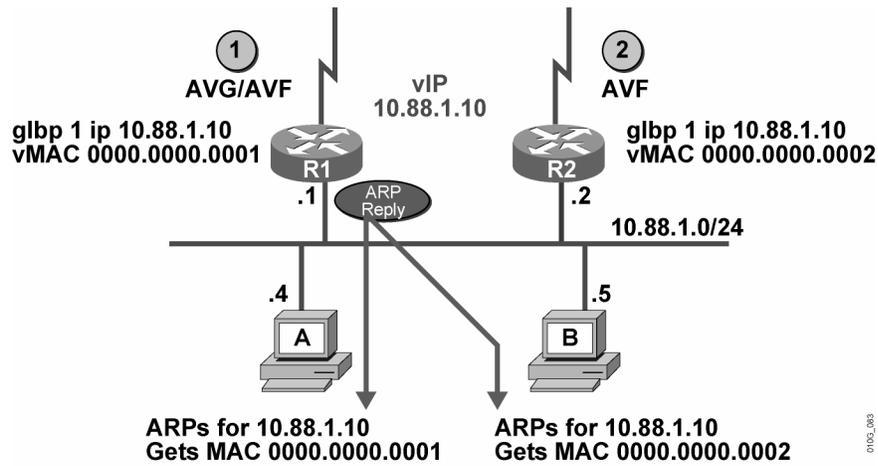
GLBP supports these operational modes for load balancing traffic across multiple default routers servicing the same default gateway IP address:

- **Weighted load-balancing algorithm:** The amount of load directed to a router is dependent upon the weighting value advertised by that router.
- **Host-dependent load-balancing algorithm:** A host is guaranteed to use the same virtual MAC address as long as that virtual MAC address is participating in the GLBP group.
- **Round-robin load-balancing algorithm:** As clients send ARP requests to resolve the MAC address of the default gateway, the reply to each client contains the MAC address of the next possible router in round-robin fashion. All routers' MAC addresses takes turns being included in address resolution replies for the default gateway IP address.

GLBP automatically manages the virtual MAC address assignment, determines who handles the forwarding, and ensures that each station has a forwarding path in the event of failures to gateways or tracked interfaces. If failures occur, the load-balancing ratio is adjusted among the remaining active virtual forwarders so that resources are used in the most efficient way.

GLBP Operation

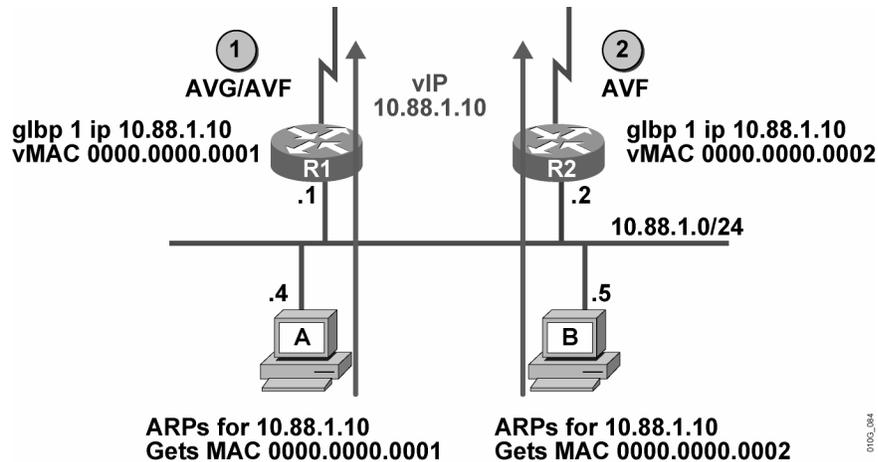
Cisco.com



As shown in the figure, by default, GLBP will attempt to balance traffic on a per-host basis using the round-robin algorithm. When a client sends an ARP message for the gateway IP address, the AVG will return the MAC address of one of the Active Virtual Forwarders. When a second device sends an ARP message, the AVG returns the next virtual MAC address from the list.

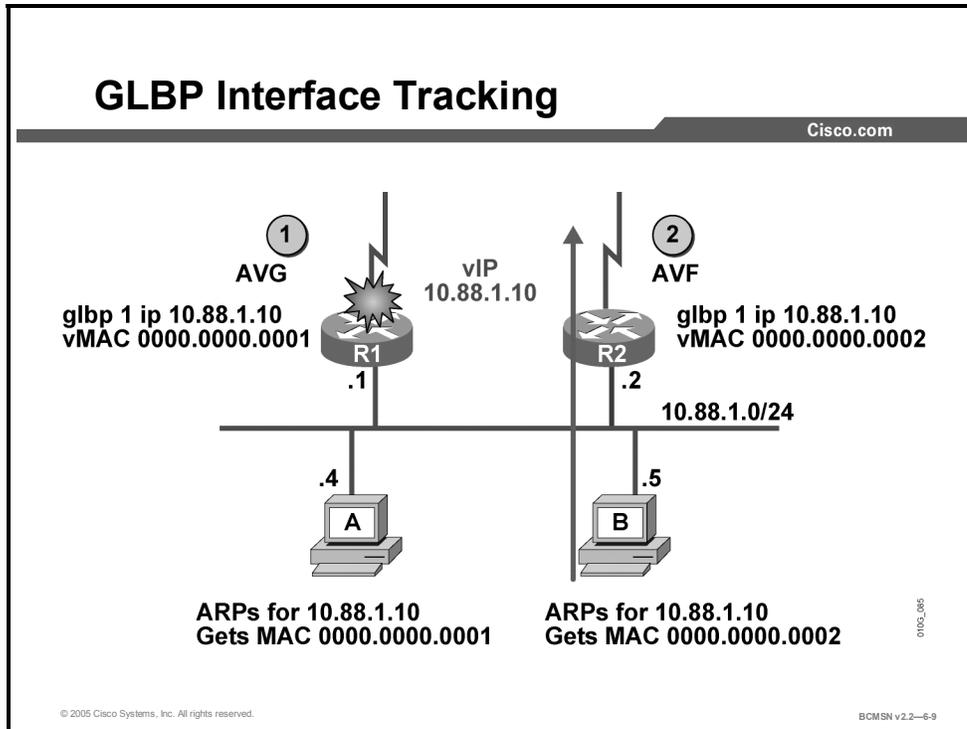
GLBP Operation (Cont.)

Cisco.com

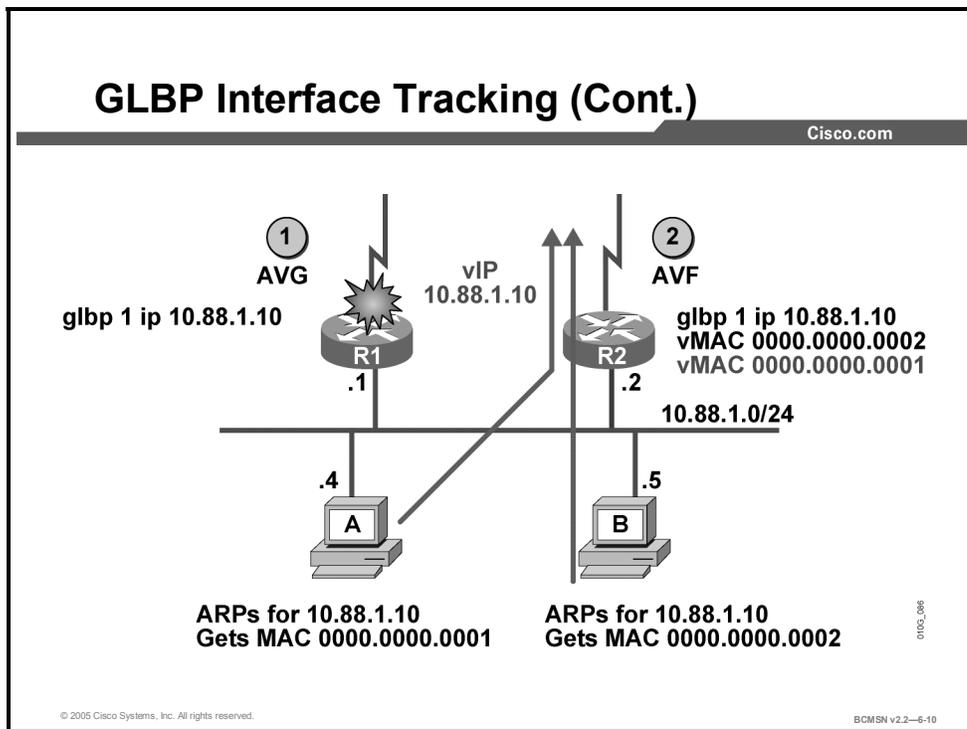


Having each resolved a different MAC address for the default gateway, clients A and B will send their routed traffic to separate routers, although they both have the same default gateway

address configured. Each GLBP router is an Active Virtual Forwarder for the virtual MAC address to which it has been assigned.



Like HSRP, GLBP can be configured to track interfaces. In the figure, the WAN link from router R1 is lost. GLBP detects the failure.



Because interface tracking was configured on R1, the job of forwarding packets for virtual MAC address 0000.0000.0001 will be taken over by the secondary virtual forwarder for the

MAC, router R2. Therefore, the client sees no disruption of service nor does it need to resolve a new MAC address for the default gateway.

VRRP and GLBP Configuration Commands

This topic lists the commands used to configure VRRP and GLBP operations.

VRRP and GLBP Configuration Commands

Cisco.com

Configure VRRP

- **vrrp 10 ip 10.1.1.1**
- **vrrp 10 priority 110**
- **vrrp 10 timers advertise 4**

Configure GLBP

- **glbp 10 ip 10.1.1.1**
- **glbp 10 priority 110**
- **glbp 10 timers msec 200 msec 700**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-6-11

VRRP Commands

Command	Description
Switch(config-if)# vrrp group ip virtual-gateway-addr	Makes the interface a member of the virtual group identified with the IP virtual address.
Switch(config-if)# vrrp group priority priority_value	Sets the priority of this router. Highest value will win election as active router. Default is 100. If routers have the same VRRP priority, the gateway with the highest real IP address is elected to become the master virtual router.
Switch(config-if)# vrrp group timers advertise timer-value	Master router configures this parameter to advertise value to the other group members. Others configure <i>timers learn</i> to accept.
Switch(config-if)# vrrp group-number timers learn	Configures non-master members to learn timer values from master.

GLBP Commands

Command	Description
Switch(config-if) # glbp group ip virtual-gateway-addr	Makes the interface a member of the virtual group identified with the IP virtual address.
Switch(config-if) # glbp group priority priority_value	Sets the priority of this router. Highest value will win election as active router. Default is 100. If routers have the same GLBP priority, the gateway with the highest real IP address will become the active virtual gateway.
Switch(config-if) # glbp group timers hello-value holdtime-value	Adjusts the hello and hold timers in seconds. Place the argument <i>msec</i> before the values to enter subsecond values.

How to Enable VRRP and GLBP

This topic describes the commands used to configure the VRRP and GLBP operations.

Configuring VRRP on an Interface

Cisco.com

```
interface vlan10
ip address 10.1.10.5
vrrp 10 ip 10.1.10.1
vrrp 10 priority 150
vrrp 10 timers advertise 4
```

10.1.10.5 Virtual Router
10.1.10.1

```
Switch# show running
```

Enable VRRP on an interface and display the configuration.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-12

VRRP and GLBP are supported on select Catalyst platforms and, when supported, can be configured using these commands:

VRRP Implementation

Step	Description
1.	To enable VRRP on an interface: <code>Switch(config-if)#vrrp <i>group-number</i> ip <i>virtual-gateway-address</i></code>
2.	To set a VRRP priority for this router for this VRRP group: <code>Switch(config-if)#vrrp <i>group-number</i> priority <i>priority-value</i></code>
3.	To change timer and indicate if it should advertise (master) or learn (backup): <code>Switch(config-if)#vrrp <i>group-number</i> timers advertise <i>timer-value</i></code> <code>Switch(config-if)#vrrp <i>group-number</i> timers learn</code>

Example: VRRP Implementation

```
Switch(config)# interface vlan10
Switch(config-if)# ip address 10.1.10.5 255.255.255.0
Switch(config-if)# vrrp 10 ip 10.1.10.1
Switch(config-if)# vrrp 10 priority 150
Switch(config-if)# vrrp 10 timer advertise 4
```

GLBP Implementation

This subtopic describes the process used to implement GLBP.

Configuring GLBP on an Interface

Cisco.com

```
interface vlan7
ip address 10.1.7.5
glbp 7 ip 10.1.7.1
glpb 7 priority 150
glbp 7 timers msec 200 msec 700
```

10.1.7.5

Virtual Router
10.1.7.1

```
Switch# show running-config
```

Enable GLBP on an interface and display the configuration.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-6-13

Below are the steps required to configure GLBP.

Step	Description
1.	Enable GLBP on an interface. Switch(config-if)# glbp group-number ip virtual-gateway-address
2.	Set a GLBP priority for this router for this GLBP group. Switch(config-if)# glbp group-number priority priority-value
3.	Change timer values for hello interval and holdtime. Switch(config-if)# glbp group-number timers hello holdtime

Example: GLBP Implementation

```
Switch(config)# interface vlan7
Switch(config-if)# ip address 10.1.7.5 255.255.255.0
Switch(config-if)# glbp 7 ip 10.1.7.1
Switch(config-if)# glbp 7 priority 150
Switch(config-if)# glbp 7 timers msec 250 msec 750
```

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **VRRP provides router redundancy in a manner similar to HSRP.**
- **VRRP supports a master and a backup router.**
- **GLBP provides router redundancy and load balancing.**
- **GLBP balances traffic on a per-host basis using a round-robin scheme.**
- **Specific commands are used to configure and verify VRRP and GLBP.**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-14

References

For additional information, refer to this resource:

- Cisco Systems, Inc., *Cisco IOS IP Command Reference, Volume 1 of 4: Addressing and Services*:
http://www.cisco.com/application/pdf/en/us/guest/products/ps5207/c2001/ccmigration_09186a0080238b7d.pdf

Implementing Hardware and Software Redundancy on Modular Switches

Overview

Redundancy in a switched network can be provided through device, topology, or hardware redundancy in combination with software components. There are benefits and drawbacks to implementing fault tolerance through device replication, as there are when redundancy is implemented by iterating the network topology.

Objectives

Upon completing this lesson, you will be able to configure redundancy features so that a second supervisor engine or power supply is enabled to provide failover capabilities with a Catalyst switch. This ability includes being able to meet these objectives:

- Define RPR+
- Configure and verify redundant supervisor engines
- Identify failover facilities that are specific to Cisco Catalyst 6500 Series switches
- Describe stateless switchover
- Describe SRM
- Enumerate the failure process with SRM and SSO
- Configure and verify SRM with SSO
- Describe Cisco Nonstop Forwarding
- List NSF aware protocols
- Enumerate the failover process of NSF with SSO
- Configure NSF for EIGRP and OSPF
- Configure and verify the operation of redundant power supplies

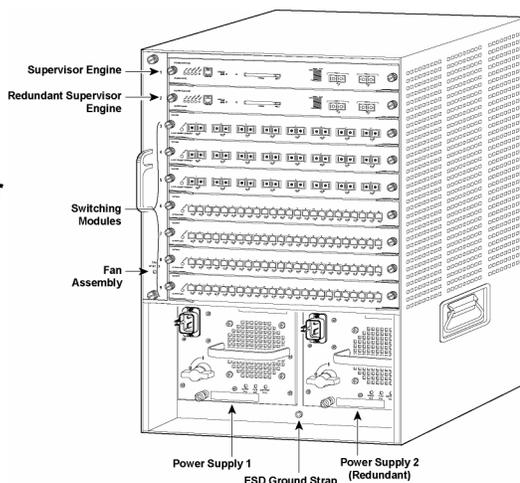
What Is RPR+?

This topic will define Route Processor Redundancy and explain why it is used.

RPR and Supervisor Redundancy

Cisco.com

- **Single supervisor engine, single point of failure**
- **Redundant supervisor engine, automatic failover**
- **RPR and RPR + address software failover in supervisor module**
- **RPR takes about 2-4 minutes to recover.**
- **RPR+ takes about 30+ seconds to recover.**



The diagram shows a Catalyst switch chassis with the following components labeled:

- Supervisor Engine
- Redundant Supervisor Engine
- Switching Modules
- Fan Assembly
- Power Supply 1
- ESD Ground Strap Connector
- Power Supply 2 (Redundant)

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

A Catalyst switch can allow a standby supervisor engine to take over if the primary supervisor engine fails. This allows the switch to resume operation quickly and efficiently in the event of a supervisor engine failure. This capability is called supervisor engine redundancy. In software, this capability is enabled by a feature called Route Processor Redundancy (RPR).

When RPR+ mode is used, the redundant supervisor engine is fully initialized and configured, and the Multilayer Switch Feature Card (MSFC) and the Policy Feature Cards (PFCs) are fully operational. This facilitates a faster failover time than RPR, in which the inactive supervisor engine is only partially booted.

The active supervisor engine checks the IOS version of the redundant supervisor engine when it boots. If the image on the redundant supervisor engine does not match the image on the active supervisor engine, RPR redundancy mode is used rather than RPR+.

There are several differences between the two RPR modes.

- RPR leaves the standby MSFC and PFC nonoperational until a failover occurs.
- RPR + places the standby MSFC and PFC in an operational mode upon boot, thereby providing faster failover.
- RPR+ maintains synchronization of the running-configuration file between the two supervisor engines.
- Both RPR and RPR+ maintain synchronization of the startup configuration file between the two supervisor engines.

Redundant Supervisor Engine Configuration Commands

This topic describes the commands used to implement RPR and RPR+.

RPR Configuration Commands

Cisco.com

Configure

- **redundancy**
- **mode rpr-plus**

Verify

- **show redundancy states**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—6-4

The following commands are used to enable and verify RPR or RPR+.

Command	Description
Switch(config)# redundancy	Configures RPR on the supervisor module and changes to redundancy configuration mode
Switch(config-red)# mode rpr-mode-name	Sets RPR mode rpr or rpr-plus
Switch# show redundancy states	Verifies the mode, status, and operation of RPR

How to Implement Redundant Supervisor Engines

This topic describes the sequence to be used to implement and verify RPR.

Configuring and Verifying Redundant Supervisor Engines with RPR or RPR+

Cisco.com

```
Router(config)#redundancy
```

- **Enables redundancy and enters redundancy configuration mode**

```
Router(config-red)#mode rpr
```

- **Specifies RPR as the supervisor engine redundancy mode**

```
Router(config-red)#mode rpr-plus
```

- **Specifies RPR+ as the supervisor engine redundancy mode**

```
Router#show redundancy states
```

- **Displays information about supervisor engine redundancy**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-6-5

These are the commands for implementing RPR+:

```
Router(config)#redundancy
```

```
Router(config-red)#mode rpr-plus
```

```
Switch#show redundancy states  
my state = 13 -ACTIVE  
peer state = 1 -DISABLED  
Mode = Simplex  
Unit = Primary  
Unit ID = 1
```

```
Redundancy Mode (Operational) = Route Processor Redundancy Plus  
Redundancy Mode (Configured) = Route Processor Redundancy Plus  
Split Mode = Disabled  
Manual Swact = Disabled Reason: Simplex mode  
Communications = Down Reason: Simplex mode
```

```
client count = 11  
client_notification_TMR = 30000 milliseconds  
keep_alive TMR = 4000 milliseconds  
keep_alive count = 0  
keep_alive threshold = 7  
RF debug mask = 0x0
```

Cisco Catalyst 6500 Switch

This topic describes Layer 3 capability on a Catalyst 6500 using an MSFC.

About the Catalyst 6500

Cisco.com

- **MSFC adds Layer 3 capabilities to Catalyst.**
- **Configured through Cisco IOS interface**
- **MSFC is daughter card on Supervisor Engine.**



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

The Catalyst 6500 platform provides Layer 3 functionality through an MSFC residing on the supervisor engine module. As of this writing, the current iteration is the MSFC3, which is an integral part of the Supervisor Engine 720. The MSFC3 adds high-performance multilayer switching and routing intelligence to the Catalyst. Equipped with a high-performance processor, the MSFC3 runs Layer 2 protocols on one CPU and Layer 3 protocols on the second CPU. These protocols include VLAN Trunking Protocol, routing protocols, multimedia services, security services—nearly any protocol capable of running on the high-end Cisco routing platforms.

The MSFC builds the Cisco Express Forwarding information base table in software and downloads this table to the hardware or ASIC on the PFC and any installed Distributed Forwarding Card (DFC).

An MSFC3 with PFC3 on a Supervisor 720 adds stateful switchover (SSO) and Nonstop Forwarding (NSF) to the arsenal of Catalyst fault tolerance features.

What Is Stateful Switchover?

This topic describes stateful switchover (SSO), supported on Catalyst 4500/6500 switches.

What Is Stateful Switchover?

Cisco.com

- **Supported on Catalyst 4500/6500 switches**
- **Redundant supervisor engine starts fully initialized**
- **Maintains Layer 2 protocol states**
- **Supported in 12.2(20)EWA and later releases**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-6-7

When a redundant supervisor engine runs in SSO mode, the redundant supervisor engine starts up in a fully initialized state and synchronizes with the persistent configuration and the running configuration of the active supervisor engine. It subsequently maintains the state of the Layer 2 protocols; all changes in hardware and software states for features that support stateful switchover are kept in sync. Consequently, it offers zero interruption to Layer 2 sessions in a redundant supervisor engine configuration. SSO is supported in 12.2(20)EWA and later releases.

Because the redundant supervisor engine recognizes the hardware link status of every link, ports that were active before the switchover will remain active, including the uplink ports. However, because uplink ports are physically on the supervisor engine, they will be disconnected only if the supervisor engine is removed.

If the active supervisor engine fails, the redundant supervisor engine becomes active. This newly active supervisor engine uses existing Layer 2 switching information to continue forwarding traffic. Layer 3 forwarding will be delayed until the routing tables have been repopulated in the newly active supervisor engine.

What Is Single Router Mode?

This topic explains the purpose and features of single router mode (SRM).

Single Router Mode

Cisco.com

- **SRM provides router redundancy using a single Catalyst chassis.**
- **One MSFC is designated.**
- **Standby MSFC is hot with current configuration.**
- **Standby MSFC has VLAN interfaces in a line-down state.**
- **Routing table is populated after standby becomes active.**
- **SRM requirements:**
 - **Both MSFCs run the same IOS image.**
 - **High availability configured on the supervisor engine.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—6-8

In SRM redundancy, only the designated router (MSFC) is visible to the network at any given time. Dual router mode (DRM) had both MSFCs active and used Hot Standby Router Protocol (HSRP) to maintain an active and secondary relationship. DRM had the problems of extra complexity and routing protocol peering, which are overcome by using SRM. The nondesignated router is booted up completely and participates in configuration synchronization, which is automatically enabled when entering SRM. The configuration of the nondesignated router is exactly the same as that of the designated router, but its interfaces are kept in a "line-down" state and are not visible to the network. Processes such as routing protocols are created on the nondesignated router and the designated router, but all nondesignated router interfaces are in a line-down state; they do not send or receive updates from the network.

When the designated router fails, the nondesignated router changes its state to become the designated router, and its interface state changes to "link up." It builds its routing table while the existing supervisor engine switch processor entries are used to forward Layer 3 traffic. After the newly designated router builds its routing table, the entries in the switch processor are updated.

Because only one MSFC is visible to the network at a given time, multiple Border Gateway Protocol (BGP) peering sessions do not have to exist between two MSFCs. If the designated MSFC fails, the nondesignated MSFC reestablishes BGP peering. Therefore, it always appears as a single BGP peer to the network and simplifies the network design, but it gives the same level of redundancy in case of an MSFC failure.

Failure with SRM and SSO

This topic describes failover scenarios for SRM and its associated MSFCs.

SRM Failure

Cisco.com

- **Failure event occurs on primary supervisor engine.**
- **Secondary supervisor takes over.**
- **Line cards and Layer 2 protocols keep forwarding.**
- **Active supervisor is reloaded.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—6-9

When the switch is powered on, SRM with SSO runs between the two supervisor engines. The supervisor engine that boots first becomes the active supervisor. The MSFC3 and PFC3 become fully operational.

If the active supervisor engine 720 or MSFC3 fails, the redundant supervisor engine 720 and MSFC3 become active. The newly active supervisor engine 720 uses the existing PFC3 Layer 3 switching information to forward traffic while the newly active MSFC3 builds its routing table.

The routing protocols have to establish connectivity with their neighbor or peers and the routing information base is built. During this time, packet forwarding cannot take place.

How to Configure and Verify SRM with SSO

This topic describes the commands associated with SRM configuration.

SRM Configuration and Verification Commands

Cisco.com

```
Router(config)#redundancy
```

- Enters redundancy configuration mode

```
Router(config-red)#mode sso
```

- Configures SRM with SSO

```
Router#show running-config
```

- Displays the current configuration

```
Router#show redundancy states
```

- Displays redundancy configuration information

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-6-10

Use these commands to configure the MSFC for SRM with SSO operation.

Command	Description
Router(config)# redundancy	Enters redundancy configuration mode.
Router(config-red)# mode sso	Configures SRM with SSO. When this command is entered, the redundant supervisor engine 720 is reloaded and begins to work in SRM with SSO mode.

Note The **sso** keyword is supported in 12.2(17b)SXA and later releases.

Configure the MSFCs for SRM with SSO

Enable SRM on the designated router first and then enable SRM on the nondesignated router as follows:

```
Router> enable
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# redundancy
Router(config-red)# mode sso
Router(config-red)# end
```

Verify SRM Configuration and Operation

This subtopic describes how to verify SRM operation.

Verifying SRM

Cisco.com

```
Router#show running-config
```

- **Displays the current configuration**

```
Router#show redundancy states
```

- **Displays redundancy configuration information**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-11

This command will verify that SRM with SSO has been configured correctly on each MSFC3:

```
Switch# show running-config
```

This command will verify that SRM with SSO is currently operational on the MSFC3:

```
Switch# show redundancy states
```

```
Router# show redundancy states
```

```
my state = 13 -ACTIVE
peer state = 8 -STANDBY HOT
Mode = Duplex
Unit = Primary
Unit ID = 5
Redundancy Mode (Operational) = Stateful Switchover
Redundancy Mode (Configured) = Stateful Switchover
Split Mode = Disabled
Manual Swact = Enabled
Communications = Up
client count = 29
client_notification_TMR = 30000 milliseconds
keep_alive TMR = 9000 milliseconds
keep_alive count = 1
keep_alive threshold = 18
```

RF debug mask = 0x0

Router#

Note Enter the **copy running-config** command on the designated router to ensure that SRM is enabled on the nondesignated router startup configuration.

Caution Before going from DRM to SRM redundancy, Cisco recommends that you use the **copy running-config** command on the MSFCs to save the non-SRM configuration to boot Flash memory. When you go to SRM redundancy, the alternative configuration (the configuration following the **alt** keyword) is lost. Therefore, before enabling SRM redundancy, save the DRM configuration to boot Flash memory by entering the following command on both MSFCs: **copy running-config bootflash:nosrm_dual_router_config**.

What Is Nonstop Forwarding?

This topic describes NSF, supported only on Catalyst 6500.

What Is Nonstop Forwarding?

Cisco.com

- **Runs with SSO**
- **Provides Layer 3 redundancy**
- **Continues forwarding packets in event of supervisor engine failover**
- **Supported by BGP, OSPF, IS-IS, and EIGRP**
- **Maintains packet forwarding through CEF table while routing tables rebuild**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—6-12

Cisco NSF always runs with SSO and provides redundancy for Layer 3 traffic. NSF works with SSO to minimize the amount of time that a network is unavailable to its users following a switchover. The main purpose of NSF is to continue forwarding IP packets after a supervisor engine switchover and the subsequent establishment of the routing protocols' peering relationships.

Cisco NSF is supported by the BGP, Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS), and Enhanced Interior Gateway Routing (EIGRP) protocols for routing and by Cisco Express Forwarding (CEF) for forwarding. The routing protocols have been enhanced with NSF capability and awareness, which means that remote routers running these protocols and configured for NSF can detect a switchover, recover route information from the peer devices, and take the necessary actions to continue forwarding network traffic. The IS-IS protocol can be configured to use state information that has been synchronized between the active and the redundant supervisor engines to recover route information following a switchover, rather than using information received from peer devices.

A networking device is NSF-aware if it is running NSF-compatible software. A device is NSF-capable if it has been configured to support NSF; it will rebuild routing information from NSF-aware or NSF-capable neighbors.

Each protocol depends on CEF to continue forwarding packets during switchover while the routing protocols rebuild the routing information base tables. After the routing protocols have converged, CEF updates the Forwarding Information Base (FIB) table and removes stale route entries. CEF then updates the line cards with the new FIB information.

Cisco NSF provides these benefits:

- Network availability is improved.
- Network stability may be improved with the reduction in the number of route flaps.
- Because the interfaces remain up throughout a switchover, neighboring routers do not detect a link flap (in other words, the link does not go down and come back up).
- User sessions established before the switchover are maintained.

Caution A number of caveats and restrictions apply to this process. Consult the Cisco documentation for your particular hardware and software combination to find more specific information.

Identifying NSF-Aware Protocols

This topic describes NSF-aware protocols supported by Catalyst 4500 and 6500.

NSF-Aware Protocols

Cisco.com

- **BGP**
- **OSPF**
- **IS-IS**
- **EIGRP**
- **Peer or neighbor routers need to run NSF-aware protocols to maintain stateful information.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—6-13

Routing protocols run only on the MSFC of the active supervisor engine, and they receive routing updates from their neighbor routers. Routing protocols do not run on the MSFC of the redundant supervisor engine. After a switchover, the routing protocols request that the NSF-aware neighbor devices send state information to help rebuild the routing tables. Alternatively, the IS-IS protocol can be configured to synchronize state information from the active to the redundant supervisor engine to help rebuild the routing table on the NSF-capable device in environments where neighbor devices are not NSF-aware. Cisco NSF supports the BGP, OSPF, IS-IS, and EIGRP protocols

Note For NSF operation, the routing protocols depend on CEF to continue forwarding packets while the routing protocols rebuild the routing information.

EIGRP Operation

When an EIGRP NSF-capable router initially comes back up from an NSF restart, it has no neighbors, and its topology table is empty. The router is notified by the redundant (now active) supervisor engine when it needs to bring up the interfaces, reacquire neighbors, and rebuild the topology and routing tables. The restarting router and its peers must accomplish these tasks without interrupting the data traffic directed toward the restarting router. EIGRP peer routers maintain the routes learned from the restarting router and continue forwarding traffic through the NSF restart process.

BGP Operation

When an NSF-capable router begins a BGP session with a BGP peer, it sends an OPEN message to the peer. Included in the message is a statement that the NSF-capable device has "graceful" restart capability. Graceful restart is the mechanism by which BGP routing peers avoid a routing flap after a switchover. If the BGP peer has received this capability, it is aware that the device sending the message is NSF-capable. Both the NSF-capable router and its BGP peers need to exchange the graceful restart capability in their OPEN messages at the time of session establishment. If both peers do not exchange the graceful restart capability, the session will not be graceful-restart-capable.

OSPF Operation

When an OSPF NSF-capable router performs a supervisor engine switchover, it must complete the following tasks to resynchronize its link state database with its OSPF neighbors:

- Relearn the available OSPF neighbors on the network without causing a reset of the neighbor relationship
- Reacquire the contents of the link state database for the network

As quickly as possible after a supervisor engine switchover, the NSF-capable router sends an OSPF NSF signal to neighboring NSF-aware devices. Neighbor networking devices recognize this signal as an indicator that the neighbor relationship with this router should not be reset. As the NSF-capable router receives signals from other routers on the network, it can begin to rebuild its neighbor list.

IS-IS Operation

When an IS-IS NSF-capable router performs a supervisor engine switchover, it must complete the following tasks to resynchronize its link state database with its IS-IS neighbors:

- Relearn the available IS-IS neighbors on the network without causing a reset of the neighbor relationship
- Reacquire the contents of the link state database for the network

The IS-IS NSF feature offers two options when you configure NSF:

- Internet Engineering Task Force (IETF) IS-IS
- Cisco IS-IS

Failover with NFS and SSO

This topic describes NFS and SSO failover.

NSF Failover

Cisco.com

- **Catalyst 4500 and 6500 feature**
- **Hardware failure on active supervisor engine**
- **Clock synchronization failure**
- **Manual switchover to second supervisor engine**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—6-14

The Catalyst 4500 and 6500 switches support fault resistance by allowing a redundant supervisor engine to take over if the primary supervisor engine fails. Cisco NSF works with SSO to minimize the amount of time the routing protocols require to rebuild their tables following a switchover. Catalyst 4500 and 6500 Series switches also support route processor redundancy (RPR), route processor redundancy plus (RPR+), and single router mode with stateful switchover (SRM with SSO) for redundancy.

The following events cause a switchover:

- A hardware failure on the active supervisor engine
- Clock synchronization failure between supervisor engines
- A manual switchover

How to Configure NSF

This topic describes how to configure NSF using Cisco IOS software.

Configuring NSF

Cisco.com

- **Configure SSO**
- **Configure NSF support for each protocol in use:**
 - **EIGRP**
 - **BGP**
 - **OSPF**
 - **IS-IS**
 - **Multicast MLS**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—6-15

In addition to configuring SSO, you will need to consider the following procedures for protocols configured on the Multilayer Switching (MLS) when SSO with NSF is implemented:

- Configure and verify multicast MLS NSF with SSO
- Configure and verify BGP NSF
- Configure and verify OSPF NSF
- Configure and verify IS-IS NSF
- Configure and verify EIGRP NSF

Caution Full NSF support requires NSF configuration steps for all routing protocols running at the time of failover.

The table illustrates how to configure NSF for EIGRP..

Example: NSF Configuration for EIGRP

To configure NSF for EIGRP, use the following commands, beginning in privileged EXEC mode.

Configuring EIGRP for NSF

Step	Action	Notes
1.	Router(config)# router eigrp <i>as-number</i>	Enables an EIGRP routing process, which places the router in router configuration mode
2.	Router(config-router)# nsf	Enables NSF for EIGRP Command to use on the restarting router and all of its peers

Verifying EIGRP for NSF

To verify NSF for EIGRP, you must check that the NSF function is configured on the SSO-enabled networking device. To verify EIGRP NSF, follow these steps.

Step 1 Verify that "nsf" appears in the EIGRP configuration of the SSO-enabled device by entering the **show running-config** command:

```
Router# show running-config
.
router eigrp 100
  auto-summary
  nsf
.
```

Step 2 Use the **show ip protocols** command to verify that NSF is enabled on the device:

```
router#show ip protocols

*** IP Routing is NSF aware ***
Routing Protocol is "eigrp 100"
  Outgoing update filter list for all interfaces is not set
.
.
.
  EIGRP NSF-aware route hold timer is 240s
  EIGRP NSF enabled
    NSF signal timer is 20s
    NSF converge timer is 120s
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
  Routing Information Sources:
    Gateway         Distance      Last Update
  Distance: internal 90 external 170
```

References

For more information on configuring NSF with the 6500, see Cisco Systems, Inc., *Configuring NSF with SSO Supervisor Engine Redundancy*:

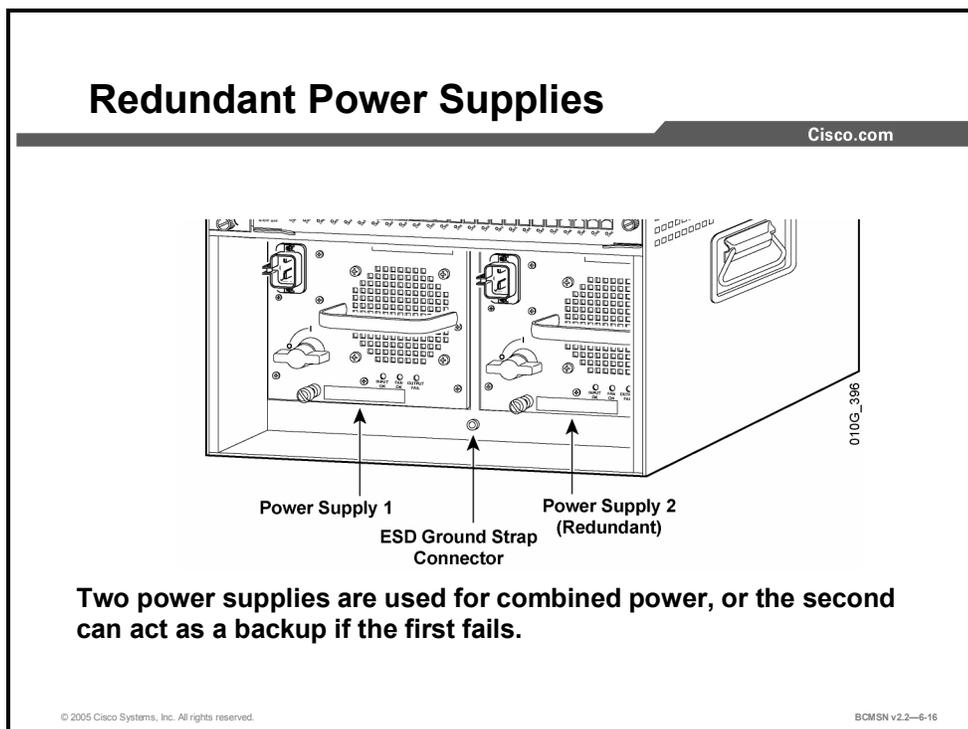
http://cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a008027e4cd.html#wp1119697

For more information on configuring NSF-aware functionality with the 4500, see Cisco Systems, Inc., *Configuring Supervisor Engine Redundancy Using RPR and SSO*:

http://cisco.com/en/US/products/hw/switches/ps4324/products_configuration_guide_chapter09186a00802c3052.html

Redundant Power Supply Configuration

This topic explains how and why to implement redundant power supplies.



The Catalyst 6000 Series of switches can be configured with dual power supplies. When dual power supplies are installed, they share the power requirements of the switch. In some cases, a specific combination of modules can exceed the maximum power output of a single power supply, and two power supplies must be installed to provide adequate power to the switch. Care should be taken if this option is chosen, as the redundancy provided by dual power supplies would be compromised.

This table lists commands used to implement redundant power supplies.

Command	Description
Switch(config)# power redundancy-mode mode	Configures redundancy on the power supply. Modes are as follows: combined (combines the resources of both power supplies) redundancy (allows one supply to back up the other)
Switch# show power	Displays information about the state of each power supply and if redundancy has been invoked

Note It is very important to maintain power budgets, especially when using inline power modules to support IP telephony.

How to Configure Redundant Power Supplies

This topic describes the commands to configure and verify redundant power supplies.

Configuring Redundant Power Supplies

Cisco.com

```
Switch(config)#power redundancy-mode combined | redundant
```

- Configures power as combined or redundant

```
Switch#show power
```

- Displays information about power supplies

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-17

The following command is used to implement power supply redundancy:

```
Switch(config)# power redundancy-mode mode
```

- **Combined mode** disables redundancy. The power available to the system is the combined power capability of both power supplies. The system powers up as many modules as the combined capacity allows. If one supply should fail and there is not enough power for all previously powered-up modules, the system powers down those modules for which there is not enough power.
- **Redundant mode** enables redundancy. In a redundant configuration, the total power drawn is at no time greater than the capability of one supply. If one supply malfunctions, the other supply can take over the entire system load. During normal operation with two power supplies, each provides approximately half of the required power to the system. Load sharing and redundancy are enabled automatically; no software configuration is required.

The **show power** command displays the current state of modules and the total power available:

```
Switch#show power
system power redundancy mode = redundant
system power total = 27.460A
system power used = -6.990A
system power available = 20.470A
FRU-type      #      current      admin state oper
power-supply  1      27.460A    on           on
module        1      -4.300A    on           on
module        2      -4.300A    off          off (admin request)
module        5      -2.690A    on           on
```

Turn Off or Cycle Power to Modules

Power to a given module can be removed using this command;

```
Switch(config) no power enable module slot
```

Note When the **no power enable module slot** command is used to power down a module, the module configuration is not saved.

To turn power back on for a module that was previously powered down, issue this command:

```
Switch(config) power enable module slot
```

The power to a module can be cycled (reset). The power cycle will turn the module power off for 5 seconds and then back on. To power cycle a module, issue this command:

```
Switch(config) power cycle module slot
```

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **RPR+ provides optimal supervisor engine failover.**
- **Specific commands are used to implement RPR+.**
- **The Catalyst 6500 optimizes hardware failover.**
- **Stateful switchover optimizes hardware failover.**
- **Single router mode allows only a single MSFC to be visible at one time.**
- **SRM failover is optimized when used with SSO.**
- **Specific commands are used to implement SRM with SSO.**
- **Nonstop forwarding ensures routing protocol operation during MLS hardware failover.**
- **Specific configuration parameters are required for NFS to operate with EIGRP and OSPF.**
- **Redundant power supplies can exist in a single catalyst chassis for the purpose of failover.**
- **Specific commands are used to configure power supply redundancy.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—6-18

Designing High Availability in a Multilayer Switch

Overview

Redundancy in a switched network can be provided through device, topology, or hardware replication in combination with software components. There are benefits and drawbacks to implementing fault tolerance through device replication as there are when redundancy is implemented by iterating the network topology.

Objectives

Upon completing this lesson, you will be able to describe features and best practices to increase network availability in a multilayer switch. This ability includes being able to meet these objectives:

- Define redundancy in a switched network
- Identify the benefits and drawbacks of providing high availability solely through fault-tolerant devices
- Identify the benefits and drawbacks of providing high availability through redundancy in the network topology
- Differentiate the connection options between switches when switch stacks are implemented in the access layer of the Campus Infrastructure module
- List best practice recommendations for supporting highly available and deterministic network paths at the access layer of the Campus Infrastructure module
- List best practice recommendations for supporting highly available and deterministic network paths at the distribution layer of the Campus Infrastructure module
- Define a strategy that provides Layer 2 and Layer 3 failover and recovery paths that are aligned in the access and distribution layers of the Campus Infrastructure module
- List best practice recommendations for supporting highly available and deterministic network paths at the core layer of the Campus Infrastructure module

What Is Redundancy in a Switched Network?

This topic introduces hardware redundancy.

High-Availability Network Components

Cisco.com

- **Reliable, fault tolerant network devices**
- **Device and link redundancy**
- **Resilient network technologies**
- **Optimized network design**
- **Best practices**

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

Providing hardware redundancy in a switched network can be accomplished by implementing redundant modules within devices or by deploying redundant devices.

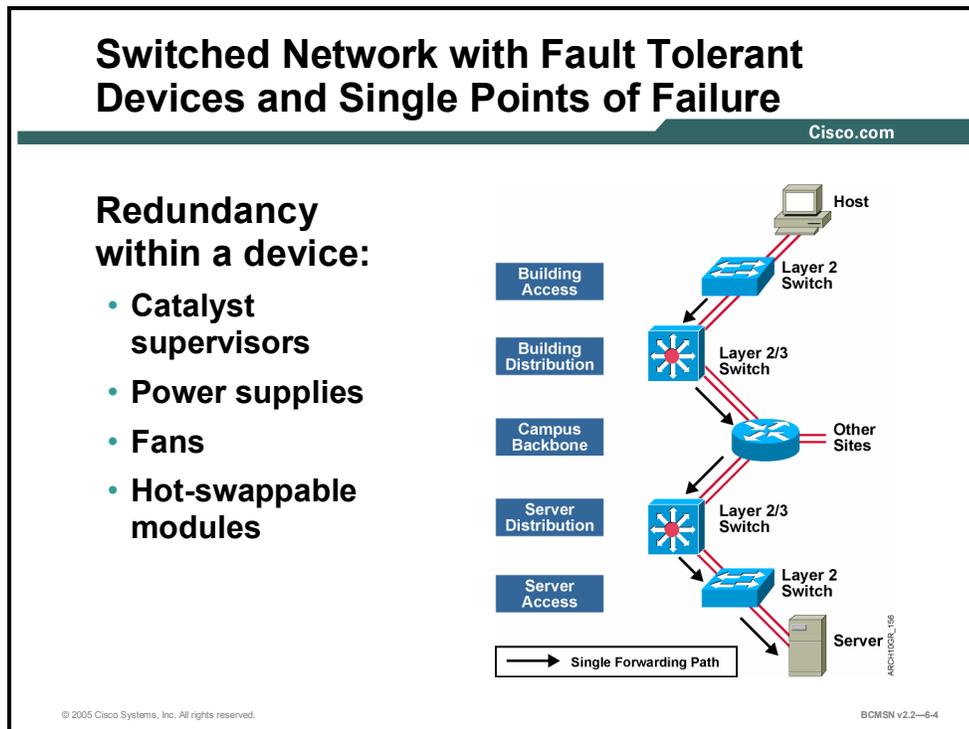
To achieve network availability as close to 100 percent of the time as possible, these network components are required:

- **Reliable, fault tolerant network devices:** Hardware and software reliability to automatically identify and overcome failures
- **Device and link redundancy:** May include entire devices, modules within devices, and links
- **Resilient network technologies:** Intelligence that ensures fast recovery around any device or link failure
- **Optimized network design:** Well-defined network topologies and configurations designed to ensure that there is no single point of failure
- **Best practices:** Documented procedures for deploying and maintaining a robust e-commerce network infrastructure

Network fault tolerance indicates the ability of a device or network to recover from the failure of a component or device. Achieving high availability relies on eliminating any single point of failure and on distributing intelligence throughout the architecture. You can increase availability by adding redundant components, including redundant network devices and connections to redundant Internet services. With the proper design, no single point of failure will have an impact on the availability of the overall system.

Benefits and Drawbacks of Device-Level Fault Tolerance

Here is an explanation of how to use multiple devices to create network fault tolerance.



One approach to building highly available networks is to replicate all devices to create a fault tolerant network. To achieve high end-to-end availability, each key network infrastructure device exists in duplicate. Fault tolerance through device replication offers several benefits.

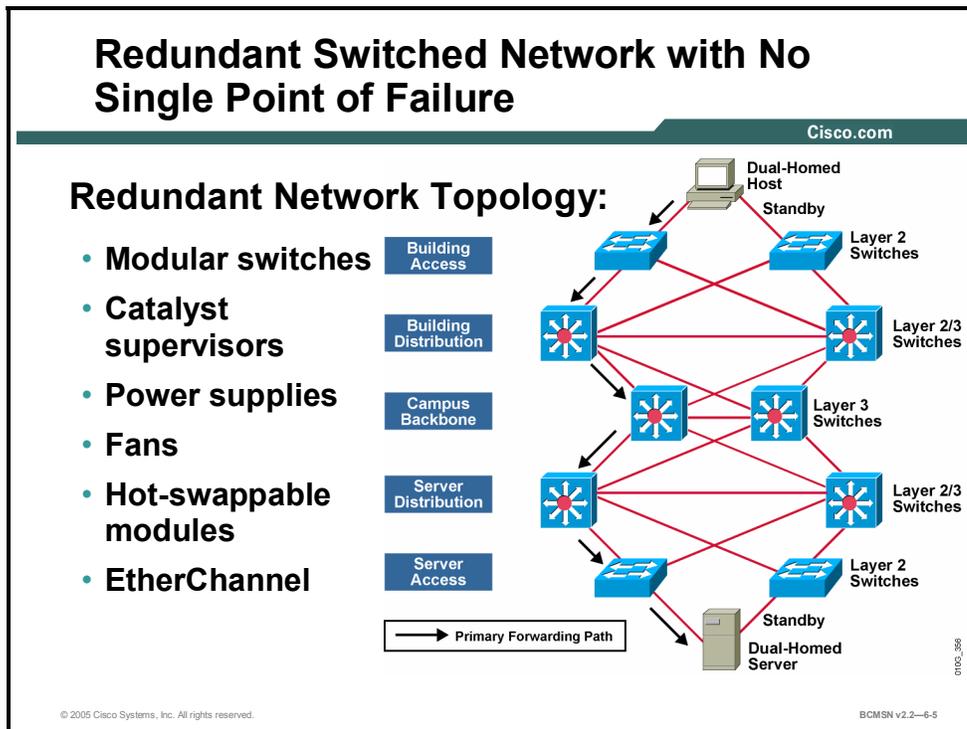
- Minimizes time periods during which the system is nonresponsive to requests (for example, while the system is being reconfigured because of a component failure or recovery)
- Eliminates all single points of failure that would cause the system to stop
- Provides disaster protection by allowing the major system components to be separated geographically

Trying to achieve high network availability solely through device-level fault tolerance has a number of drawbacks.

- Massive redundancy within each device adds significantly to its cost. Massive redundancy also reduces the physical capacity of each device by consuming slots that could otherwise house network interfaces or provide useful network services.
- Redundant subsystems within devices are often maintained in a hot standby mode. In hot standby mode, such redundant subsystems cannot contribute additional performance because they are only fully activated when the primary component fails.
- Focusing on device-level hardware reliability may result in a number of other failure mechanisms being overlooked. Network elements are not standalone devices; they are components of a network system whose internal operations and system-level interactions are governed by software and configuration parameters.

Benefits and Drawbacks of Redundant Network Topology

A redundant network topology has benefits.



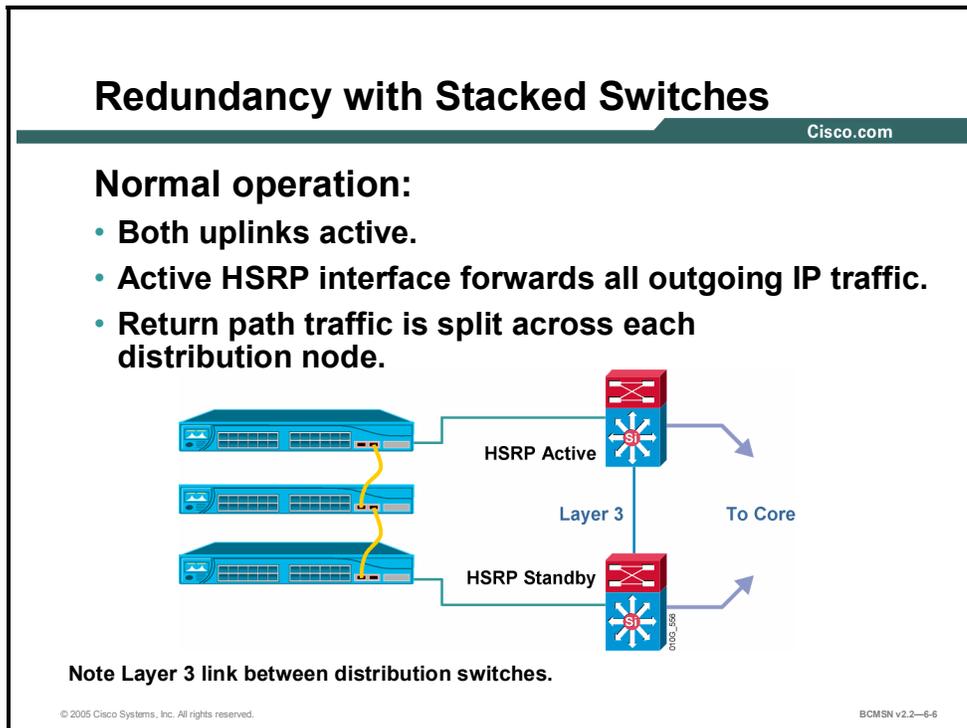
A complementary way to build highly available networks is to provide redundancy in the links between devices in the network topology. In the campus network design shown in the figure, there is a backup for every link and for every network device in the path between the client and server. Using network links to supplement devices' fault tolerance has these advantages:

- The network elements providing redundancy can be geographically disparate. This reduces the probability that problems with the physical environment will interrupt service.
- Software errors and changes can be dealt with separately in the primary and secondary forwarding paths without completely interrupting service.
- Device-level fault tolerance can be concentrated in the Building Core and Building Distribution Layers of the network, where a hardware failure would affect a larger number of users. By partially relaxing the requirement for device-level fault tolerance, the cost per network device is reduced. To some degree, this offsets the requirement for more devices.
- Redundant links provided for fault tolerance can be used to balance the traffic load between the respective layers of the network topology (that is, the Building Access Layer to the Building Distribution Layer, also distribution to core) during times of normal operation. Therefore, network-level redundancy can also provide increased aggregate performance and capacity.
- Redundant resources can be configured to failover from primary to secondary facilities automatically. Failover times can be as low as subsecond, depending on the type of failure (for example, Link failure, Node failure, and so on).

- Fast EtherChannel and Gigabit EtherChannel provide both fault tolerance and high-speed links between switches with minimal convergence times in the event of link loss.

Redundancy with Stacked Switches

This topic describes the use of stacked switches in the Campus Infrastructure module.



Stacking access switches has become commonplace. Behavior of a switch stack in the event of a failure depends on its application. Stacked switches are sometimes used to implement hardware redundancy and high port density at the access layer.

Rather than having redundant uplinks between each access and distribution device, the stack as a whole represents a single logical switch with redundant links between the stack and the distribution layer.

The figure shows a deployment of stacked switches in the Campus Infrastructure module with Hot Standby Router Protocol (HSRP) being used. When the link between distribution devices is a Layer 3 link, no Layer 2 can exist, so the uplinks to both distribution switches are active.

Layer 3 Failure with Stacked Switches

This subtopic describes Layer 3 failure when a link in the switch stack fails.

Layer 3 Failure with Stacked Switches

Cisco.com

Upon failure of a stack switch or cable:

- **HSRP becomes active on both distribution switches.**
- **Return path traffic is split across distribution switches and potentially isolated from destination hosts.**

Note Layer 3 link between distribution switches.

© 2005 Cisco Systems, Inc. All rights reserved. BCM5N v2.2-6-7

Consider a failure of either a middle switch or cable in the switch stack.

The stack maintains the Layer 2 connectivity between the distribution switches. When a link between switches in the stack fails, HSRP packets are no longer sent between the two distribution switches. This causes the standby HSRP router to transition to active and advertise itself as the default gateway. Traffic from SW1 uses distribution SWA as the active gateway, but now traffic from SW2 and SW3 uses distribution SWB as the default gateway.

As the distribution switches announce routes into the core, the VLAN interfaces on both distribution switches will advertise reachability to the IP subnet of the switch stack. These VLAN interfaces will present themselves as equal-cost paths to the subnet. When return path, potentially load-balanced traffic arrives at each distribution VLAN interface, some percentage will not be able to reach the originating end system because it is on the wrong side of the failure.

Loopback Cable to Maintain Layer 2 Path

This subtopic discusses adding a cable between each switch in the stack to provide an alternative path in the event of a failure.

Loopback Cable for Fault Tolerance

Cisco.com

Avoiding problem due to failure:

- Loopback cable across stack provides connectivity.
- STP is required for reconvergence.

Note Layer 3 link between distribution switches.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-6-8

If a loopback cable is installed between the end switches of the switch stack, the Layer 2 path in the segment has redundancy that can be maintained by Spanning Tree Protocol (STP). HSRP communication can now be maintained between the distribution switches if a Layer 2 link occurs. Ideally, the HSRP and STP failover times would be closely associated, so that there is little time when connectivity is compromised. RST (Rapid Spanning Tree) should be implemented if at all possible.

Note Stack redundancy software and hardware solutions such as Stackwise in the Cisco Catalyst 3750 Series can avoid some of the issues associated with stacked switches with no common backplane.

High Availability: Access Layer Best Practices

This topic describes best practices at the access layer.

High Availability: Access Layer Best Practices

Cisco.com

- Ensure that trunk parameters match at each end.
- Limit VLANs to single closet.
- VTP should be in transparent mode.
- Prune unused VLANs.
- Set trunks ON/ON.
- Rapid PVST+
- Establish STP failover plan.

Access Layer

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-6-9

When deploying the Campus Infrastructure module, adopting best practice recommendations at the access layer means providing a highly available and deterministic Layer 2 network. It is generally assumed that high availability in the access layer will be accomplished through the implementation of link redundancy between the access and distribution layers, with STP managing the use of those links. Redundant links to individual user devices are not typical.

These are best practices to follow when establishing highly available access layer devices.

- If two different versions of Cisco software exist at either end of a trunk link, ensure that trunk parameters are manually set to match one another.
- Limit VLANs to a single access switch or switch stack. Spanning VLANs across switches may be necessary in some instances but should be avoided if at all possible.
- VLAN Trunk Protocol (VTP) can be disabled or run in transparent mode only.
- Use the range command to prune unused VLANs from trunks. This will ensure that both Layer 2 and Layer 3 convergence occurs and that the traffic traverses the intended paths through the network.
- Set trunks permanently on to avoid autonegotiation and security issues.
- Rapid PVST+ is preferred to keep convergence times to 1 to 2 seconds. Leave STP active even if there are no redundant Layer 2 links in the network. This will guard against the attachment of rogue switches.
- Establish a STP Failover Plan to share the primary root evenly between the two distribution switches.

Note Consider using multilayer switches and routing at the access layer to avoid the use of spanning tree and minimize convergence time.

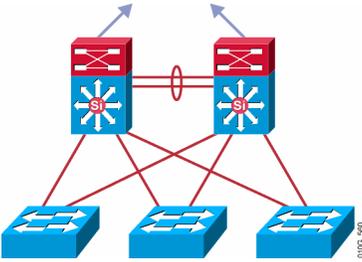
High Availability: Distribution Layer Best Practices

This topic describes high availability best practices at the distribution layer.

High Availability: Distribution Layer Best Practices

Cisco.com

- Use equal-cost redundant paths to core.
- Use a Layer 3 link between distribution switches where possible.
- Optimize EIGRP or OSPF HSRP/GLBP convergence with summarization and passive interfaces.
- Tune EtherChannel and CEF.



The diagram illustrates a network topology with two distribution switches (DS) at the top, each represented by a red and blue square with a white 'S' and a star. They are connected to three access switches (AS) at the bottom, represented by blue squares with a white 'A' and a star. A red line connects the two DS switches, representing a Layer 3 link. Red lines also connect each DS switch to all three AS switches, representing equal-cost redundant paths. Blue arrows point upwards from each DS switch, indicating traffic flow towards the core. The diagram is labeled '01002_500' in the bottom right corner.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-10

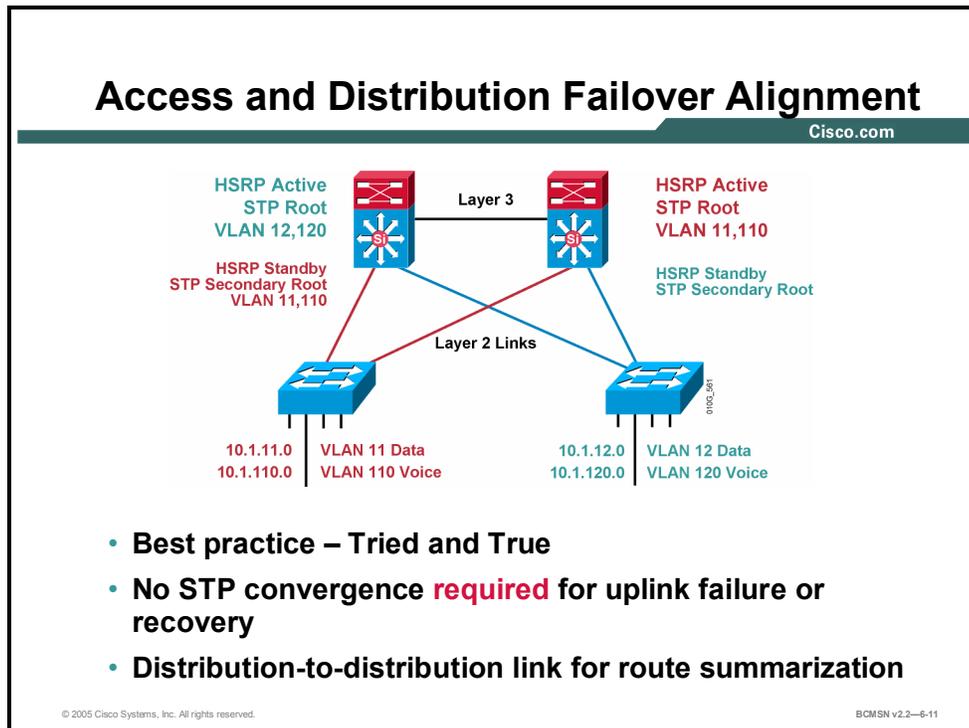
Adopting these best practice recommendations at the distribution layer will support the intention of providing a highly available and deterministic network.

- Connect distribution switches with a Layer 3 EtherChannel link.
- Use equal-cost redundant connections between the distribution and core for fastest convergence and to avoid black holes.
- Summarization is required to facilitate optimum Enhanced Interior Gateway Routing Protocol (EIGRP) or Open Shortest Path First (OSPF) convergence. If summarization is implemented at the distribution layer, the distribution nodes must be linked, or routing black holes occur.
- Use Gateway Load Balancing Protocol (GLBP) or HSRP millisecond timers. Convergence around a link or node failure in the Layer 2-Layer 3 distribution boundary model depends on default gateway redundancy and failover. Millisecond timers can reliably be implemented to achieve subsecond (800 ms) convergence based on HSRP or GLBP failover.
- Tune GLBP or HSRP preempt delay to avoid black holes. HSRP and GLBP standby peers so that traffic is not dropped while connectivity to the core is established. The delay should be adjusted to ensure that the node is ready to forward traffic before it preempts.
- The hierarchical campus model implements multiple Layer 3 equal-cost paths, and traffic should be load balanced across these paths from the access layer to the distribution and core layers. The Cisco Express Forwarding (CEF) hashing algorithm should be tuned at the core and distribution layers to vary decision input and avoid CEF polarization, which can

result in under-utilization of redundant paths. Use the default Layer information for the core nodes and use Layer 3 with Layer 4 information for the distribution nodes.

Layers 2 and 3 Redundancy Alignment

This topic discusses best practices for Layers 2 and 3 failover alignment.



When you are implementing strategies for failover at the access and distribution layers, it is important that the failover paths and timers are aligned between Layer 2 failover protocol (STP) and Layer 3 failover protocol (HSRP or GLBP). This is most significant if the link between the distribution switches is a Layer 2 link and is therefore hosting a redundant Layer 2 path for the VLANs in the access layer. Although the link is a Layer 3 link, alignment of the protocols is still a best practice in case a rogue switch is placed on the network.

In the figure, the left-hand distribution switch is configured as the HSRP active router for VLANs 12 and 120 and is also configured as the STP primary root for the same VLANs. The right-hand distribution switch serves as the HSRP standby and STP secondary root for those VLANs.

Likewise, the right-hand distribution switch is configured as the HSRP active router for VLANs 11 and 110 and is also configured as the STP primary root for the same VLANs. The left-hand distribution switch serves as the HSRP standby and STP secondary root for VLANs 11 and 110.

It is important that the timers of STP and HSRP agree, providing failover and recovery at the nearly the same time. This would require the implementation of RSTP on all access and distribution switches.

Autostate Layer 3 Convergence During Layer 2 Failure

This subtopic discusses alignment between Layer 3 routing protocol convergence and Layer 2 link failure.

Using Autostate to Ensure Layer 3 Convergence

Cisco.com

- VLAN interface on distribution switch should shut down when no Layer 2 path exists to local VLAN.
- Autostate requires that at least one port offers an active connection to the VLAN.
- Autostate senses the STP forwarding state of ports associated with VLAN.

3. No direct connection route to VLAN10

4. New routed path onto VLAN11

2. Interface VLAN11 shuts down

1. Physical failure on only Layer 2 link to VLAN11

Layer 3

Layer 2 Links

VLAN 11

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-6-12

The autostate feature notifies a switch or routing module VLAN interface (Layer 3 interface) to transition to “up” status when at least one Layer 2 port becomes active in that VLAN.

Autostate also senses the STP forwarding state of ports associated with the VLAN ID. This will prevent routing protocols and other features from using the VLAN interface as if it were fully operational.

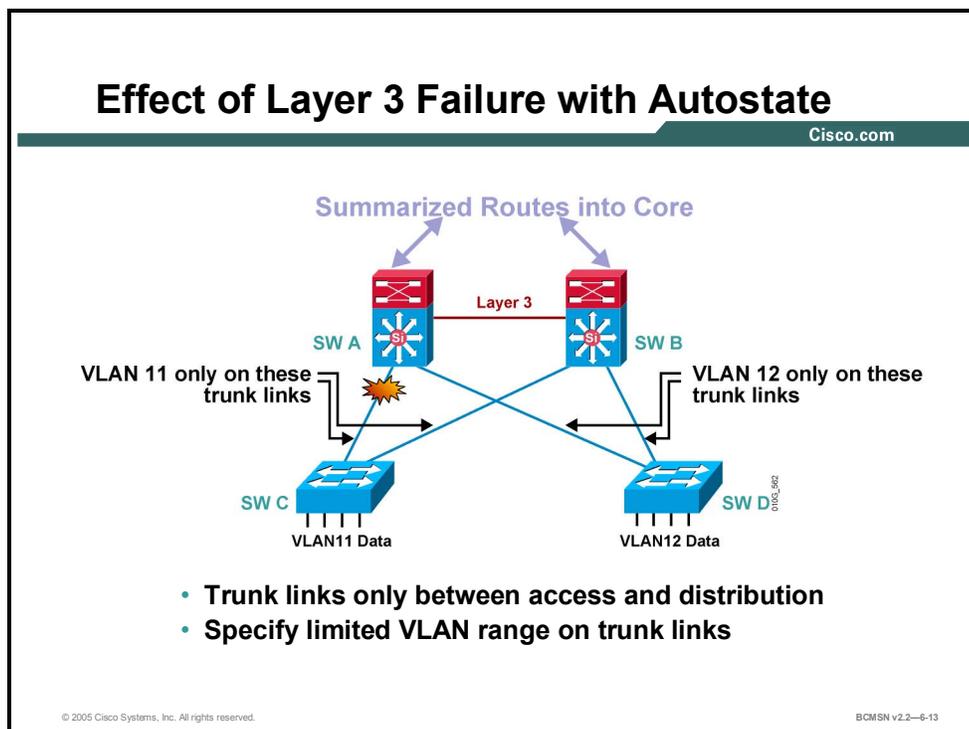
To operate correctly, there should be no local ports with the VLAN ID that are not offering a connection directly to the access switch that has that VLAN configured.

- Trunk links that have the VLAN ID are assumed to provide a path to the VLAN and will keep the interface up.
- Access ports with the VLAN ID will also keep the VLAN interface up.

An example of a problem: A trunk link to an access switch has only VLAN 12 and 14 associated with it, but the trunk is configured to carry all VLANs. This trunk would appear to the autostate process to provide a path every active VLAN, and hence local VLAN interfaces for 12 and 14 would never be shut down because this trunk appears to provide a path.

Affect of Layer 3 Failure with Autostate

This subtopic describes the affect of Layer 3 failure when autostate is in use.



Using the **trunk range** command will ensure an appropriate reaction by the VLAN interface to a loss of physical connectivity. Having discussed the process of autostate, we can now discuss the effects of a failure upon IP traffic. For the following discussion, we will assume that the distribution nodes are summarizing.

As seen in the figure, when the Layer 2 trunk between SW A and SW C fails, physical connectivity to VLAN11 is lost on SW A. This is because the trunks are properly configured so autostate will detect that there are no longer any ports active for VLAN11. The VLAN11 interface will shut down on SW A, and the directly connected route to VLAN11 will be removed from the routing table.

Benefits of this process:

- The distribution switch will replace its directly connected route to VLAN11, with the route to VLAN11 being advertised by SW B across the Layer 3 link.
- When return path traffic arrives on the distribution switch SW A, destined for VLAN 11, it will be routed toward the access layer through SW B.
- Because summarization is taking place, no external network routing update has been propagated into the core.

If the VLAN interface had not shut down, then the IP return path traffic would have been lost at SW A. This is sometimes referred to as being “black holed.”

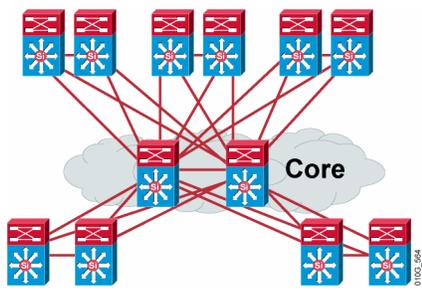
High Availability: Core Layer Best Practices

This topic discusses best practices for high availability at the core layer.

Core Layer Tuning

Cisco.com

- For optimum convergence build triangles rather than squares.
- Use point-to-point link technologies to speed up propagation of link events.
- Routed core preferred:
 - Faster convergence
 - Increased scalability
 - Efficient bandwidth utilization.



© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—6-14

These best practices are recommended for optimum core layer convergence.

- Build triangles, not squares, to take advantage of equal-cost redundant paths for the best deterministic convergence.
- Design the core layer as a high-speed Layer 3 switching environment using only hardware-accelerated services. Layer 3 core designs are superior to Layer 2 and other alternatives because they provide
 - Faster convergence around a link or node failure
 - Increased scalability because neighbor relationships and meshing are reduced
 - More efficient bandwidth utilization
- With high availability in the core, it is assumed that point-to-point links such as direct Ethernet links exist between the core and distribution and between core devices. Link-up or -down topology changes can be propagated almost immediately. With topologies that rely on indirect notification and timer-based detection such as SONET, convergence is nondeterministic and is measured in seconds

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Redundancy is required in a switched network.**
- **Redundancy is not effective when provided solely through fault tolerant devices.**
- **Redundancy is not effective when provided solely through redundancy in network topology.**
- **Switch stacks provide high port density and hardware redundancy at the access layer.**
- **Best practices should be followed when implementing redundancy in access, distribution, and core layers.**
- **Layer 2 and Layer 3 failover strategies should be aligned in the access and distribution layers.**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—6-15

Module Summary

This topic summarizes the key points discussed in this module.

Summary

Cisco.com

- **HSRP is enabled so that redundant routers can provide default gateway functionality.**
- **HSRP can be tuned to provide subsecond failover to a standby router.**
- **VRRP or GLBP can provide Layer 3 router failover as well as load balancing at the distribution layer.**
- **SSO and NSF provide the capability for a second supervisor engine or power supply to provide failover capabilities within a Catalyst switch.**
- **Specific technologies and best practices are required to maintain high availability in the Campus Infrastructure module.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—6-1

Device, link, or hardware component redundancy at strategic points in the network leads to high availability. HSRP provides router redundancy to network hosts and can be optimized in several ways. VRRP and GLBP were derived from HSRP, providing additional redundancy features. Cisco uses redundant supervisor engines as well as software features to optimize network resiliency upon component failure. SRM is a Cisco high availability solution facilitating failover to a second MSFC. Highly available networks are those configured to provide infrastructure paths and access to key devices at all times.

References

For additional information, refer to these resources:

- Cisco Systems, Inc., *Hot Standby Router Protocol Features and Functionality*:
http://cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a0080094a91.shtml#hsrpdebug
- Cisco Systems, Inc., *Cisco IOS IP Command Reference, Volume 1 of 4: Addressing and Services*:
http://www.cisco.com/application/pdf/en/us/guest/products/ps5207/c2001/ccmigration_09186a0080238b7d.pdf

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) During which three HSRP states do routers send hello messages? (Choose three.) (Source: Configuring Layer 3 Redundancy with HSRP)
- A) initial
 - B) listen
 - C) speak
 - D) active
 - E) standby
- Q2) Which of the following two are Layer 3 redundancy protocols? (Choose two.) (Source: Implementing Redundancy in the Routing Layer)
- A) HSRP
 - B) VRRP
 - C) GRRP
 - D) All of the above are correct.
- Q3) HSRP provides Layer 3 redundancy through the use of a _____? (Source: Configuring Layer 3 Redundancy with HSRP)
- A) virtual router
 - B) virtual switch
 - C) logical router
 - D) logical switch
- Q4) In VRRP, the active router is also known as the _____? (Source: Configuring Layer 3 Redundancy with VRRP and GLBP)
- A) master virtual router
 - B) master router
 - C) virtual router
 - D) master VRRP router
- Q5) In which state are the nondesignated router interfaces when SRM is configured? (Source: Implementing Hardware and Software Redundancy in Modular Switches)
- A) visible
 - B) link-up
 - C) invisible
 - D) line-down
- Q6) Which two of the following routing protocols are NSF aware? (Choose two.) (Source: Implementing Hardware and Software Redundancy in Modular Switches)
- A) RIPv2
 - B) EIGRP
 - C) IS-IS
 - D) IGRP

Module Self-Check Answer Key

Q1) D, E, F

Q2) A, B

Q3) A

Q4) A

Q5) D

Q6) B, C

Minimizing Service Loss and Data Theft in a Switched Network

Overview

Today's networks are vulnerable to more than just cable infrastructure failures and power outages. Because of the ubiquity of Internet connectivity, strategic partnerships, and the dependence on electronic information coming from external entities, organizations must create a well-defined security posture for their networks that spells out possible vulnerabilities, defines the threats, and describes the countermeasures that should be implemented to mitigate the associated risks. In this module, a number of network security vulnerabilities are discussed, and actual threat types are defined. Cisco Systems has implemented a number of device-level countermeasures to defend the individual devices as well as the entire network from security threats. This module describes some measures to secure Cisco's multilayer switches based on the Cisco Security Architecture for Enterprise (SAFE) blueprint.

Module Objectives

Upon completing this module, you will be able to secure switches in the Campus Infrastructure module against data theft and service loss in the event of a network compromise. This ability includes being able to meet these objectives:

- Identify key security issues in a switched network
- Identify configuration features that mitigate VLAN attacks
- Identify switch configuration features that can mitigate spoof attacks
- Implement AAA services to enforce secure authentication, authorization and accounting
- Identify switch security risks and list best practices for placing new switches into service

Understanding Switch Security Issues

Overview

Basic security measures should be taken to guard against a host of attacks that can be launched at a switch and its ports. Specific measures can be taken to guard against MAC flooding, which is a common Layer 2 malicious activity.

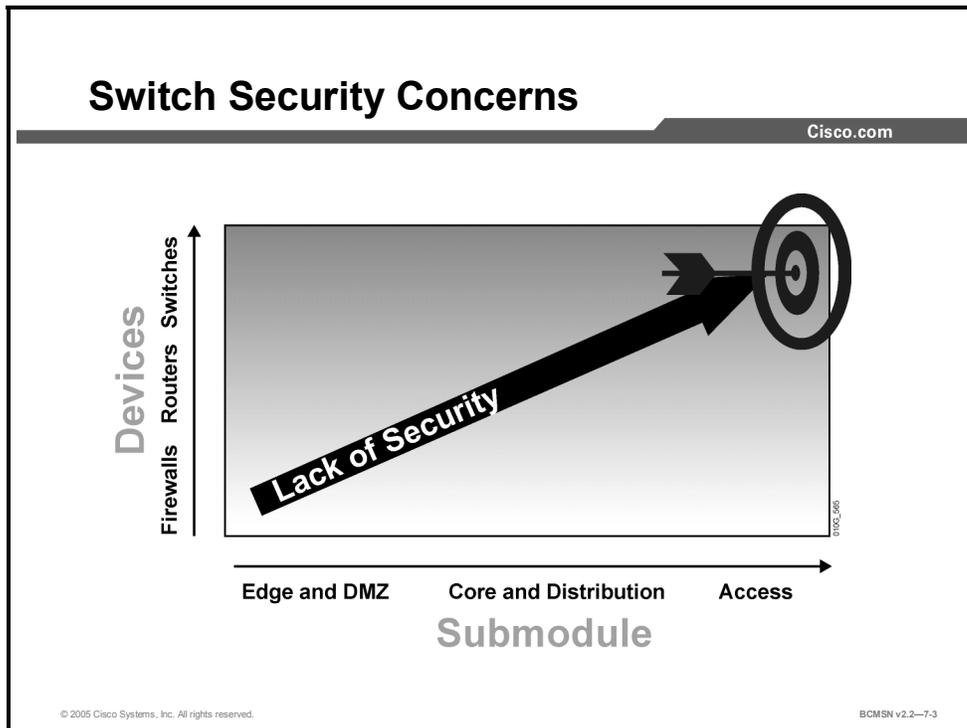
Objectives

Upon completing this lesson, you will be able to identify key security issues in a switched network. This ability includes being able to meet these objectives:

- Describe switch and Layer 2 security as a subset of an overall network security plan
- Categorize switch attack types and list mitigation options
- Describe a MAC flooding attack
- Describe how port security is used to block input from devices
- Configure port security on a switch
- Configure the “sticky” MAC option with port security

Switch Security Concerns

This topic describes security concerns that are specific to switches.



Much industry attention surrounds security attacks from outside the walls of an organization and at the upper Open Systems Interconnection (OSI) layers. Network security often focuses on edge-routing devices and the filtering of packets based upon Layers 3 and 4 headers, ports, stateful packet inspection, and so forth. This includes all issues surrounding Layer 3 and above, as traffic makes its way into the campus network from the Internet. Campus access devices and Layer 2 communication are left largely unconsidered in most security discussions.

The default state of networking equipment highlights this focus on external protection and internal open communication. Firewalls, placed at the organizational borders, arrive in a secure operational mode and allow no communication, until configured to do so. Routers and switches placed internal to an organization and designed to accommodate communication, delivering needful campus traffic, have a default operational mode that forwards all traffic unless configured otherwise. Their function as devices to facilitate communication often results in minimal security configuration and renders them targets for malicious attacks. If an attack is launched at Layer 2 on an internal campus device, the rest of the network can be quickly compromised, often without detection.

Many security features are available for switches and routers, but they must be enabled to be effective. As with Layer 3, where security had to be tightened on devices within the campus as malicious activity increased that compromised this layer, now security measures must be taken to guard against malicious activity at Layer 2. A new security focus centers on attacks launched by maliciously leveraging normal Layer 2 switch operations. Security features exist to protect switches and Layer 2 operations, but, as with access control lists (ACLs) for upper-layer security, a policy must be established and appropriate features configured to protect against potential malicious acts while maintaining daily network operations.

Switch Attack Categories

This topic categorizes the types of switch security issues.

Switch Attack Categories

Cisco.com

- **MAC layer attacks**
- **VLAN attacks**
- **Spoofing attacks**
- **Attacks on switch devices**

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

Layer 2 malicious attacks are typically launched by a device connected to the campus network. This can be a physical rogue device placed on the network for malicious purposes or an external intrusion that takes control of and launches attacks from a trusted device. In either case, the network sees all traffic as originating from a legitimate connected device.

Attacks launched against switches and at Layer 2 can be grouped as follows:

- MAC layer attacks
- VLAN attacks
- Spoof attacks
- Attacks on switch devices

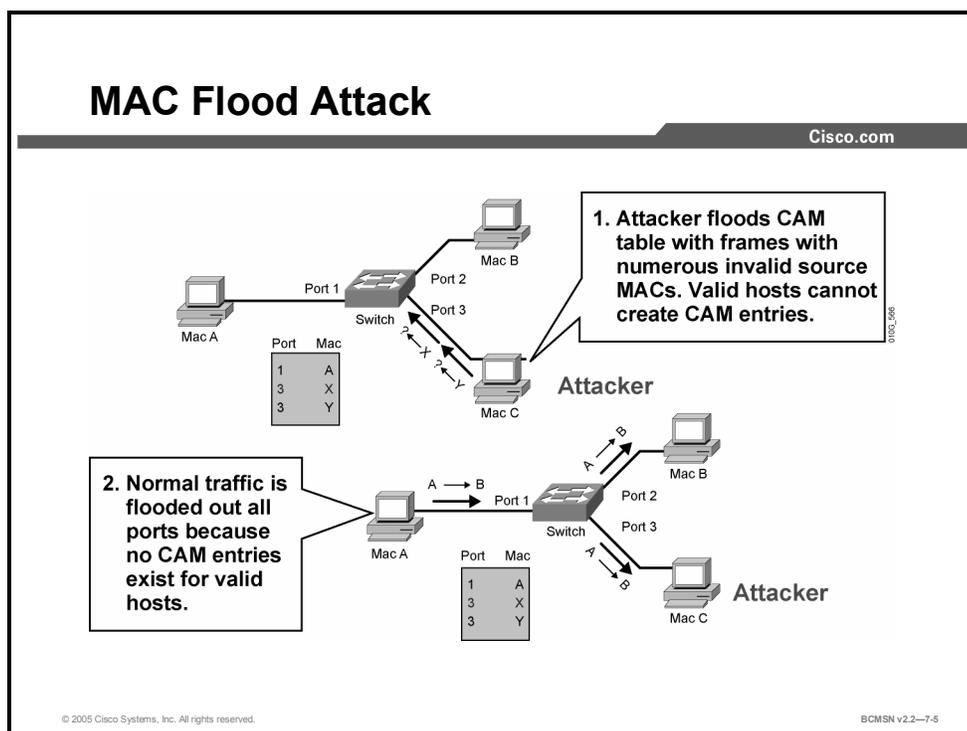
Significant attacks in these categories, known as of this writing, are discussed in more detail in subsequent sections of the course. Each attack method is accompanied by a standard measure for mitigating the security compromise.

Switch Security Concerns and Mitigation Steps

Attack Method	Description	Steps to Mitigation
MAC LAYER ATTACKS		
MAC address flooding	Frames with unique, invalid source MAC addresses flood the switch, exhausting content addressable memory (CAM) table space, disallowing new entries from valid hosts. Traffic to valid hosts is subsequently flooded out all ports.	Port security MAC address VLAN access maps
VLAN ATTACKS		
VLAN hopping	By altering the VLAN ID on packets encapsulated for trunking, an attacking device can send or receive packets on various VLANs, bypassing Layer 3 security measures.	Tighten up trunk configurations and the negotiation state of unused ports. Place unused ports in a common VLAN.
Attacks between devices on a common VLAN	Devices may need protection from one another, even though they are on a common VLAN. This is especially true on service provider segments supporting devices from multiple customers.	Implement private VLANs (PVLANS).
SPOOFING ATTACKS		
DHCP spoofing	An attacking device can exhaust the address space available to the DHCP servers for a period of time or establish itself as a DHCP server in man-in-the-middle attacks.	Use DHCP snooping.
Spanning tree compromises	Attacking device spoofs the root bridge in the Spanning Tree Protocol (STP) topology. If successful, the network attacker can see a variety of frames.	Proactively configure the primary and backup root devices. Enable root guard.
MAC spoofing	Attacking device spoofs the MAC address of a valid host currently in the CAM table. Switch then forwards frames destined for the valid host to the attacking device.	Use DHCP snooping, port security.
Address Resolution Protocol (ARP) spoofing	Attacking device crafts ARP replies intended for valid hosts. The attacking device's MAC address then becomes the destination address found in the Layer 2 frames sent by the valid network device.	Use Dynamic ARP Inspection DHCP snooping, port security.
SWITCH DEVICE ATTACKS		
Cisco Discovery Protocol (CDP) manipulation	Information sent through CDP is transmitted in clear text and unauthenticated, allowing it to be captured and divulge network topology information.	Disable CDP on all ports where it is not intentionally used.
Secure Shell (SSH) and Telnet attacks	Telnet packets can be read in clear text. SSH is an option but has security issues in version 1.	Use SSH version 2. Use Telnet with VTY ACLs

Describing a MAC Flooding Attack

This topic describes a MAC flooding attack.



A common Layer 2 or switch attack as of this writing is MAC flooding, resulting in CAM table overflow that causes flooding of regular data frames out all switch ports. This attack can be launched for the malicious purpose of collecting a broad sample of traffic or as a denial of service (DoS) attack.

CAM tables are limited in size, and therefore can contain only a limited number of entries at any one time. A network intruder can maliciously flood a switch with a large number of frames from a range of invalid source MAC addresses. If enough new entries are made before old ones expire, new valid entries will not be accepted. Then, when traffic arrives at the switch for a legitimate device that is located on one of the switch ports that was not able to create a CAM table entry, the switch must flood frames to that address out all ports. This has two adverse effects:

- The switch traffic forwarding is inefficient and voluminous.
- An intruding device can be connected to any switch port and capture traffic not normally seen on that port.

If the attack is launched before the beginning of the day, the CAM table would be full as the majority of devices are powered on. Then frames from those legitimate devices are unable to create CAM table entries as they power on. If this represents a large number of network devices, the number of MAC addresses for which traffic will be flooded will be high, and any switch port will carry flooded frames from a large number of devices.

If the initial flood of invalid CAM table entries is a one-time event, eventually, the switch will age out older, invalid CAM table entries, allowing new, legitimate devices to create an entry.

Traffic flooding will cease and may never be detected, as the intruder captured a significant amount of data from the network.

As the figure shows, MAC flooding occurs in the following progression:

MAC Flood Attack Progression

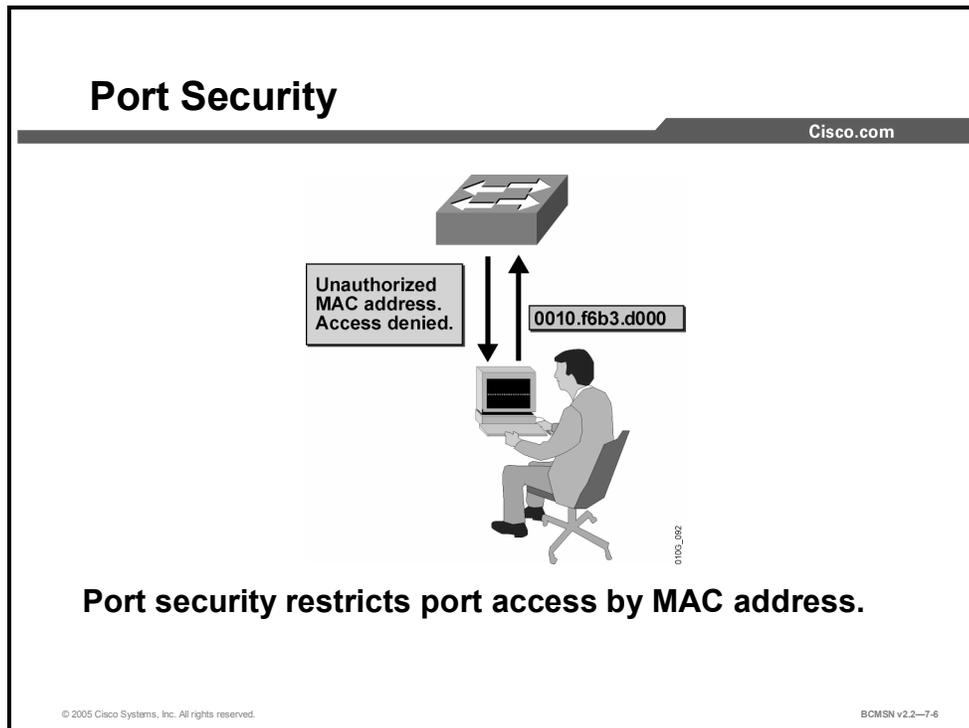
Step	Description
1.	Switch forwards traffic based on valid CAM table entries.
2.	Attacker (MAC address C) sends out multiple packets with various source MAC addresses.
3.	Over a short period of time, the CAM table in the switch fills up until it cannot accept new entries. As long as the attack is running, the CAM table on the switch will remain full.
4.	Switch begins to flood all packets that it receives out of every port so that frames sent from host A to host B are also flooded out of port 3 on the switch.

Suggested Mitigation for MAC Flood Attacks

Configure port security to define the number of MAC addresses that are allowed on a given port. Port security can also specify what MAC address is allowed on a given port.

Describing Port Security

Switch access security can be provided through port security. This topic describes port security.



Port security is a feature supported on Cisco Catalyst switches that restricts a switch port to a specific set or number of MAC addresses. Those addresses can be learned dynamically or configured statically. The port will then provide access only to frames from those addresses. If, however, the number of addresses is limited to four but no specific MAC addresses are configured, the port will allow any four MAC addresses to be learned dynamically and port access will be limited to those four dynamically learned addresses.

A port security feature called “sticky learning,” available on some switch platforms, combines the features of dynamically learned and statically configured addresses. When this feature is configured on an interface, the interface converts dynamically learned addresses to “sticky secure” addresses. This adds them to the running configuration as if they were configured using the **switchport port-security mac-address** command.

Scenario

Imagine five individuals whose laptops are allowed to connect to a specific switch port when they visit an area of the building. We want to restrict switch port access to the MAC addresses of those five laptops and allow no addresses to be learned dynamically on that port.

Process

Here is the process that can achieve the desired results for this scenario.

Implementing Port Security

Step	Action	Notes
1.	Configure port security.	Configure port security to allow only five connections on that port. Configure an entry for each of the five allowed MAC addresses. This, in effect, populates the MAC address table with five entries for that port and allows no additional entries to be learned dynamically.
2.	Allowed frames are processed.	When frames arrive on the switch port, their source MAC address is checked against the MAC address table. If the frame source MAC address matches an entry in the table for that port, the frames are forwarded to the switch to be processed like any other frames on the switch.
3.	New addresses are not allowed to create a new MAC address table entry.	When frames with a nonallowed MAC address arrive on the port, the switch determines that the address is not in the current MAC address table and does not create a dynamic entry for that new MAC address because the number of allowed addresses has been limited.
4.	Switch takes action in response to nonallowed frames.	The switch will disallow access to the port and take one of the following configuration-dependent actions: (a) the entire switch port can be shut down; (b) access can be denied for that MAC address only and a log error can be generated; (c) access can be denied for that MAC address but no log message generated.

Note Port security cannot be applied to trunk ports where addresses might change frequently. Implementations of port security vary by Catalyst platform. Check documentation to see if and how particular hardware supports this feature.

References

For additional information, refer to these resources:

http://cisco.com/en/US/products/hw/switches/ps663/products_configuration_guide_chapter09186a00801162fc.html

http://cisco.com/en/US/products/hw/switches/ps4324/products_configuration_guide_chapter09186a008019d0de.html

Port Security Configuration Commands

This topic describes port security configuration commands.

Port Security Configuration Commands

Cisco.com

```
Switch(config-if)#switchport port-security [maximum value]
violation {protect | restrict | shutdown}
```

Enables port security and specifies the maximum number of MAC addresses that can be supported by this port

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

The table reviews the commands for implementing, verifying, and removing port security on a switch.

Note Port security is configured only on access ports.

Port Security Commands

Command	Description
Switch(config-if)# switchport port security [maximum value]	Enables port security on an interface and can also specify the maximum number of MAC addresses that can be supported by this port. Set maximum MAC addresses for the port.
Switch(config-if)# switchport port-security mac-address	Specifies access by a specific MAC address.
Switch(config-if)# switchport port-security violation	Specifies what action is to be taken when access is attempted by a MAC address not in the MAC address table for the interface.
Switch# show port-security	Displays the port security settings defined for an interface or for the switch.
Switch# clear port-security dynamic	Deletes a specific dynamic secure address or all the dynamic secure addresses on an interface from the MAC address table.

How to Configure Port Security on a Switch

This topic explains how to configure port security on a Catalyst switch.

Configuring Port Security on a Switch

Cisco.com

- **Enable port security.**
- **Set MAC address limit.**
- **Specify allowable MAC addresses.**
- **Define violation actions.**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-7.8

Here are the steps to set up port security that will limit switch port access to a finite number and a specific set of end-device MAC addresses.

Port Security Configuration Steps

Step	Description
1.	Enables port security <code>Switch(config-if)#switchport port-security</code>
2.	Sets a maximum number of MAC addresses that will be allowed on this port. Default is one. <code>Switch(config-if)#switchport port-security maximum value</code>
3.	Specifies which MAC addresses will be allowed on this port (optional). <code>Switch(config-if)#switchport port-security mac-address mac-address</code> <code>Switch(config-if)#switchport port-security mac-address mac-address</code>
4.	Defines what action an interface will take if a nonallowed MAC address attempts access <code>Switch(config-if)#switchport port-security violation {shutdown restrict protect}</code>

Caveats to Port Security Configuration Steps

- Step 1** Port security is enabled on a port-by-port basis.
- Step 2** By default, only one MAC address is allowed access through a given switch port when port security is enabled. This parameter increases that number. It implies no restriction on specific MAC addresses, just on the total number of addresses that can be learned by the port. Learned addresses are not aged out by default but can be configured to do so after a specified time. The *value* parameter can be any number from 1 to 1024, with some restrictions having to do with the number of ports on a given switch with port security enabled.
- Step 3** Access to the switch port can be restricted to one or more specific MAC addresses. If the number of specific MAC addresses assigned using this command is lower than the *value* parameter set in Step 2, then the remaining allowed addresses can be learned dynamically. If you specify a set of MAC addresses that is equal to the maximum number allowed, access is limited to that set of MAC addresses.
- Step 4** By default, if the maximum number of connections is achieved and a new MAC address attempts to access the port, the switch must take one of the following actions:
- **Protect:** Frames from the nonallowed address are dropped, but there is no log of the violation.

Note The *protect* argument is platform or version dependent.

- **Restrict:** Frames from the nonallowed address are dropped, and a log message is created.
- **Shut down:** If any frames are seen from a nonallowed address, the interface is errdisabled, a log entry is made, and manual intervention or errdisable recovery must be used to make the interface usable.

How to Verify Port Security

This subtopic describes how to verify port security

Verifying Port Security

Cisco.com

```
Switch#show port-security
```

- **Displays security information for all interfaces**

```
Switch#show port-security
Secure Port      MaxSecureAddr  CurrentAddr  SecurityViolation  Security Action
              (Count)          (Count)      (Count)
-----
Fa5/1             11              11           0                  Shutdown
Fa5/5             15              5            0                  Restrict
Fa5/11            5               4            0                  Protect
-----

Total Addresses in System: 21
Max Addresses limit in System: 128
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-7.9

Use **show** commands to verify the configuration of port security.

Verifying Network Access Security

The **show port-security** command can be used to verify the ports on which port security has been enabled. It also displays count information and security actions to be taken per interface.

The full command syntax is as follows:

```
Switch# show port-security [interface interface_id] address
```

Arguments are provided to view port security status by interface or view the addresses associated with port security on all interfaces.

Example: show port-security Command Output

This example displays output from the **show port-security** command when you do not enter an interface:

```
Switch#show port-security
Secure Port      MaxSecureAddr  CurrentAddr  SecurityViolation
Security Action  (Count)        (Count)      (Count)
-----
Fa5/1            11             11           0
Shutdown
Fa5/5            15             5            0
Restrict
Fa5/11           5              4            0
Protect
-----
Total Addresses in System: 21
Max Addresses limit in System: 128
```

Verifying Port Security (Cont.)

Cisco.com

```
Switch#show port-security interface type mod/port
```

- **Displays security information for a specific interface**

```
Switch#show port-security interface fastethernet 5/1
Port Security: Enabled
Port status: SecureUp
Violation mode: Shutdown
Maximum MAC Addresses: 11
Total MAC Addresses: 11
Configured MAC Addresses: 3
Aging time: 20 mins
Aging type: Inactivity
SecureStatic address aging: Enabled
Security Violation count: 0
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2--7-10

Use the *interface* argument to provide output for a specific interface.

Example: show port-security Command for a Specific Interface

This example displays output from the `show port-security` command for a specified interface:

```
Switch#show port-security interface fastethernet 5/1
```

```
Port Security: Enabled
Port status: SecureUp
Violation mode: Shutdown
Maximum MAC Addresses: 11
Total MAC Addresses: 11
Configured MAC Addresses: 3
Aging time: 20 mins
Aging type: Inactivity
SecureStatic address aging: Enabled
Security Violation count: 0
```

Verifying Port Security (Cont.)

Cisco.com

```
Switch#show port-security address
```

- **Displays MAC address table security information**

```
Switch#show port-security address
Secure Mac Address Table
-----
Vlan    Mac Address      Type                Ports    Remaining Age
        (mins)
-----
1       0001.0001.0001   SecureDynamic       Fa5/1    15 (I)
1       0001.0001.0002   SecureDynamic       Fa5/1    15 (I)
1       0001.0001.1111   SecureConfigured    Fa5/1    16 (I)
1       0001.0001.1112   SecureConfigured    Fa5/1    -
1       0001.0001.1113   SecureConfigured    Fa5/1    -
1       0005.0005.0001   SecureConfigured    Fa5/5    23
1       0005.0005.0002   SecureConfigured    Fa5/5    23
1       0005.0005.0003   SecureConfigured    Fa5/5    23
1       0011.0011.0001   SecureConfigured    Fa5/11   25 (I)
1       0011.0011.0002   SecureConfigured    Fa5/11   25 (I)
-----
Total Addresses in System: 10
Max Addresses limit in System: 128
```

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—7-11

Use the *address* argument to display MAC address table security information. The remaining age column will only be populated if specifically configured for a given interface.

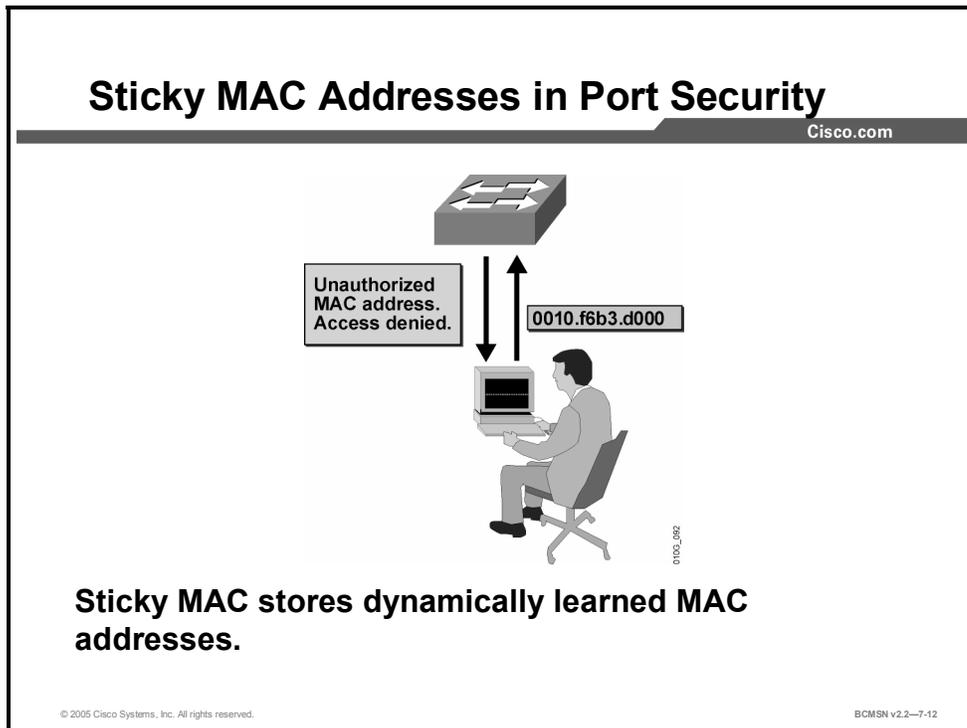
Example: Displaying MAC Address Table Security Information

This example displays output from the **show port-security address** privileged EXEC command:

```
Switch#show port-security address
Secure Mac Address Table
-----
Vlan      Mac Address      Type                Ports      Remaining Age
          (mins)
-----
1         0001.0001.0001   SecureDynamic       Fa5/1      15 (I)
1         0001.0001.0002   SecureDynamic       Fa5/1      15 (I)
1         0001.0001.1111   SecureConfigured    Fa5/1      16 (I)
1         0001.0001.1112   SecureConfigured    Fa5/1      -
1         0001.0001.1113   SecureConfigured    Fa5/1      -
1         0005.0005.0001   SecureConfigured    Fa5/5      23
1         0005.0005.0002   SecureConfigured    Fa5/5      23
1         0005.0005.0003   SecureConfigured    Fa5/5      23
1         0011.0011.0001   SecureConfigured    Fa5/11     25 (I)
1         0011.0011.0002   SecureConfigured    Fa5/11     25 (I)
-----
Total Addresses in System: 10
Max Addresses limit in System: 128
```

Port Security with Sticky MAC Addresses

This topic describes the sticky MAC address feature of port security.



Port security can be used to mitigate spoof attacks by limiting access through each switch port to a single MAC address. This prevents intruders from using multiple MAC addresses over a short period of time but does not limit port access to a specific MAC address. The most restrictive port security implementation would specify the exact MAC address of the single device that is to gain access through each port. Implementing this level of security, however, requires considerable administrative overhead.

Port security has a feature called “sticky MAC addresses” that can limit switch port access to a single, specific MAC address without the network administrator having to gather the MAC address of every legitimate device and manually associate it with a particular switch port.

When sticky MAC addresses are used, the switch port will convert dynamically learned MAC addresses to sticky MAC addresses and subsequently add them to the running configuration as if they were static entries for a single MAC address to be allowed by port security. Sticky secure MAC addresses will be added to the running configuration but will not become part of the startup configuration file unless the running configuration is copied to the startup configuration after addresses have been learned. If they are saved in the startup configuration, they will not have to be relearned upon switch reboot, and this provides a higher level of network security.

The following command will convert all dynamic port security–learned MAC addresses to sticky secure MAC addresses.

```
switchport port-security mac-address sticky
```

This command cannot be used on ports where voice VLANs are configured.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Security measures must be taken to protect switch functions used for Layer 2 access.**
- **Switch attacks fall into four main categories.**
- **MAC flood attacks are launched against Layer 2 access switches.**
- **Port security can be configured to mitigate MAC flood attacks.**
- **Sticky MAC addresses allow port security to limit access to a specific, dynamically learned MAC address.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—7-13

References

For additional information, refer to these resources:

- Cisco Systems, Inc., *SAFE Layer 2 Security In-depth Version 2*:
http://www.cisco.com/en/US/netsol/ns340/ns394/ns171/ns128/networking_solutions_white_paper09186a008014870f.shtml
- *Packet Magazine*, First Quarter 2003, “Layer2: The Weakest Link”:
http://cisco.com/en/US/partner/about/ac123/ac114/ac173/ac222/about_cisco_packet_feature09186a0080142deb.html

Mitigating VLAN Attacks

Overview

On networks using trunking protocols, there is a possibility of rogue traffic “hopping” from one VLAN to another, thereby creating security vulnerabilities. These attacks are best mitigated by close control of trunk links. Private VLANs (PVLANS) can be configured to establish security regions within a single VLAN without subnetting.

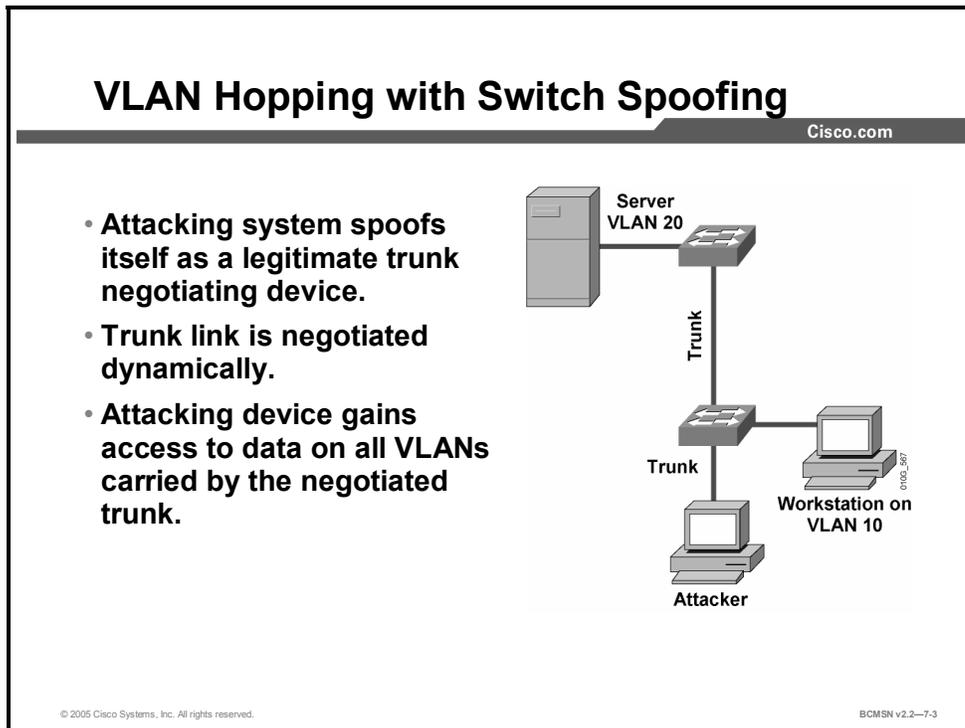
Objectives

Upon completing this lesson, you will be able to configure various features to prevent VLAN hopping and address VLAN security issues. This ability includes being able to meet these objectives:

- Describe how VLAN hopping occurs and why this is a security vulnerability
- List the commands used to mitigate VLAN attacks
- Configure a switch to mitigate VLAN attacks
- Define the purpose of a PVLAN
- Configure PVLANS as a means of network security
- Define the purpose of VACLs

What Is VLAN Hopping?

This topic describes the process of VLAN hopping.



VLAN hopping is a network attack whereby an end system sends packets to, or collects packets from, a VLAN that should not be accessible to that end system. This is accomplished by tagging the invasive traffic with a specific VLAN ID or by negotiating a trunk link in order to send or receive traffic on penetrated VLANs. VLAN hopping can be accomplished by switch spoofing or double tagging.

Switch Spoofing

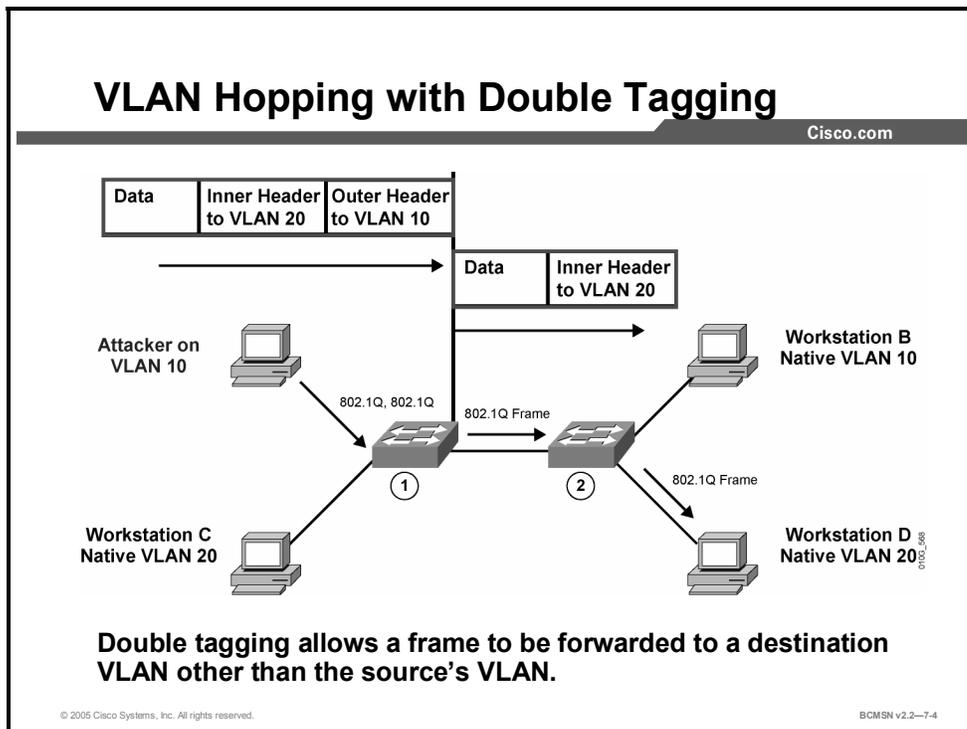
In a switch spoofing attack, the network attacker configures a system to spoof itself as a switch by emulating Inter-Switch Link (ISL) or 802.1Q signaling along with Dynamic Trunk Protocol (DTP) signaling an attempt to establish a trunk connection to the switch. Any switch port configured as DTP auto, upon receipt of a DTP packet generated by the attacking device, may become a trunk port and thereby accept traffic destined for any VLAN supported on that trunk. The malicious device can then send packets to, or collect packets from, any VLAN carried on the negotiated trunk.

Switch Spoofing Sequence of Events

Step	Description
1	Attacker gains access to a switch port and sends DTP negotiation frames toward a switch with DTP running and <i>auto</i> negotiation turned on (often, the default settings).
2	Attacker and switch negotiate trunking over the port.
3	Switch allows all VLANs (default) to traverse the trunk link.
4	Attacker sends data to, or collects it from, all VLANs carried on that trunk.

Double Tagging

This subtopic describes double tagging as a means of VLAN hopping.



Another method of VLAN hopping is for any workstation to generate frames with two 802.1Q headers in order to get the switch to forward the frames onto a VLAN that would be inaccessible to the attacker through legitimate means.

The first switch to encounter the double-tagged frame strips the first tag off the frame as it enters the switch because it matches the access ports native VLAN and then forwards the frame. The result is that the frame is forwarded with the inner 802.1Q tag out all the switch ports including trunk ports configured with the native VLAN of the network attacker. The second switch then forwards the packet to the destination based on the VLAN identifier in the second 802.1Q header. Should the trunk not match the native VLAN of the attacker, the frame would be untagged and flooded only to the original VLAN.

Double Tagging Method of VLAN Hopping

Step	Description
1	Workstation A (native VLAN 10) sends a frame with two 802.1Q headers to switch 1.
2	Switch 1 strips the outer tag and forwards the frame to all ports within same native VLAN.
3	Switch 2 interprets the frame according to information in the inner tag marked with VLAN ID 20.
4	Switch 2 forwards the frame out all ports associated with VLAN 20, including trunk ports.

How to Mitigate VLAN Hopping

This topic describes how to mitigate VLAN hopping attacks.

Defending Against VLAN Attacks

Cisco.com

```
Switch(config)# interface-range type mod/port-port
```

- **Selects a range of interfaces to configure**

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

- **Configures the ports as access ports and turns off DTP**

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

- **Statically assigns the ports to specific unused VLAN**

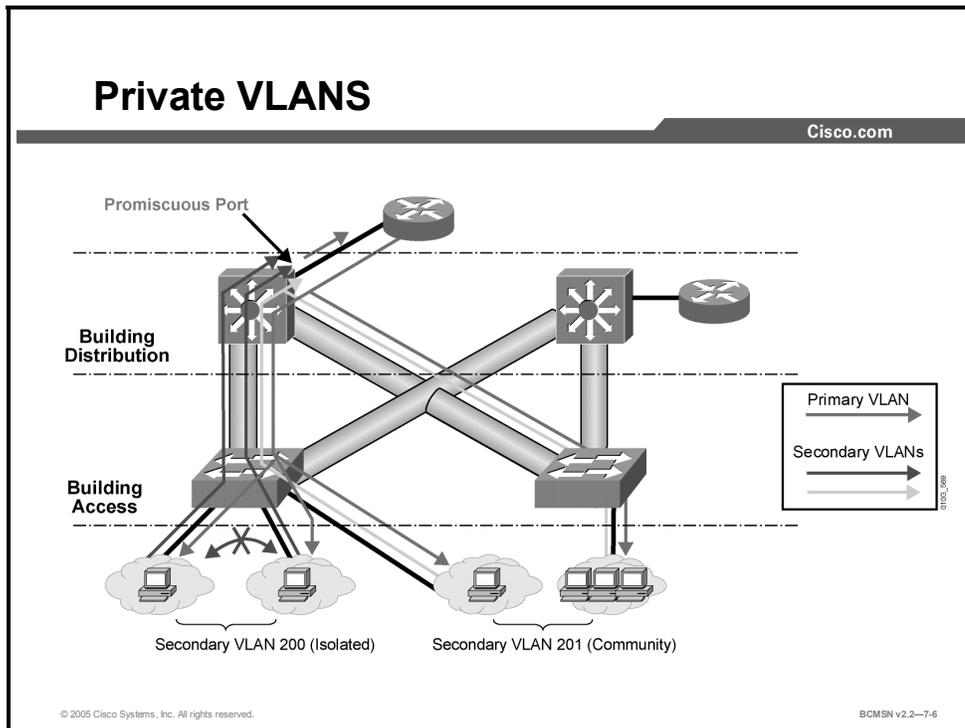
© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—7-5

The measures to defend the network from VLAN hopping are a series of best practices for all switch ports and parameters to follow when establishing a trunk port.

- Configure all unused ports as access ports so that trunking cannot be negotiated across those links.
- Place all unused ports in the shutdown state and associate with a VLAN designed only for unused ports, carrying no user data traffic.
- When establishing a trunk link, purposefully configure the following:
 - the native VLAN to be different from any data VLANs
 - trunking as “on,” rather than negotiated
 - the specific VLAN range to be carried on the trunk

What Is a Private VLAN?

This topic describes a private VLAN.



Service providers often have devices from multiple clients, as well as their own servers, on a single Demilitarized Zone (DMZ) segment or VLAN. As security issues proliferate, it becomes necessary to provide traffic isolation between devices even though they may exist on the same Layer 3 segment and VLAN. Catalyst 6500/4500 switches implement private VLANs (PVLANS) to keep some switch ports shared and some switch ports isolated, although all ports exist on the same VLAN. The 2950 and 3550 support “protected ports,” which is functionality similar to PVLANS on a per-switch basis.

The traditional solution to address these Internet service provider (ISP) requirements is to provide one VLAN per customer, with each VLAN having its own IP subnet. A Layer 3 device then provides interconnectivity between VLANs and Internet destinations.

Here are the challenges with this traditional solution:

- Supporting a separate VLAN per customer may require a high number of interfaces on service provider network devices.
- Spanning tree becomes more complicated with many VLAN iterations.
- Network address space must be divided into many subnets, which wastes space and increases management complexity.
- Multiple ACL applications are required to maintain security on multiple VLANs, resulting in increased management complexity.

PVLANS provide Layer 2 isolation between ports within the same VLAN. This isolation eliminates the need for a separate VLAN and IP subnet per customer.

PVLAN Port Types

This subtopic discusses PVLAN port types.

PVLAN Port Types

Cisco.com

- **Isolated: Communicate only with promiscuous ports**
- **Promiscuous: Communicate with all other ports**
- **Community: Communicate with other members of community and all promiscuous ports**

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

A port in a PVLAN can be one of three types:

- **Isolated:** An isolated port has complete Layer 2 separation from other ports within the same PVLAN except for the promiscuous port. PVLANS block all traffic to isolated ports, except the traffic from promiscuous ports. Traffic received from an isolated port is forwarded only to promiscuous ports.
- **Promiscuous:** A promiscuous port can communicate with all ports within the PVLAN, including the community and isolated ports. The default gateway for the segment would likely be hosted on a promiscuous port, given that all devices in the PVLAN will need to communicate with that port.
- **Community:** Community ports communicate among themselves and with their promiscuous ports. These interfaces are isolated at Layer 2 from all other interfaces in other communities, or in isolated ports within their PVLAN.

Note Because trunks can support the VLANs carrying traffic between isolated, community, and promiscuous ports, isolated and community port traffic might enter or leave the switch through a trunk interface.

Resources

For additional information, refer to these resources:

Configuring Private VLANs (6500 Series):

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a0080160a44.html

Configuring Private VLANs (4500 Series):

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/12_1_13/config/pvlans.htm

Configuring PVLANS

This topic describes configuring PVLANS.

Configuring Private VLANs

Cisco.com

```
Switch(config-vlan)#private-vlan [primary | isolated | community]
```

- **Configures a VLAN as a private VLAN**

```
Switch(config-vlan)#private-vlan association {secondary_vlan_list | add svl | remove svl}
```

- **Associates secondary VLANs with the primary VLAN**

```
Switch#show vlan private-vlan type
```

- **Verifies private VLAN configuration**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—7-8

To configure a PVLAN, follow these steps.

Step 1 Set VTP mode to transparent.

Step 2 Create the secondary VLANs.

Note Isolated and community VLANs are secondary VLANs.

Step 3 Create the primary VLAN.

Step 4 Associate the secondary VLAN with the primary VLAN. Only one isolated VLAN can be mapped to a primary VLAN, but more than one community VLAN can be mapped to a primary VLAN.

Step 5 Configure an interface as an isolated or community port.

Step 6 Associate the isolated port or community port with the primary-secondary VLAN pair.

Step 7 Configure an interface as a promiscuous port.

Step 8 Map the promiscuous port to the primary-secondary VLAN pair.

Use these commands to configure a VLAN as a PVLAN:

```
Switch(config)#vlan vlan_ID
Switch(config-vlan)#[no] private-vlan {isolated | primary}
```

Example: PVLAN Configurations

This example shows how to configure VLAN202 as a primary VLAN and verify the configuration:

```
Switch#configure terminal
Switch(config)#vlan 202
Switch(config-vlan)#private-vlan primary
Switch(config-vlan)#end
Switch#show vlan private-vlan type
```

Primary	Secondary	Type	Interfaces
202		primary	

This example shows how to configure VLAN 200 as an isolated VLAN and verify the configuration:

```
Switch#configure terminal
Switch(config)#vlan 200
Switch(config-vlan)#private-vlan isolated
Switch(config-vlan)#end
Switch#show vlan private-vlan type
```

Primary	Secondary	Type	Interfaces
202		primary	
200		isolated	

To associate secondary VLANs with a primary VLAN, perform this procedure:

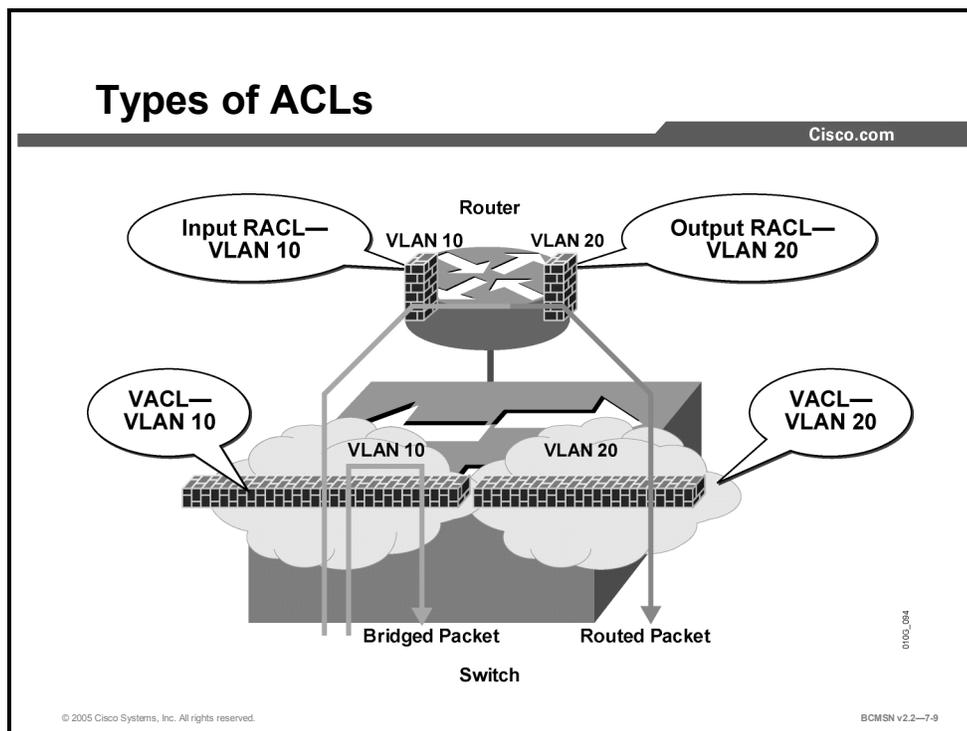
```
Switch(config)#vlan primary_vlan_ID
Switch(config-vlan)#[no] private-vlan association {secondary_vlan_list
| add secondary_vlan_list | remove secondary_vlan_list}
```

When you associate secondary VLANs with a primary VLAN, note the following:

- The *secondary_vlan_list* parameter contains only one isolated VLAN ID.
- Use the **remove** keyword with the *secondary_vlan_list* variable to clear the association between the secondary VLAN and the primary VLAN. The list can contain only one VLAN.
- Use the **no** keyword to clear all associations from the primary VLAN.
- The command does not take effect until you exit VLAN configuration submode.

Configuring VLAN Security Using Access Lists

ACLs are useful for controlling access in a multilayer switched network. This topic explains how to configure security with ACLs.



Cisco multilayer switches support three types of ACLs:

- **Router access control lists (RACLs):** Supported in the ternary content addressable memory (TCAM) hardware on Cisco multilayer switches
- **Quality of service (QoS) access control lists:** Supported in the TCAM hardware on Cisco multilayer switches
- **VLAN access control lists (VACLs):** Supported in software on Cisco multilayer switches

Catalyst switches support four ACL lookups per packet: input and output security ACL, and input and output QoS ACL.

Catalyst switches use two methods of performing a merge: order independent and order dependent. With order independent merge, ACLs are transformed from a series of order dependent actions to a set of order independent masks and patterns. The resulting access control entry can be very large. The merge is processor- and memory-intensive.

Order-dependent merge is a recent improvement on some Catalyst switches in which ACLs retain their order-dependent aspect. The computation is much faster and is less processor-intensive.

RACLs are supported in hardware through IP standard ACLs and IP extended ACLs, with permit and deny actions. ACL processing is an intrinsic part of the packet forwarding process. ACL entries are programmed in hardware. Lookups occur in the pipeline whether ACLs are configured or not. With RACLs, access list statistics and logging are not supported.

Configuring VACLs

Cisco.com

```
Switch(config)#vlan access-map map_name [seq#]
```

- **Defines a VLAN access map**

```
Switch(config-access-map)# match {ip address {1-199 |  
1300-2699 | acl_name} | ipx address {800-999 | acl_name} |  
mac address acl_name}
```

- **Configures the match clause in a VLAN access map sequence**

```
Switch(config-access-map)#action {drop [log]} | {forward  
[capture]} | {redirect {type slot/port} | {port-channel  
channel_id}}
```

- **Configures the action clause in a VLAN access map sequence**

```
Switch(config)#vlan filter map_name vlan_list list
```

- **Applies the VLAN access map to the specified VLANs**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—7-10

VACLs (also called VLAN access maps in Cisco IOS software) apply to all traffic on the VLAN. They filter on the basis of Ethertype and MAC address traffic.

VACLs follow route-map conventions, in which map sequences are checked in order.

When a matching permit access control entry (ACE) is encountered, the switch takes the action. When a matching deny ACE is encountered, the switch checks the next ACL in the sequence or checks the next sequence.

Three VACL actions are permitted:

- **Permit** (with capture, Catalyst 6500 only)
- **Redirect** (Catalyst 6500 only)
- **Deny** (with logging, Catalyst 6500 only)

The VACL capture option copies traffic to specified capture ports. VACL ACEs installed in hardware are merged with RACLs and other features.

Two features are supported only on the Catalyst 6500:

- **VACL capture:** Forwarded packets are captured on capture ports. The capture option is only on permit ACEs. The capture port can be an IDS monitor port or any Ethernet port. The capture port must be in an output VLAN for Layer 3–switched traffic.
- **VACL redirect:** Matching packets are redirected to specified ports. You can configure up to five redirect ports. Redirect ports must be in a VLAN where VACL is applied.

To configure VACLs, complete these steps:

Step	Description
1.	Define a VLAN access map. <code>Switch(config)#vlan access-map map_name [seq#]</code>
2.	Configure a match clause. <code>Switch(config-access-map)# match {ip address {1-199 1300-2699 acl_name} ipx address {800-999 acl_name} mac address acl_name}</code>
3.	Configure an action clause. <code>Switch(config-access-map)#action {drop [log]} {forward [capture]} {redirect {{fastethernet gigabitethernet tengigabitethernet} slot/port} {port-channel channel_id}}</code>
4.	Apply a map to VLANs. <code>Switch(config)#vlan filter map_name vlan_list list</code>
5.	Verify the VACL configuration. <code>Switch#show vlan access-map map_name</code> <code>Switch#show vlan filter [access-map map_name vlan_id]</code>

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **VLAN hopping can allow Layer 2 unauthorized access to another VLAN.**
- **VLAN hopping can be mitigated by**
 - Properly configuring 802.1Q trunks
 - Turning off trunk negotiation
- **PVLANs are configured to allow traffic flows to be restricted between ports within the same VLAN.**
- **Access lists can be applied in a specific manner in order to protect VLANs.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—7-11

References

For additional information, refer to these resources:

- Cisco Systems, Inc., *VLAN Security White Paper*
http://cisco.com/en/US/partner/products/hw/switches/ps708/products_white_paper09186a008013159f.shtml
- Cisco Systems, Inc., Configuring Private VLANs (4500 series)
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/12_1_13/config/pvlans.htm

Mitigating Spoof Attacks

Overview

DHCP, MAC, and Address Resolution Protocol (ARP) spoofing are all methods used to gain unauthorized access to a network or to redirect traffic for malicious purposes. DHCP snooping, port security, and dynamic ARP inspection can be configured to guard against these threats.

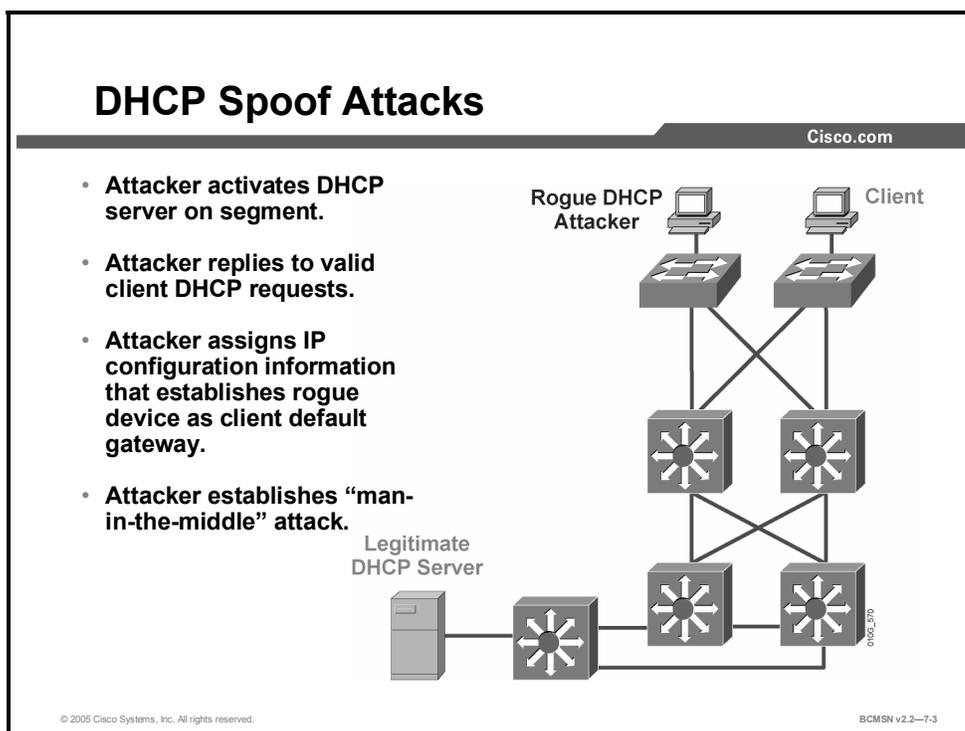
Objectives

Upon completing this lesson, you will be able to configure DHCP snooping, and enable port security and dynamic ARP inspection to mitigate DHCP, MAC, and ARP spoofing. This ability includes being able to meet these objectives:

- Describe a DHCP spoof attack
- Describe how DHCP snooping provides security against DHCP spoof attacks
- Identify the commands used to configure DHCP snooping
- Configure DHCP snooping on a Catalyst switch
- Describe a MAC spoof attack
- Describe an ARP spoofing attack
- Describe how DAI functions to defend against ARP spoofing
- Configure DAI

Describing a DHCP Spoof Attack

This topic describes a DHCP spoof attack.



One of the ways an attacker can gain access to network traffic is to spoof responses that would be sent by a valid DHCP server. The DHCP spoofing device replies to client DHCP requests. The legitimate server may reply as well, but if the spoofing device is on the same segment as the client, its reply to the client may arrive first. The intruder’s DHCP reply offers an IP address and supporting information that designates the intruder as the default gateway or Domain Name System (DNS) server. In the case of a gateway, the clients will then forward all packets to the attacking device, which will in turn send them to the desired destination. This is referred to as a “man-in-the-middle” attack, and it may go entirely undetected as the intruder intercepts the data flow through the network.

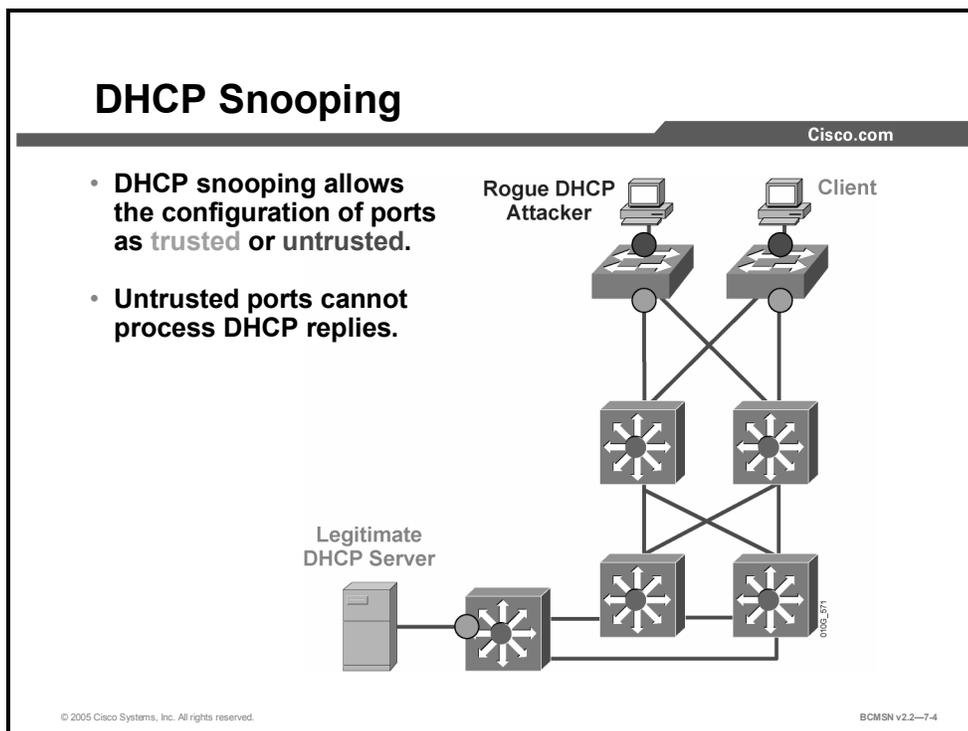
This table describes the DHCP spoofing attack sequence, as shown in the figure.

DHCP Spoof Attack Sequence

Sequence of Events	Description
1.	Attacker hosts a rogue DHCP server off a switch port.
2.	Client broadcasts a request for DHCP configuration information.
3.	The rogue DHCP server responds before the legitimate DHCP server, assigning attacker-defined IP configuration information.
4.	Host packets are redirected to the attacker’s address as it emulates a default gateway for the erroneous DHCP address provided to the client.

Describing DHCP Snooping

This topic describes the DHCP snooping feature of the Catalyst family of switches.



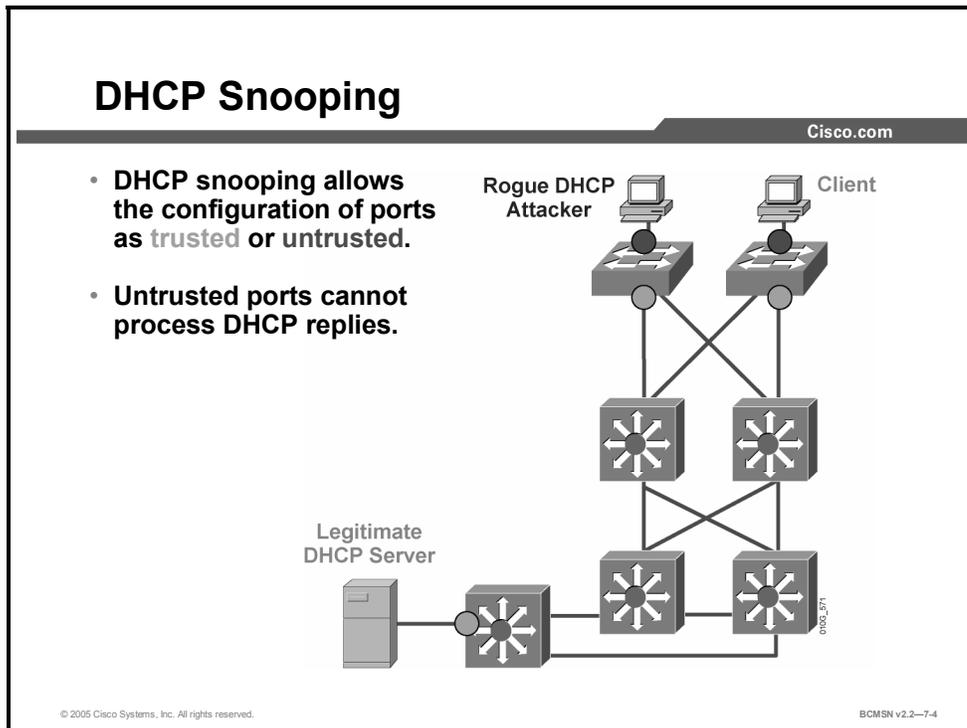
DHCP snooping is a Catalyst feature that determines which switch ports can respond to DHCP requests. Ports are identified as trusted and untrusted. Trusted ports can source all DHCP messages, while untrusted ports can source requests only. Trusted ports host a DHCP server or can be an uplink toward the DHCP server. If a rogue device on an untrusted port attempts to send a DHCP response packet into the network, the port is shut down. This feature can be coupled with DHCP Option 82, in which switch information, such as the port ID of the DHCP request, can be inserted into the DHCP request packet.

Untrusted ports are those not explicitly configured as trusted. A DHCP binding table is built for untrusted ports. Each entry contains client MAC address, IP address, lease time, binding type, VLAN number, and port ID recorded as clients make DHCP requests. The table is then used to filter subsequent DHCP traffic. From a DHCP snooping perspective, untrusted access ports should not send any DHCP server responses, such as DHCP Offer, DHCPACK, or DHCPNAK.

Sequence of Configuration	Description
1.	Configure global DHCP snooping
2.	Configure trusted ports
3.	Configure Option 82 insertion off (default enabled by step 2)
4.	Configure rate limiting on untrusted ports
5.	Configure DHCP snooping for the selected VLANs

DHCP Snooping Configuration Commands

This topic describes the commands used to configure DHCP snooping.



To configure DHCP snooping, use these commands:

DHCP Snooping Commands

Command	Description
Switch(config) # ip dhcp snooping	Enables DHCP snooping globally
Switch(config-if) # ip dhcp snooping trust	Configures an interface as trusted
Switch(config) # ip dhcp snooping limit rate rate	Configures the number of packets per second (pps) that an interface can receive
Switch(config) # ip dhcp snooping vlan number [number]	Enables DHCP snooping on a VLAN or a range of VLANs

References

For additional information, refer to this resource:

Understanding and Configuring DHCP Snooping:

http://cisco.com/en/US/partner/products/hw/switches/ps663/products_configuration_guide_chapter09186a00800dde9f.html#30724

How to Configure DHCP Snooping

This topic describes how to use the commands to configure DHCP snooping on a Catalyst switch.

Securing Against DHCP Snooping Attacks

Cisco.com

```
Switch(config)# ip dhcp snooping
```

- Enables DHCP snooping globally

```
Switch(config)# ip dhcp snooping information option
```

- Enables DHCP Option 82 data insertion

```
Switch(config-if)# ip dhcp snooping trust
```

- Configures a trusted interface

```
Switch(config)# ip dhcp snooping vlan number [number]
```

- Enables DHCP snooping on your VLANs

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—7-5

Steps for Enabling DHCP Snooping

Step	Comments
1. Enable DHCP snooping globally. <pre>Switch(config)# ip dhcp snooping</pre>	By default the feature is not enabled.
2. Enable DHCP Option 82. <pre>Switch(config)# ip dhcp snooping information option</pre>	This is optional for the packet to contain information on the switch port where it originated.
3. Configure DHCP server interfaces or uplink ports as trusted. <pre>Switch(config-if)# ip dhcp snooping trust</pre>	At least one trusted port must be configured. Use the no keyword to revert to untrusted.
4. Configure the number of DHCP packets per second that are acceptable on the port. <pre>Switch(config-if)# ip dhcp snooping limit rate rate</pre>	Configure the number of DHCP pps that an interface can receive. Normally, the rate limit applies to untrusted interfaces.
5. Enable DHCP snooping on specific VLAN(s). <pre>Switch(config)# ip dhcp snooping vlan number [number]</pre>	This is required to identify those VLANs that will be subject to DHCP snooping.
7. Verify the configuration. <pre>Switch# show ip dhcp snooping</pre>	Verify the configuration.

Verifying the DHCP Snooping Configuration

Verifying DHCP Snooping

Cisco.com

```
Switch# show ip dhcp snooping
```

- Verifies the DHCP snooping configuration

```
Switch# show ip dhcp snooping
Switch DHCP snooping is enabled
DHCP Snooping is configured on the following VLANs:
  10 30-40 100 200-220
Insertion of option 82 information is enabled.
Interface          Trusted          Rate limit (pps)
-----
FastEthernet2/1    yes             none
FastEthernet2/2    yes             none
FastEthernet3/1    no              20
Switch#
```

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-7.6

This example shows how to display the DHCP snooping configuration for a switch:

```
Switch# show ip dhcp snooping
```

```
Switch DHCP snooping is enabled.
```

```
DHCP Snooping is configured on the following VLANs:
```

```
  10 30-40 100 200-220
```

```
Insertion of option 82 information is enabled.
```

```
Interface          Trusted          Rate limit (pps)
```

```
-----
```

```
FastEthernet2/1    yes             none
```

```
FastEthernet3/1    yes             none
```

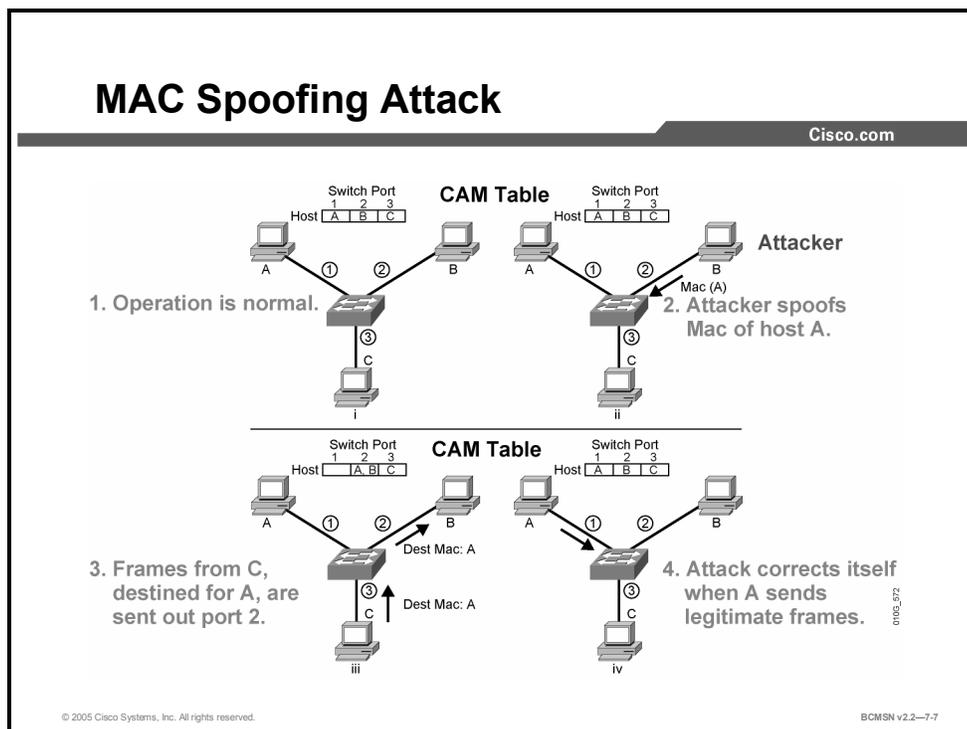
```
GigabitEthernet1/1 no              20
```

```
Switch#
```

Only ports that are trusted or that have a rate limit applied will be shown in the output. All other ports are untrusted and are not displayed.

Describing a MAC Spoof Attack

This topic describes what occurs during a MAC spoof attack.



MAC spoofing attacks occur when a device spoofs the MAC address of a valid network device to gain access to frames not normally forwarded out the switch port of the attacker. The attacker generates a single frame with a source MAC address of the valid device. The switch overwrites the valid content addressable memory (CAM) table entry with an entry for the same MAC address out the port of the attacking device. This causes the switch to forward frames destined for the valid MAC address out the port of the network attacker. Once the valid host sends additional frames, the spoofed CAM table entry is overwritten, so forwarding to that MAC address resumes on the legitimate port.

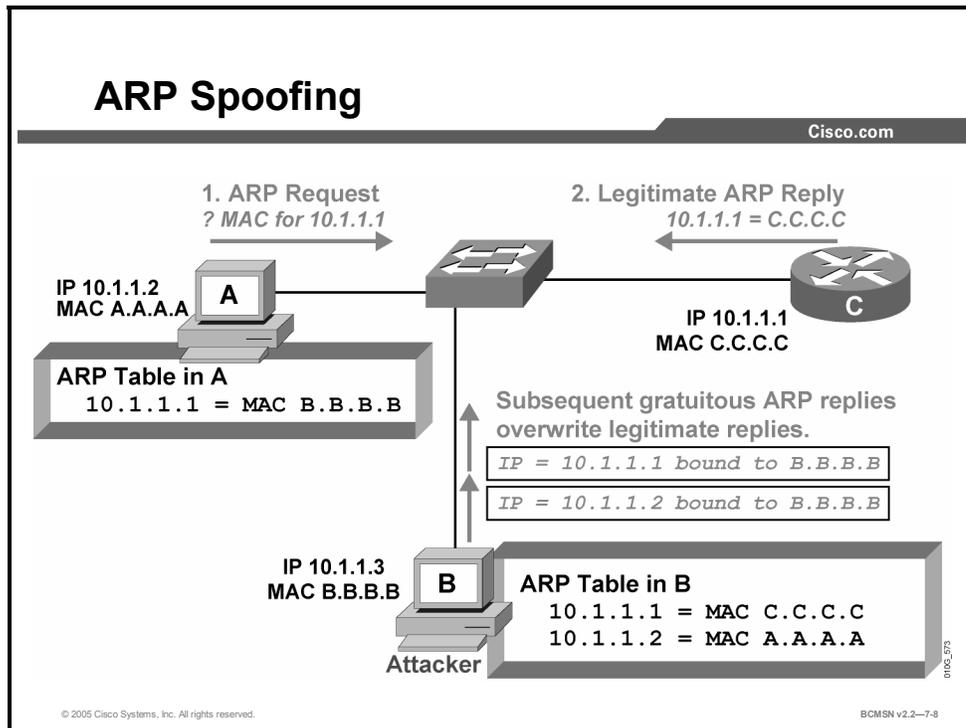
A MAC spoofing attack follows the sequence shown in the figure.

MAC Spoof Attack

Sequence of Events	Description
1.	The switch has learned that host A is on port 1, host B is on port 2, and host C is on port 3.
2.	Host B sends out a packet identifying itself as host B's IP address but with host A's MAC address or another packet with the same IP address and MAC address combination.
3.	This traffic causes the switch to move the location of host A in its CAM table from port 1 to port 2. Traffic from host C destined for host A is now visible to host B.
4.	To correct this situation, host A must send out traffic on the switch port so the switch can "relearn" the location of the MAC address for host A.

Describing ARP Spoofing

This topic describes ARP spoofing.



In normal ARP operation, a host sends a broadcast to determine the MAC address of a host with a particular IP address. The device at that IP address replies with its MAC address. The originating host caches the ARP response, using it to populate the destination Layer 2 header of packets sent to that IP address. By spoofing an ARP reply from a legitimate device with a gratuitous ARP, an attacking device appears to be the destination host sought by the senders. The ARP reply from the attacker causes the sender to store the attacking system's MAC address in its ARP cache. All packets destined for those IP addresses will be forwarded through the attacker system.

As illustrated in the figure, this is the sequence of events in an ARP spoofing attack.

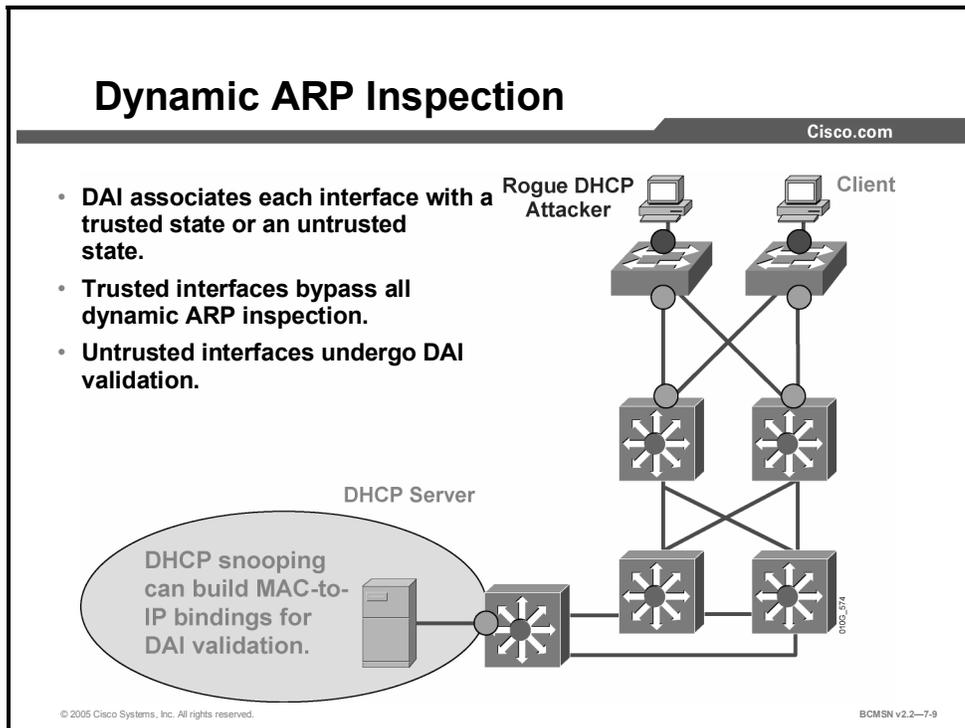
ARP Spoof Attack

Step or Sequence Number	Description
1.	Host A sends an ARP request for C's MAC address.
2.	Router C replies with its MAC and IP addresses. C also updates its ARP cache.
3.	Host A binds C's MAC address to its IP address in its ARP cache.
4.	Host B sends ARP binding B's MAC address to C's IP address.
5.	Host A updates ARP cache with B's MAC address bound to C's IP address.
6.	Host B sends ARP binding B's MAC address to A's IP address.
7.	Router C updates ARP cache with B's MAC address bound to A's IP address.

Step or Sequence Number	Description
8.	Packets are now diverted through attacker (B).

What Is Dynamic ARP Inspection?

This describes dynamic ARP inspection.



To prevent ARP spoofing or “poisoning,” a switch must ensure that only valid ARP requests and responses are relayed. Dynamic ARP inspection (DAI) prevents these attacks by intercepting and validating all ARP requests and responses. Each intercepted ARP reply is verified for valid MAC address-to-IP address bindings before it is forwarded to a PC to update the ARP cache. ARP replies coming from invalid devices are dropped.

DAI validates ARP replies coming from statically configured IP addresses or for a set of MAC addresses defined as in a VLAN access control list. DAI can also determine the validity of an ARP reply based on bindings stored in a DHCP snooping database. To ensure that only valid ARP requests and responses are relayed, DAI takes the following actions:

- Forwards ARP packets received on a trusted interface without any checks
- Intercepts all ARP packets on untrusted ports
- Verifies that each intercepted packet has a valid IP-to-MAC address binding before forwarding packets that can update the local ARP cache
- Drops, logs, or drops and logs ARP packets with invalid IP-to-MAC address bindings

Configure all access switch ports as untrusted and all switch ports connected to other switches as trusted. In this case, all ARP packets entering the network would be from an upstream distribution or core switch, bypassing the security check and requiring no further validation.

References

For additional information, refer to this resource:

Cisco Systems, Inc., *Configuring DAI (4500)*:

http://cisco.com/en/US/partner/products/hw/switches/ps4324/products_configuration_guide_chapter09186a008019d0ca.html

How to Configure Dynamic ARP Inspection

This topic describes the commands used to configure dynamic ARP inspection.

Configuring Dynamic ARP Inspection

Cisco.com

```
Switch(config)#ip arp inspection vlan vlan_id [,vlan_id]
```

- Enables DAI on a VLAN or range of VLANs

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

- Enables DAI on an interface and sets the interface as a trusted interface

```
Switch(config-if)#ip arp inspection validate {[src-mac] [dst-mac] [ip]}
```

- Configures DAI to drop ARP packets when the IP addresses are invalid

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—7-10

Here are the commands used to configure DAI.

Dynamic ARP Inspection Commands

Command	Description
Switch(config) # <code>ip arp inspection vlan <i>vlan_id</i> [,<i>vlan_id</i>]</code>	Enables DAI on a VLAN or range of VLANs
Switch(config-if) # <code>ip arp inspection trust</code>	Enables DAI on an interface and sets the interface as a trusted interface
Switch(config) # <code>ip arp inspection validate {<i>[src-mac]</i> [<i>dst-mac</i>] [<i>ip</i>]}</code>	Configures DAI to drop ARP packets when the IP addresses are invalid, or when the MAC addresses in the body of the ARP packets do not match the addresses specified in the Ethernet header

It is generally advisable to configure all access switch ports as untrusted and to configure all uplink ports connected to other switches as trusted. The following example of DAI implementation illustrates the configuration required on switch 2 with port FastEthernet 3/3 as the uplink port toward the DHCP server.

Example: DAI Implementation

This example shows how to configure dynamic ARP inspection for hosts on VLAN1, where client devices are located for switch 2. All client ports are untrusted by default. Only port 3/3 is trusted, as this is the only port where DHCP replies would be expected.

```
Switch S2(config)#ip arp inspection vlan 1
Switch S2(config)#interface fastethernet 3/3
Switch S2(config-if)#ip arp inspection trust
```

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **DHCP spoof attacks send unauthorized replies to DHCP queries.**
- **DHCP snooping is used to counter a DHCP spoof attack.**
- **MAC spoof attacks provide an unauthorized device access to frames intended for a valid network host.**
- **Port security is used to counter MAC spoof attacks.**
- **ARP spoofing can be used to redirect traffic to an unauthorized device on the network.**
- **Dynamic ARP inspection in conjunction with DHCP snooping can be used to counter ARP spoofing attacks.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—7-11

References

For additional information, refer to these resources:

- Cisco Systems, Inc., *Catalyst 4500 Series Switch, Understanding and Configuring DHCP Snooping*:
http://cisco.com/en/US/partner/products/hw/switches/ps663/products_configuration_guide_chapter09186a00800dde9f.html#30724
- Cisco Systems, Inc., *Catalyst 4500 Series Switch, Configuring Dynamic ARP*:
http://cisco.com/en/US/partner/products/hw/switches/ps4324/products_configuration_guide_chapter09186a008019d0ca.html

Implementing AAA

Overview

In networks in which there are a number of infrastructure devices, it is sometimes desirable to have all devices use a centralized server for authentication, authorization, and accounting (AAA) services for all the infrastructure devices. Cisco's asynchronous communications (ACS) provides both RADIUS and Terminal Access Controller Access Control System (TACACS+) services to provide this functionality.

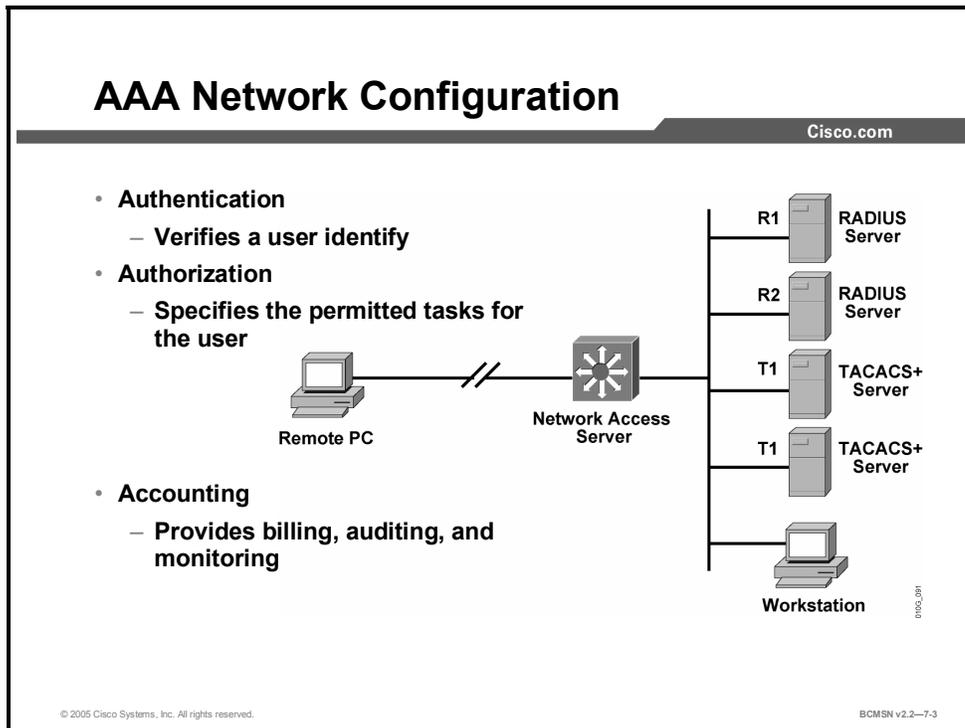
Objectives

Upon completing this lesson, you will be able to implement AAA with appropriate authentication, authorization, and accounting methods. This ability includes being able to meet these objectives:

- Describe AAA services
- Describe how AAA provides network security services
- Identify the AAA authentication and authorization methods
- Configure AAA features on a Catalyst switch
- Describe port-based authentication using 802.1X

Authentication, Authorization, and Accounting

This topic describes security in a multilayer switched network.



Authentication, authorization, and accounting (AAA) network security services provide the primary framework through which access control is set up on a switch. AAA is an architectural framework for configuring a set of three independent security functions in a consistent manner. AAA provides a modular way of performing these services.

- **Authentication:** Provides the method for identifying users, including login and password dialog, challenge and response, messaging support, and, depending on the security protocol, encryption.

Authentication is the way a user is identified before being allowed access to the network and network services. AAA authentication is configured by defining a list of named authentication methods and then applying that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they will be performed; it must be applied to a specific interface before any of the defined authentication methods will be performed. The only exception is the default method list (which is named “default”). The default method list is automatically applied to all interfaces if no other method list is defined. A defined method list overrides the default method list.

All authentication methods must be defined through AAA, with the exception of local, line password, and enable authentication.

- **Authorization:** Provides the method for remote access control, including one-time authorization or authorization for each service, per-user account list and profile, user group support, and support of IP, Internetwork Packet Exchange (IPX), AppleTalk Remote Access (ARA), and Telnet.

AAA authorization works by assembling a set of attributes that describes what the user is authorized to perform, such as access to different parts of the network. These attributes are compared to the information contained in a database for a given user, and the result is returned to AAA to determine the actual capabilities and restrictions of the user. The database can be located on the multilayer switch, or it can be hosted remotely on a RADIUS or TACACS+ security server. Remote security servers, such as RADIUS and TACACS+, authorize users for specific rights through attribute-value pairs, which associate those rights with the appropriate user. All authorization methods must be defined through AAA.

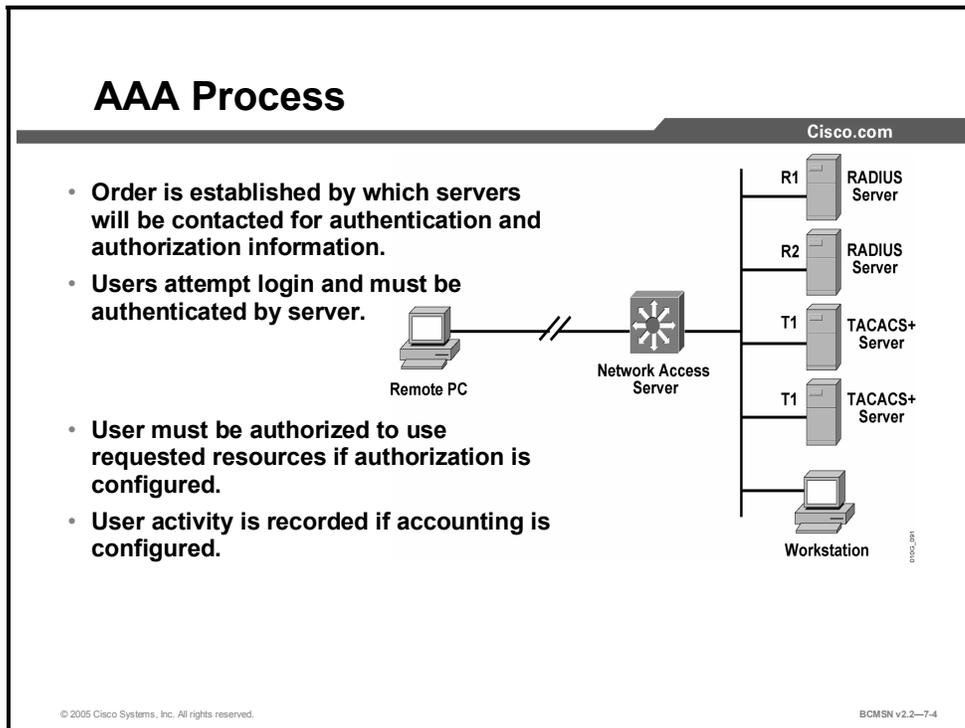
As with authentication, configure AAA authorization by defining a list of authorization methods and then applying that list to various interfaces.

- **Accounting:** Provides a method for collecting and sending security server information used for billing, auditing, and reporting. This information may include user identities, start and stop times, executed commands (such as PPP), number of packets, and number of bytes. Security experts can use the information gained from accounting to audit and improve security.

In many circumstances, AAA uses protocols such as RADIUS, TACACS+, or 802.1X to administer its security functions. If the switch is acting as a network access server, AAA is the means through which a switch establishes communication between the network access server and the RADIUS, TACACS+, or 802.1X security server.

Describing the AAA Process

This topic describes how the AAA process provides network security services.



AAA enables dynamic configuration of the type of authentication and authorization on a per-line (per-user) or per-service—for example, IP, IPX, or virtual private dial-up network (VPDN)—basis. Define the type of authentication and authorization by creating method lists and then applying those method lists to specific services or interfaces.

A method list is a sequential list of the methods for authenticating a user. Method lists enable you to designate one or more security protocols to be used for authentication, thus ensuring a backup system in case the initial method fails. Cisco IOS software uses the first method listed to authenticate users; if that method does not respond, the software selects the next authentication method in the list. This process continues until there is successful communication with a listed authentication method or until the authentication method list is exhausted, in which case authentication fails.

Note Cisco IOS software attempts authentication with the next listed authentication method only when there is no response from the previous method. If any device denies authentication, the authentication process stops; no other authentication methods are attempted.

First, decide what kind of security solution should be implemented. Assess the security risks in the particular network and decide on the appropriate means to prevent unauthorized entry and attack.

This table illustrates the AAA process based on the graphic.

AAA Authorization Process

Step	Description
1.	The system administrator has defined a method list in which R1 will be contacted first for authentication information, then R2, T1, T2, and finally the local username database on the access server itself. For authorization and accounting, R1 and R2 make up the group of RADIUS servers. T1 and T2 make up the group of TACACS+ servers.
2.	A remote user attempts to dial in to the network. The network access server first queries R1 for authentication information.
3.	If R1 authenticates the user, it issues a PASS response to the network access server, and the user is allowed to access the network. If R1 returns a FAIL response, the user is denied access and the session is terminated.
4.	If R1 does not respond, then the network access server processes that as an ERROR and queries R2 for authentication information. This pattern continues through the remaining designated methods until the user is either authenticated or rejected, or until the session is terminated. If all of the authentication methods return errors, the network access server will process the session as a failure, and the session will be terminated.
5.	Once authenticated, the remote user attempts to access a network resource such as an e-mail server or a database server.
6.	If authorization is configured, the network access server requests authorization from the appropriate authorization resource, per the AAA authorization method lists configured.
7.	Method lists for authorization define the ways that authorization will be performed and the sequence in which these methods will be performed. Cisco IOS software uses the first method listed to authorize users for specific network services; if that method fails to respond, the software selects the next method listed in the list. This process continues until there is successful communication with a listed authorization method or until all methods defined are exhausted.
8.	If accounting is configured, the network access server reports user activity to the appropriate accounting resource, per the AAA accounting method lists configured.
9.	Like authentication and authorization method lists, method lists for accounting define the way accounting will be performed and the sequence in which these methods will be performed. Cisco IOS software uses the first method listed to support accounting; if that method fails to respond, the software selects the next accounting method in the list. This process continues until there is successful communication with a listed accounting method or until all methods defined are exhausted.

Authentication and Authorization Methods

This topic describes the methods that can be used for authentication and authorization in AAA.

Authentication Methods

Cisco.com

```
Switch(config)#aaa authentication login {default |  
list-name} method1 [method2...]
```

- Creates a local authentication list

Cisco IOS AAA supports these authentication methods:

- enable password
- Kerberos 5
- Kerberos 5-Telnet authentication
- line password
- local database
- local database with case sensitivity
- no authentication
- RADIUS
- TACACS+

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-7-5

The AAA security services facilitate a variety of login authentication methods. Use the **aaa authentication login** command to enable AAA authentication no matter which of the supported login authentication methods are in use. With the **aaa authentication login** command, it is possible to create one or more lists of authentication methods that are tried at login. These lists are applied using the **login authentication** line configuration command.

The *list-name* is a character string used to name the list being created. The method argument refers to the actual method the authentication algorithm tries. The additional methods of authentication are used only if the previous method returns an error, not if it fails.

For example, to specify RADIUS as the default method for user authentication during login, enter the following command:

```
aaa authentication login default group radius
```

Cisco IOS AAA supports these authentication methods: enable, krb5, krb5-telnet, line, local, local-case, none, group radius, group tacacs+.

Specific commands for configuring authentication are covered later in this lesson.

Authorization Methods

This subtopic describes the methods that can be used for remote access control.

Authorization Methods

Cisco.com

```
Switch(config)#  
aaa authorization commands level1[method1 [method2...]]
```

- **Creates an authorization method list and enables authorization**

Cisco IOS AAA supports five different methods of authorization:

- TACACS+
- RADIUS
- If-Authenticated
- Local database
- None

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-7-6

AAA authorization enables limitation of the services available to a user. When AAA authorization is enabled, the network access server uses information retrieved from the profile of the user (which is located either in the local user database or on the security server) to configure the user session. Once this is done, the user will be granted access to a requested service only if the information in the user profile allows it.

Cisco IOS AAA supports five different methods of authorization: TACACS+, RADIUS, If-Authenticated, Local, and None.

The commands and sequence used to configure authorization are covered later in the lesson.

Configuring AAA

This topic describes how to configure AAA.

Configuring AAA

Cisco.com

1. **Enable AAA new model.**
2. **Consider external security servers.**
3. **Define methods for authentication.**
4. **Apply methods to interface or line.**
5. **Configure authorization as needed.**
6. **Configure accounting as needed.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-7.7

Configuring AAA is relatively simple once the basic process is understood.

Basic Process for Configuring AAA

Step	Description
1.	Enable AAA by using the aaa new-model global configuration command.
2.	If a separate security server is used, configure security protocol parameters, such as RADIUS, TACACS+, or Kerberos.
3.	Define the method lists for authentication by using an AAA authentication command.
4.	Apply the method lists to a particular interface or line, if required.
5.	(Optional) Configure authorization using the AAA authentication command.
6.	(Optional) Configure accounting using the AAA accounting command.

Configuring Authentication

This subtopic describes the process for configuring AAA authentication.

Configuring Authentication

Cisco.com

```
Switch(config)#aaa new-model
```

- Enables AAA globally

```
Switch(config)#aaa authentication login {default | list-name} method1 [method2...]
```

- Creates a local authentication list

```
Switch(config)#line [aux | console | tty | vty] line-number [ending-line-number]
```

- Enters line configuration mode

```
Switch(config-line)#login authentication {default / list-name}
```

- Applies the authentication list to a line

© 2005 Cisco Systems, Inc. All rights reserved.
BCMSN v2.2-7-8

The AAA security services facilitate a variety of login authentication methods. Use the **aaa authentication login** command to enable AAA authentication. With the **aaa authentication login** command, create one or more lists of authentication methods that are tried at login. These lists are applied using the **login authentication** line configuration command.

To configure login authentication by using AAA, use the commands in this table, beginning in global configuration mode.

Step	Description
1.	Enable AAA globally. <pre>Switch(config)#aaa new-model</pre>
2.	Create a local authentication list using appropriate authentication method options. <pre>Switch(config)#aaa authentication login {default list-name} method1 [method2...]</pre>
3.	Enter line configuration mode for the lines to which the authentication list will be applied. <pre>Switch(config)#line [aux console tty vty] line-number [ending-line-number]</pre>
4.	Apply the authentication list to a line or set of lines. <pre>Switch(config-line)#login authentication {default list-name}</pre>

Authentication Methods

Keyword	Description
enable	Uses the enable password for authentication.
krb5	Uses Kerberos 5 for authentication.
krb5-telnet	Uses Kerberos 5-Telnet authentication protocol when using Telnet to connect to the router. If selected, this keyword must be listed as the first method in the method list.
line	Uses the line password for authentication.
local	Uses the local username database for authentication.
local-case	Uses case-sensitive local username authentication.
none	Uses no authentication.
group radius	Uses the list of all RADIUS servers for authentication.
group tacacs+	Uses the list of all TACACS+ servers for authentication.
group <i>group-name</i>	Uses a subset of RADIUS or TACACS+ servers for authentication as defined by the aaa group server radius or aaa group server tacacs+ commands.

To specify that the authentication should succeed even if all methods return an error, specify *none* as the final method in the command line. For example, to specify that authentication should succeed even if, as in this example, the TACACS+ server returns an error, enter the following command:

```
aaa authentication login default group tacacs+ none
```

Note Because the **none** keyword enables *any* user logging in to successfully authenticate, it should be used only as a backup method of authentication.

To create a default list that is used when a named list is not specified in the **login authentication** command, use the **default** keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces.

Example: Configuring Authentication

The following example creates an authentication list called “myway” that uses TACACS+ as the first authentication method and local authentication as the second. The authentication list is then applied to a line.

```
Switch(config)#aaa authentication login myway tacacs+ local
Switch(config)#line con 0
Switch(config-line)#login authentication myway
Authentication
```

Configuring Authorization

This subtopic describes the process for configuring AAA authorization.

Configuring Authorization

Cisco.com

```
Switch(config)#aaa authorization {auth-proxy | network |  
exec | commands level | reverse-access | configuration |  
ipmobile} {default | list-name} [method1 [method2...]]
```

- **Creates an authorization method list and enables authorization**

```
Switch(config)#interface interface-type interface-number
```

- **Enters interface configuration mode**

```
Switch(config-if)#ppp authorization {default | list-name}
```

- **Applies the named authorization method list to the interface**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-7-9

AAA authorization enables limitation of services available to a user. When AAA authorization is enabled, the multilayer switch uses information retrieved from the user profile (which is located either in the local user database on the switch or on the security server) to configure the user session. When this task is done, the user will be granted access to a requested service only if the information in the user profile allows it.

Just as with AAA authentication, authorization creates method lists to define the ways that authorization will be performed and the sequence in which these methods will be performed. Method lists are specific to the authorization type requested.

- **Auth-proxy:** Applies specific security policies on a per-user basis.
- **Commands:** Applies to the EXEC mode commands that a user issues. Command authorization attempts authorization for all EXEC mode commands, including global configuration commands, associated with a specific privilege level.
- **EXEC:** Applies to the attributes associated with a user EXEC terminal session.
- **Network:** Applies to network connections. These connections can include a PPP, Serial Line Internet Protocol (SLIP), or AppleTalk Remote Access (ARA) Protocol connection.
- **Reverse access:** Applies to reverse Telnet sessions.

When creating a named method list, define a particular list of authorization methods for the indicated authorization type.

AAA supports five different methods of authorization.

- **TACACS+:** The network access server exchanges authorization information with the TACACS+ security daemon. TACACS+ authorization defines specific rights for users by

associating attribute-value pairs, which are stored in a database on the TACACS+ security server, with the appropriate user.

- **If-Authenticated:** The user is allowed to access the requested function, provided that the user has been authenticated successfully.
- **None:** The network access server does not request authorization information; authorization is not performed over this line or interface.
- **Local:** The router or access server consults its local database—as defined by the **username** command, for example—to authorize specific rights for users. Only a limited set of functions can be controlled via the local database.
- **RADIUS:** The network access server requests authorization information from a RADIUS security server. RADIUS authorization defines specific rights for users by associating attributes.

To configure AAA authorization using named method lists, use these commands, beginning in global configuration mode.

Step	Description
1.	Enables AAA globally (if not previously enabled) Switch(config)# aaa new-model
2.	Creates an authorization method list for a particular authorization type and enables authorization Switch(config)# aaa authorization {auth-proxy network exec commands level reverse-access configuration ipmobile} {default list-name} [method1 [method2...]]
3.	Enters line configuration mode for the lines or interfaces to which the authorization list should be applied Switch(config)# line [aux console tty vty] line-number [ending-line-number] OR Switch(config)# interface interface-type interface-number
4.	Applies the authorization list to a line or set of lines, or to an interface Switch(config-line)# authorization {arap commands level exec reverse-access} {default list-name} OR Switch(config-if)# ppp authorization {default list-name}

To have the multilayer switch request authorization information via a TACACS+ security server, use the **aaa authorization** command with the *group tacacs+* value for the *method* variable.

To allow users access to the functions that they request as long as they have been authenticated, use the **aaa authorization** command with the **if-authenticated** method keyword. If this method is selected, all requested functions are automatically granted to authenticated users.

To select local authorization, which means that the router or access server consults its local user database to determine the functions that are permitted to a user, employ the **aaa authorization** command with the **local** method keyword. The functions associated with local authorization are defined by using the **username** global configuration command.

To have the network access server request authorization via a RADIUS security server, use the **radius** method keyword.

Authorization Methods

Method	Description
TACACS+	The network access server exchanges authorization information with the TACACS+ security daemon. TACACS+ authorization defines specific rights for users by associating attribute-value pairs, which are stored in a database on the TACACS+ security server, with the appropriate user.
If-Authenticated	The user is allowed to access the requested function, provided that the user has been authenticated successfully.
None	The network access server does not request authorization information; authorization is not performed over this line or interface.
Local	The router or access server consults its local database—as defined by the username command, for example—to authorize specific rights for users. Only a limited set of functions can be controlled via the local database.
RADIUS	The network access server requests authorization information from a RADIUS security server. RADIUS authorization defines specific rights for users by associating attributes.

To have the multilayer switches request authorization information via a TACACS+ security server, use the **aaa authorization** command with the *group tacacs+* value for the *method* variable.

To allow users access to the functions that they request as long as they have been authenticated, use the **aaa authorization** command with the **if-authenticated** method keyword. If this method is selected, all requested functions are automatically granted to authenticated users.

To select local authorization, which means that the router or access server consults its local user database to determine the functions that a user is permitted to use, use the **aaa authorization** command with the **local** method keyword. The functions associated with local authorization are defined by using the **username** global configuration command.

To have the network access server request authorization via a RADIUS security server, use the **radius** method keyword.

Configuring Accounting

This subtopic describes how to configure the accounting feature of AAA.

Configuring Accounting

Cisco.com

```
Switch(config)#aaa accounting {system | network | exec |  
connection | commands level} {default | list-name} {start-  
stop | stop-only | none} [method1 [method2...]]
```

- **Creates an accounting method list and enables accounting**

```
Switch(config)#interface interface-type interface-number
```

- **Enters interface configuration mode**

```
Switch(config-if)#ppp accounting {default | list-name}
```

- **Applies the named accounting method list to the interface**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—7-10

Accounting is the process of keeping track of the activity of each user who is accessing the network resources; this includes the amount of time spent in the network, the services accessed while there, and the amount of data transferred during the session. Accounting data is used for trend analysis, capacity planning, billing, auditing, and cost allocation.

AAA supports six different accounting types.

- **Network accounting:** Provides information for all PPP, SLIP, or AppleTalk Remote Access Protocol (ARAP) sessions, including packet and byte counts
- **Connection accounting:** Provides information about all outbound connections made from the network, such as Telnet and remote login (rlogin)
- **EXEC accounting:** Provides information about user EXEC terminal sessions (user shells) on the network access server, including username, date, start and stop times, the access server IP address, and (for dial-in users) the telephone number the call originated from
- **System accounting:** Provides information about all system-level events (for example, when the system reboots or when accounting is turned on or off)
- **Command accounting:** Provides information about the EXEC shell commands for a specified privilege level that are being executed on a network access server
- **Resource accounting:** Provides start and stop record support for calls that have passed user authentication

To configure AAA accounting using named method lists, use the commands in this table, beginning in global configuration mode.

Step	Description
1.	<p>Create an accounting method list and enable accounting.</p> <pre>Switch(config)#aaa accounting {system network exec connection commands level} {default list-name} {start-stop stop-only none} [method1 [method2...]]</pre>
2.	<p>Enter the line configuration mode or the interface to which the accounting method list should be applied.</p> <pre>Switch(config)#line [aux console tty vty] line-number [ending- line-number]</pre> <p>or</p> <pre>Switch(config)#interface interface-type interface-number</pre>
3.	<p>Apply the accounting method list to a line or interface.</p> <pre>Switch(config-line)#accounting {arap commands level connection exec} {default list-name}</pre> <p>or</p> <pre>Switch(config-if)#ppp accounting {default list-name}</pre>

Accounting Methods

Method	Description
group radius	Uses the list of all RADIUS servers for accounting
group tacacs+	Uses the list of all TACACS+ servers for accounting
group <i>group-name</i>	Uses a subset of RADIUS or TACACS+ servers for accounting, as defined by the aaa group server radius or aaa group server tacacs+ commands

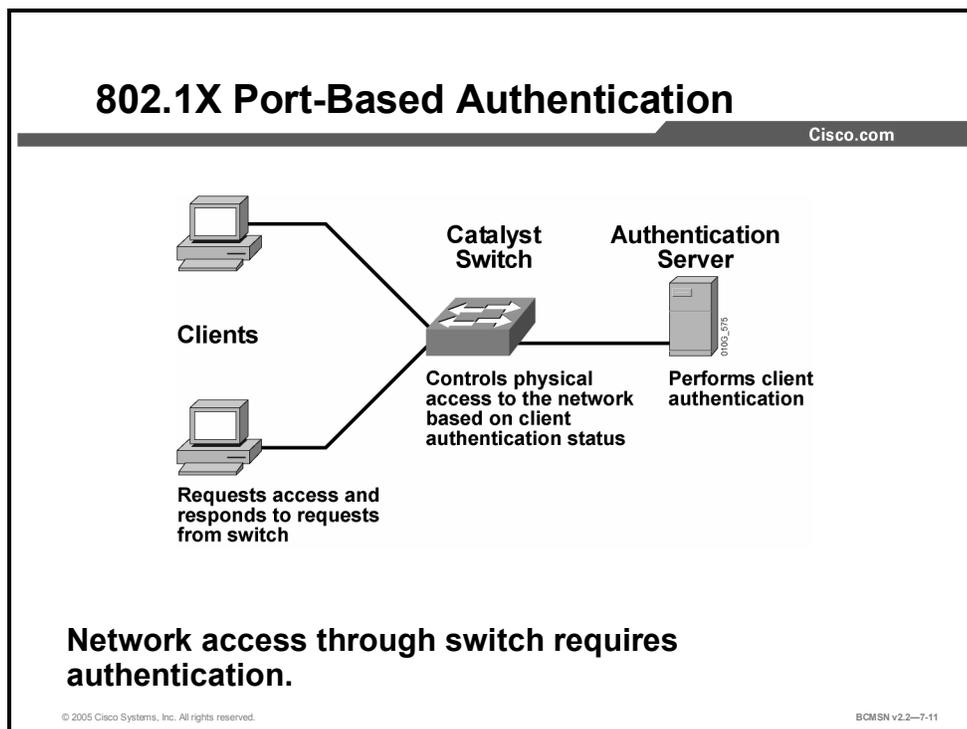
Comprehensive AAA Configuration Example

The following example shows how to configure a Cisco access device for AAA services to be provided by the RADIUS server for an access server with dialup links. If the RADIUS server fails to respond, then the local database will be queried for authentication and authorization information, and accounting services will be handled by a TACACS+ server.

```
aaa new-model
aaa authentication login admins local
aaa authentication ppp dialins group radius local
aaa authorization network scoobee group radius local
aaa accounting network charley start-stop group radius group tacacs+
!
username root password ALongPassword
!
tacacs-server host 172.31.255.0
tacacs-server key goaway
!
radius-server host 172.16.2.7
radius-server key myRaDiUSpassWoRd
!
interface group-async 1
  group-range 1 16
  encapsulation ppp
  ppp authentication chap dialins
  ppp authorization scoobee
  ppp accounting charley
!
line 1 16
  autoselect ppp
  autoselect during-login
  login authentication admins
  modem dialin
```

802.1X Port-Based Authentication

This topic describes 802.1X port-based authentication.



The IEEE 802.1X standard defines a port-based access control and authentication protocol that restricts unauthorized workstations from connecting to a LAN through publicly accessible ports. The authentication server authenticates each workstation connected to a switch port before making available any services offered by the switch or the LAN.

Until the workstation is authenticated, 802.1X access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the workstation is connected. After authentication succeeds, normal traffic can pass through the port.

With 802.1X port-based authentication, the devices in the network have specific roles, as follows:

- **Client:** The device (workstation) that requests access to the LAN and switch services, and responds to requests from the switch. The workstation must be running 802.1X-compliant client software, such as that offered in the Microsoft Windows XP operating system. (The port that the client is attached to is the supplicant [client] in the IEEE 802.1X specification.)
- **Authentication server:** Performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. The RADIUS security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server.
- **Switch (also called the authenticator):** Controls physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client (supplicant) and the authentication server, requesting identifying information

from the client, verifying that information with the authentication server, and relaying a response to the client. The switch uses a RADIUS software agent, which is responsible for encapsulating and decapsulating the EAP frames and interacting with the authentication server.

The switch port state determines whether or not the client is granted access to the network. The port starts in the unauthorized state. While in this state, the port disallows all ingress and egress traffic except for 802.1X protocol packets. When a client is successfully authenticated, the port transitions to the authorized state, allowing all traffic for the client to flow normally.

If the switch requests the client identity (authenticator initiation) and the client does not support 802.1X, the port remains in the unauthorized state, and the client is not granted access to the network.

In contrast, when an 802.1X-enabled client connects to a port and the client initiates the authentication process (supplicant initiation) by sending the EAPOL-start frame to a switch not running the 802.1X protocol, no response is received, and the client begins sending frames as if the port is in the authorized state.

You control the port authorization state by using the **dot1x port-control** interface configuration command and these keywords:

- **force-authorized:** Disables 802.1X port-based authentication and causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1X-based authentication of the client. This is the default setting.
- **force-unauthorized:** Causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
- **auto:** Enables 802.1X port-based authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port transitions from down to up (authenticator initiation) or when an EAPOL-start frame is received (supplicant initiation). The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. The switch uniquely identifies each client attempting to access the network by using the client MAC address.

If the client is successfully authenticated (receives an “accept” frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails and network access is not granted.

When a client logs off, it sends an EAPOL-logoff message, causing the switch port to transition to the unauthorized state.

Configuring 802.1X Port-Based Authentication

This subtopic describes configuring 802.1X port-based authentication.

Configuring 802.1X

Cisco.com

```
Switch(config)#aaa authentication dot1x {default} method1 [method2...]
```

- **Creates an 802.1X port-based authentication method list**

```
Switch(config)#dot1x system-auth-control
```

- **Globally enables 802.1X port-based authentication**

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

- **Enters interface configuration mode**

```
Switch(config-if)#dot1x port-control auto
```

- **Enables 802.1X port-based authentication on the interface**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2--7-12

To implement 802.1X port-based authentication follow these steps:

Step	Description
1.	Enable AAA. <pre>Switch(config)#aaa new-model</pre>
2.	Create an 802.1X port-based authentication method list. <pre>Switch(config)#aaa authentication dot1x {default} method1 [method2...]</pre>
3.	Globally enable 802.1X port-based authentication. <pre>Switch(config)#dot1x system-auth-control</pre>
4.	Enter interface configuration mode and specify the interface to be enabled for 802.1X port-based authentication. <pre>Switch(config)#interface type slot/port</pre>
5.	Enable 802.1X port-based authentication on the interface. <pre>Switch(config-if)#dot1x port-control auto</pre>
6.	Return to privileged EXEC mode. <pre>Switch(config)#end</pre>

Example

The example shows how to enable AAA and 802.1X on Fast Ethernet port 5/1:

```
Switch#configure terminal
Switch(config)#aaa new-model
Switch(config)#aaa authentication dot1x default group radius
Switch(config)#dot1x system-auth-control
Switch(config)#interface fastethernet 5/1
Switch(config-if)#dot1x port-control auto
Switch(config-if)#end
```

Reference

For additional information, refer to this resource:

http://www.cisco.com/en/US/partner/products/hw/switches/ps628/products_configuration_guide_chapter09186a00800d84b9.html

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **AAA services provide additional security for Cisco's network infrastructure devices.**
- **Different authorization levels can be configured on the AAA server and implemented on the networking device, based on the network security policy.**
- **A number of different authentication methods can be configured on a device, to protect its lines and ports based on the network security policy.**
- **Both the Catalyst switch and the AAA server need to be configured properly to implement AAA services.**
- **802.1X can be used to authenticate users on a port-by-port basis using an authentication server.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—7-13

References

For additional information, refer to these resources:

- Cisco Systems, Inc., *Authentication Protocols, Configuring Basic AAA on and Access Server*.
http://www.cisco.com/en/US/tech/tk59/technologies_tech_note09186a0080093c81.shtml
- Cisco Systems, Inc., *Catalyst 2950 Desktop Switch Software Configuration Guide, Configuring 802.1X Port-Based Authentication*:
http://www.cisco.com/en/US/partner/products/hw/switches/ps628/products_configuration_guide_chapter09186a00800d84b9.html

Defending Network Switches

Overview

The devices on any network must be secured. A number of vulnerabilities can be reduced by setting passwords on physical and virtual ports, by disabling unneeded services, by forcing the encryption of sessions, and by enabling logging at the device level.

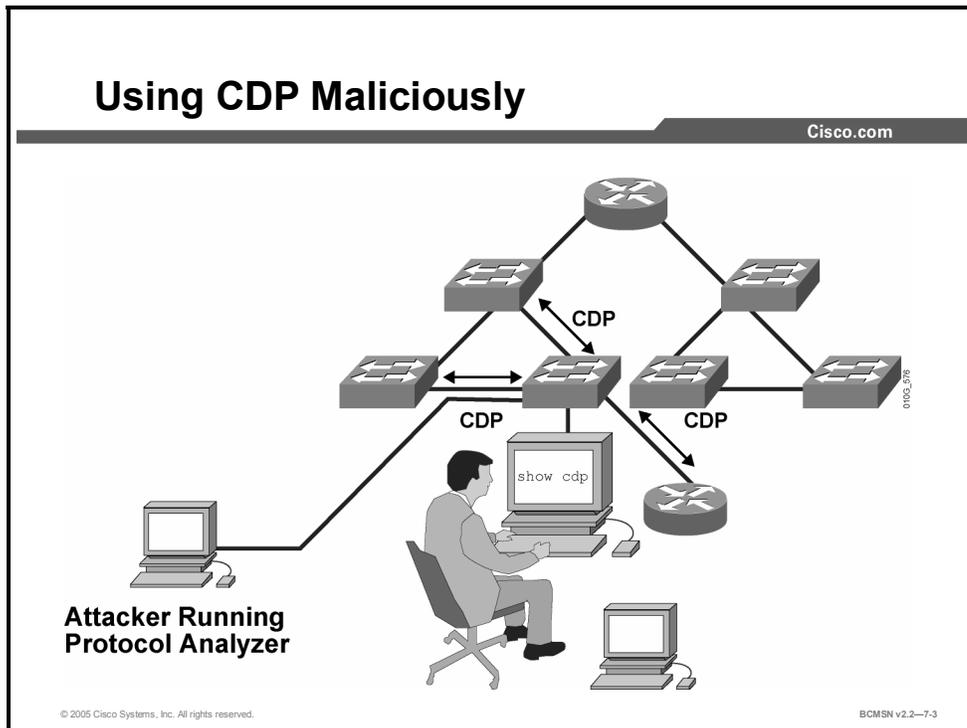
Objectives

Upon completing this lesson, you will be able to identify switch security risks and list best practices when placing new switches into service. This ability includes being able to meet these objectives:

- Describe how CDP can be used for reconnaissance
- Describe vulnerabilities inherent to the Telnet protocol
- Configure VTY ACLs to secure Telnet access to devices
- Describe benefits of the SSH Protocol
- List best practices for strengthening the security posture of a switched network
- List the configuration steps required to capture aggregate data traffic on a switched network

CDP Security Issues

This topic describes Cisco Discovery Protocol (CDP).



Attackers with knowledge of how CDP works could find ways to take advantage of the clear-text CDP packets to gain knowledge of the network. The CDP runs at Layer 2 and allows Cisco devices to identify themselves to other Cisco devices. However, the information sent through CDP is transmitted in clear text and is unauthenticated. Utilizing a packet analyzer, attackers could glean information about the network device from CDP advertisements.

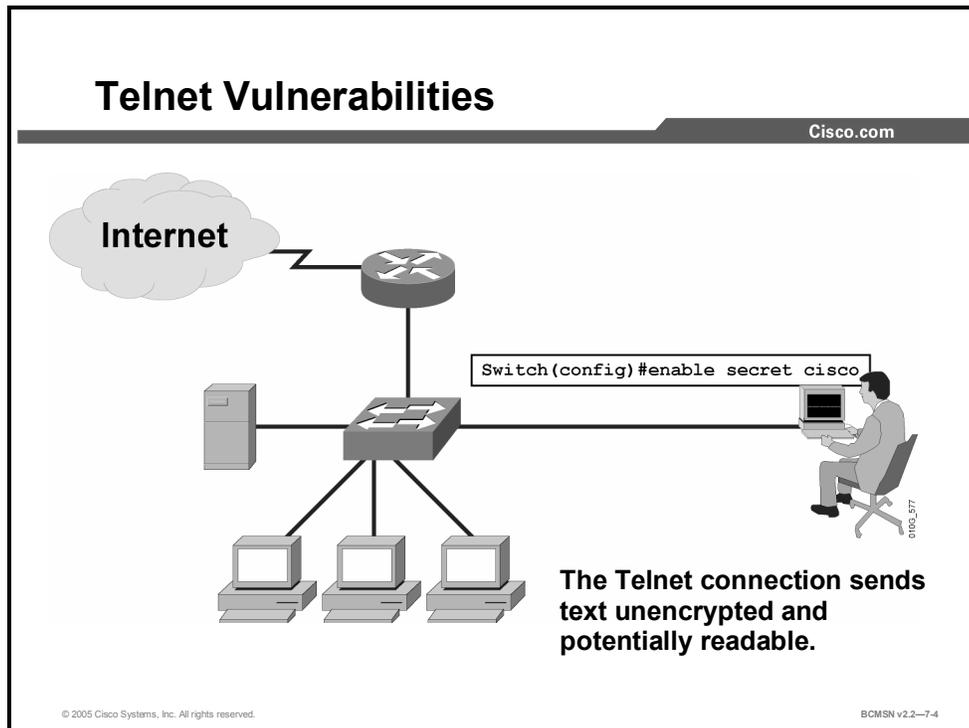
CDP is necessary for management applications and cannot be disabled without impairing some network-management applications. However, CDP can be selectively disabled on interfaces where management is not being performed.

How CDP Can Be Used Maliciously

Sequence of Events	Description
1.	System administrator uses CDP to view neighbor information.
2.	Attacker uses a packet analyzer to intercept CDP traffic.
3.	Attacker analyzes information in CDP packets to gain knowledge of network address and device information.
4.	Attacker formulates attacks based on known vulnerabilities of network platforms.

Vulnerabilities in Telnet

This topic describes vulnerabilities within the Telnet service on Cisco IOS devices.



Known Telnet vulnerabilities:

- All usernames, passwords, and data that are sent over the public network in clear text are vulnerable.
- A user with an account on the system could gain elevated privileges.
- A remote attacker could crash the Telnet service, preventing legitimate use of that service.
- A remote attacker could find an enabled guest account that may be present anywhere within the trusted domains of the server.

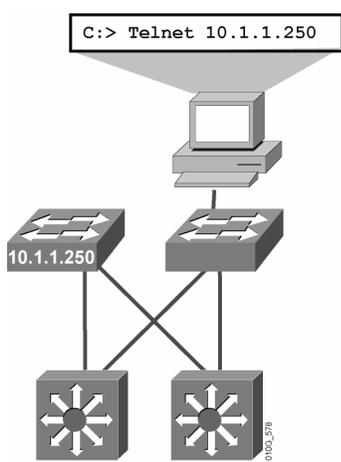
VTY ACLs

This topic describes why to apply access control lists (ACLs) on VTY ports.

Controlling VTY Access

Cisco.com

- **Set up standard IP ACL.**
- **Use line configuration mode to filter access with the access-class command.**
- **Set identical restrictions on every VTY line.**



© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-7-5

Cisco provides ACLs to permit or deny Telnet access to the VTY ports of a switch. Cisco devices vary in the number of VTY ports that are available by default. When configuring VTY ACLs, ensure that all default ports are removed or have a specific VTY ACL applied.

Telnet filtering is normally considered an extended IP ACL function because it is filtering a higher-level protocol. However, because the **access-class** command is used to filter incoming Telnet sessions by source address and to apply filtering to VTY lines, standard IP ACL statements can be used to control VTY access. The **access-class** command also applies standard IP ACL filtering to VTY lines for outgoing Telnet sessions originating from the switch.

VTY ACLs can be applied to any combination of VTY lines. The same ACL can be applied to all VTY lines, or separately to each VTY line. The most common practice is to apply the same ACL to all VTY lines.

Commands to Configure VTY ACLs

This subtopic lists commands used to configure VTY ACLs.

Configuring VTY ACLs

Cisco.com

```
Switch(config)#access-list access-list-number
{permit | deny | remark} source [mask]
```

- **Configure a standard IP access list**

```
Switch(config)#line vty {vty# | vty-range}
```

- **Enter configuration mode for a VTY or VTY range**

```
Switch(config-line)#access-class access-list-number in|out
```

- **Restrict incoming or outgoing VTY connections to addresses in the ACL**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2-7-6

To configure VTY ACLs on a Cisco switch, create a standard IP ACL and apply the ACL on the VTY interfaces. Rather than applying the ACL to a data interface, apply it to a VTY line or range of lines with the **access-class** command.

Example: VTY Access

In this example, permission is granted to any device on network 192.168.1.0 0.0.0.255 to establish a virtual terminal (Telnet) session with the switch. Of course, the user must know the appropriate passwords to enter user mode and privileged mode.

Notice that identical restrictions have been set on every VTY line because the line on which the VTY user will connect cannot be controlled.

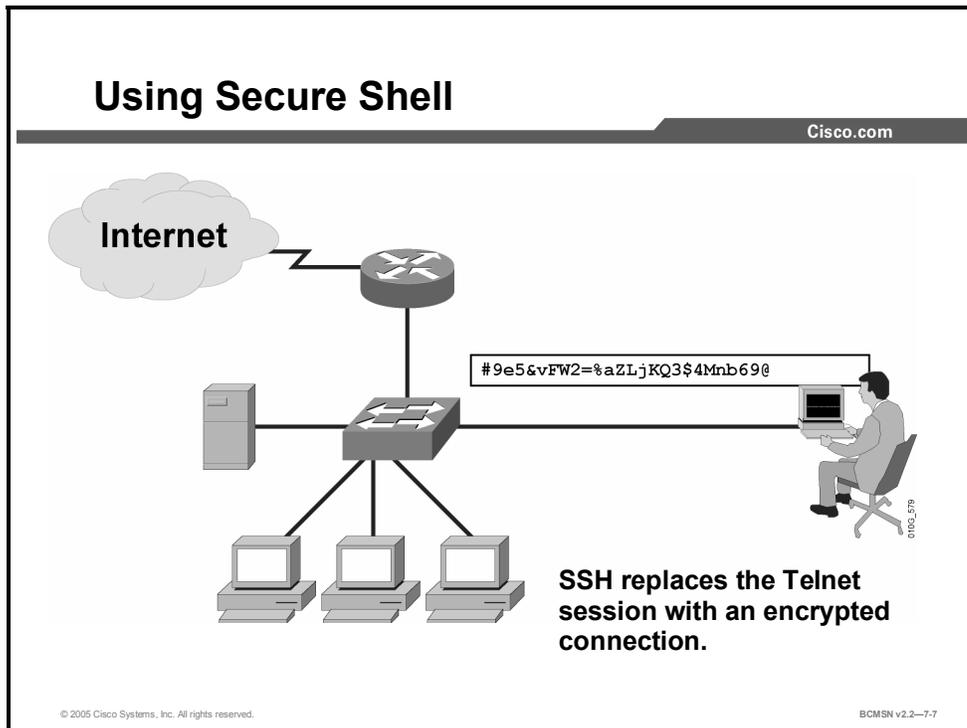
The implicit *deny any* statement at the end of the access list still applies to the ACL when it is used as an access-class entry.

```
Switch(config)# access-list 12 permit 192.168.1.0 0.0.0.255
Switch(config)# line vty 0 15
Switch (config-line)# access-class 12 in
```

Note The actual number of VTY lines depends on the platform and the IOS software being run.

Secure Shell Protocol

This topic describes security advantages of using the Secure Shell Protocol.



Secure Shell (SSH) is a program used to log in to another computer over a network, to execute commands in a remote machine, and to move files from one machine to another. It provides strong authentication and secure communications over insecure channels. It is a replacement for login, rsh, rcp, and rdist as well as Telnet.

SSH protects a network from attacks such as IP spoofing, IP source routing, and Domain Name System (DNS) spoofing. An attacker who has managed to take over a network can only force SSH to disconnect. The attacker cannot play back the traffic or hijack the connection when encryption is enabled.

When using the SSH login (instead of Telnet), the entire login session, including transmission of password, is encrypted; therefore, it is almost impossible for an outsider to collect passwords.

Caution SSH version 1 implementations are vulnerable to various security compromises. Whenever possible, use SSH version 2 instead of SSH version 1.

Best Practices: Switch Security Considerations

This topic describes general security considerations that should be applied in any switched network.

Best Practices: Switch Security

Cisco.com

Secure Switch Access:

- **Set system passwords.**
- **Secure physical access to the console.**
- **Secure access via Telnet.**
- **Use SSH when possible.**
- **Configure system warning banners.**
- **Disable unused services.**
- **Configure basic logging.**
- **Secure SNMP.**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—7-8

Network security vulnerabilities include loss of privacy, data theft, impersonation, and loss of integrity. Basic security measures should be taken on every network to mitigate adverse effects of user negligence or acts of malicious intent.

Best practices following these general steps are required whenever placing new equipment in service.

1. Consider or establish organizational security policies.
2. Secure switch devices.
3. Secure switch protocols.
4. Mitigate compromises launched through a switch.

Organizational Security Policies

It is important to consider the policies of an organization when determining what level of security and what type of security should be implemented. There is a need to balance the goal of reasonable network security against the administrative overhead that is clearly associated with extremely restrictive security measures.

A well-established security policy has these characteristics:

- Provides a process for auditing existing network security
- Provides a general security framework for implementing network security

- Defines disallowed behaviors toward electronic data
- Determines which tools and procedures are needed for the organization
- Communicates consensus among a group of key decision makers and defines responsibilities of users and administrators
- Defines a process for handling network security incidents
- Enables an enterprise-wide, all-site security implementation and enforcement plan

Secure Switch Devices

Follow these best practices for secure switch access.

- **Set system passwords:** Use the **enable secret** command to set the password that grants enabled access to the Cisco IOS system. Because the **enable secret** command simply implements a Message Digest 5 (MD5) hash on the configured password, that password still remains vulnerable to dictionary attacks. Therefore, apply standard practices in selecting a feasible password. Try to pick passwords that contain both letters and numbers as well as special characters, for example, \$pecial1\$ instead of “specials,” where the “s” has been replaced by “\$,” and the “l” has been replaced with “1”(one).
- **Secure access to the console:** Console access requires a minimum level of security both physically and logically. An individual who gains console access to a system will be able to recover or reset the system-enable password, thus allowing that person to bypass all other security implemented on that system. Consequently, it is imperative to secure access to the console.
- **Secure access to VTY lines:** These are the minimum recommended steps for securing Telnet access.
 - Apply the basic ACL for in-band access to all VTY lines.
 - Configure a line password for all configured VTY lines.
 - If the installed IOS image permits, use SSH instead of Telnet to access the device remotely.
- **Use SSH:** The SSH protocol and application provide a secure remote connection to a router. Two versions of SSH are available: SSH version 1 and SSH version 2. SSH version 1 is implemented in Cisco IOS software. It encrypts all traffic, including passwords, between a remote console and a network router across a Telnet session. Because SSH sends no traffic in clear text, network administrators can conduct remote access sessions that casual observers will not be able to view. The SSH server in IOS software will work with publicly and commercially available SSH clients.
- **Configure system-warning banners:** For both legal and administrative purposes, configuring a system-warning banner to display prior to login is a convenient and effective way of reinforcing security and general usage policies. By clearly stating the ownership, usage, access, and protection policies before a login, you provide more solid backing for potential future prosecution.
- **Disable unneeded services:** By default, Cisco devices implement multiple TCP and UDP servers to facilitate management and integration into existing environments. For most installations these services are typically not required, and disabling them can greatly reduce overall security exposure. These commands will disable the services not typically used:

```
no service tcp-small-servers
no service udp-small-servers
```

```
no service finger
no service config
```

- **Disable the integrated HTTP daemon if not in use:** Although Cisco IOS software provides an integrated HTTP server for management, it is highly recommended that it be disabled to minimize overall exposure. If HTTP access to the switch is absolutely required, use basic ACLs to permit access from only trusted subnets.
- **Configure basic logging:** To assist and simplify both problem troubleshooting and security investigations, monitor the switch subsystem information received from the logging facility. View the output in the on-system logging buffer memory. To render the on-system logging useful, increase the default buffer size.
- **Secure SNMP:** Whenever possible, avoid using Simple Network Management Protocol (SNMP) read-write features. SNMP v2c authentication consists of simple text strings communicated between devices in clear, unencrypted text. In most cases, a read-only community string may be configured. In doing so, apply the basic access list mask to allow SNMP traffic to trusted hosts only.

Secure Switch Protocols

This subtopic continues a discussion of best practices for switch security.

Best Practices: Switch Security (Cont.)

Cisco.com

Secure Switch Protocols

- **Trim CDP and use only as needed.**
- **Secure spanning tree.**

Mitigate Compromises Through a Switch

- **Proactively configure unused ports.**
- **Take precautions for trunk links.**
- **Minimize physical port access.**
- **Establish standard access port configuration for both unused and used ports.**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—7.9

Follow these best practices for switch security.

- **Cisco Discovery Protocol (CDP):** CDP does not reveal security-specific information, but it is possible for an attacker to exploit this information in a reconnaissance attack, whereby an attacker learns device and IP address information for the purpose of launching other types of attacks. Two practical guidelines should be followed for CDP.
 - If CDP is not required, or the device is located in an unsecure environment, disable CDP globally on the device.
 - If CDP is required, disable CDP on a per-interface basis on ports connected to untrusted networks. Because CDP is a link-level protocol, it is not transient across a network (unless a Layer 2 tunneling mechanism is in place). Limit it to run only between trusted devices and disable it everywhere else. However, CDP is required on any access port when you are attaching a Cisco phone to establish a trust relationship.
- **Secure the spanning tree topology:** It is important to protect the Spanning Tree Protocol (STP) process of the switches composing the infrastructure. Inadvertent or malicious introduction of STP bridge protocol data units (BPDUs) could potentially overwhelm a device or pose a denial of service (DoS) attack. The first step in stabilizing a spanning tree installation is to positively identify the intended root bridge in the design and to hard set the STP bridge priority of that bridge to an acceptable root value. Do the same for the designated backup root bridge. These actions will protect against inadvertent shifts in STP due to an uncontrolled introduction of a new switch.

On some platforms, the BPDU guard feature may be available. If so, enable it on access ports in conjunction with the PortFast feature to protect the network from unwanted BPDU traffic injection. Upon receipt of a BPDU, the feature will automatically disable the port.

Mitigating Compromises Launched Through a Switch

Follow these best practices to mitigate compromises through a switch.

- **Proactively configure unused router and switch ports.**
 - Execute the **shut** command on all unused ports and interfaces.
 - Place all unused ports in a "parking-lot" VLAN used specifically to group unused ports until they are proactively placed into service.
 - Configure all unused ports as access ports disallowing automatic trunk negotiation.
- **Considerations for trunk links:** By default, Catalyst switches running IOS software are configured to automatically negotiate trunking capabilities. This situation poses a serious hazard to the infrastructure because an unsecured third party device can be introduced to the network as a valid infrastructure component. Potential attacks include interception of traffic, redirection of traffic, DoS, and more. To avoid this risk, disable automatic negotiation of trunking and manually enable it on links that will require it. Ensure that trunks use a native VLAN dedicated exclusively to trunk links.
- **Physical device access:** Physical access to the switch should be closely monitored to avoid rogue device placement in wiring closets with direct access to switch ports.
- **Access port–based security:** Specific measures should be taken on every access port of any switch placed into service. Ensure that a policy is in place outlining the configuration of unused switch ports as well as those that are in use.

For ports enabled for end-device access, there is a macro called **switchport host**, which, when executed on a specific switch port, takes the following actions: sets the switch port mode to access, enables spanning tree PortFast, and disables channel grouping.

Note The **switchport host** macro disables EtherChannel, disables trunking, and enables STP PortFast.

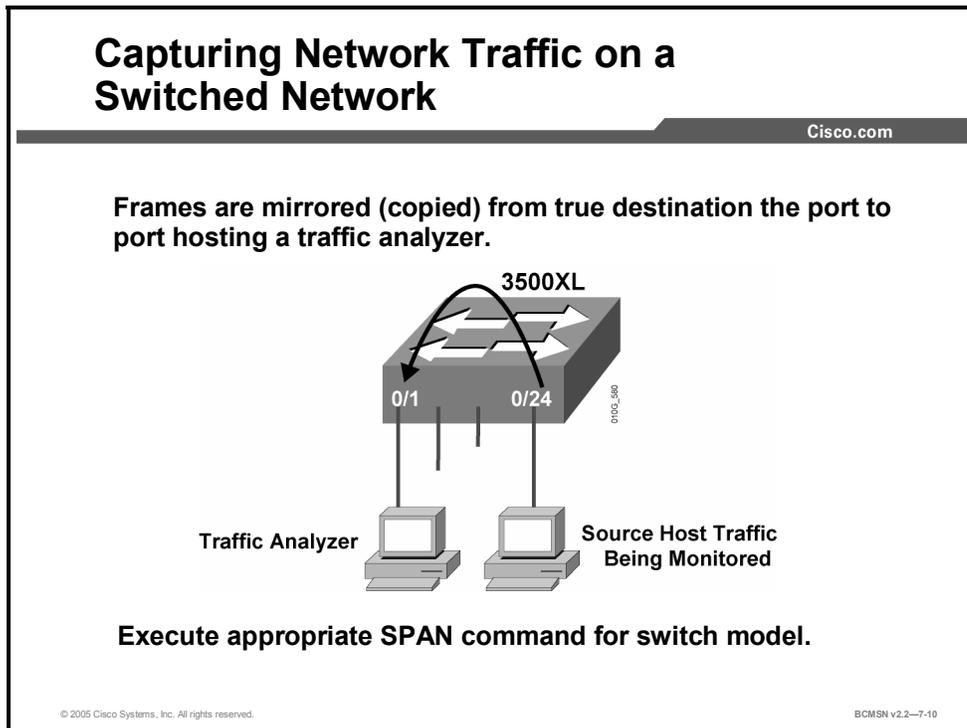
The command is a macro that executes several configuration commands. There is no command such as **no switchport host** to revoke the effect of the **switchport host** command. To return an interface to its default configuration, use the **default interface** *interface-id* global config command. This command returns all interface configurations to the default.

This example shows what occurs when the **switchport host** command is executed.

```
Switch(config-if)# switchport host
switchport mode will be set to access
spanning-tree portfast will be enabled
channel group will be disabled
Switch(config-if)#
```

Capturing Traffic in a Switched Network

This topic describes issues surrounding the capturing of data packets in a switched network.



In order to collect data traffic on a switched network, it is necessary to configure a switch port where a network analyzer will be attached.

Capturing Data in a Switched Network

Before switches were so prevalent in networks, collecting a sample of network traffic was relatively simple. A network analyzer could be plugged into any free port on a hub to capture all data frames on the shared segment because all frames were forwarded out all hub ports. In a switched network, however, frames are usually confined to the switch ports where the sending and receiving devices are located. This makes collecting an overall sample of data on a network segment more challenging.

If a network analyzer is connected to a switch port, it will by default only collect data directed to the MAC address of the analyzer and any broadcast, multicast, or unicast flooding traffic forwarded out that switch port.

To resolve this issue, Cisco switches allow one or more ports to be configured as switch port analyzer (SPAN) ports. SPAN sends a copy of frames generated on one port or an entire VLAN to another switch port hosting a network analyzer. The concept of SPAN is also called port mirroring or port monitoring.

Commands Used in Capturing Network Traffic

Various commands are used across Catalyst platforms to inform the switch about which port carries the traffic of interest and which port will have the network analyzer attached to it.

Here are three commands commonly used to invoke SPAN on Catalyst platforms.

Catalyst SPAN Commands

Catalyst Platform	Command
2950/3550/3750/4000/4500	monitor session
2900XL/3500XL	port monitor
6500 CatOs	set SPAN
8500 IOS	snoop interface

Configuring SPAN on a Local 3500XL

This subtopic shows the commands used to configure SPAN on a local 3500XL switch.

Configuring SPAN on a 3500XL

Cisco.com

```
Switch(config)#interface fastethernet 0/1
```

- Configures interface where packet analyzer is located

```
Switch(config-if)#port monitor fastethernet 0/24
```

- Instructs the monitor port to mirror traffic from the source port

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—7-11

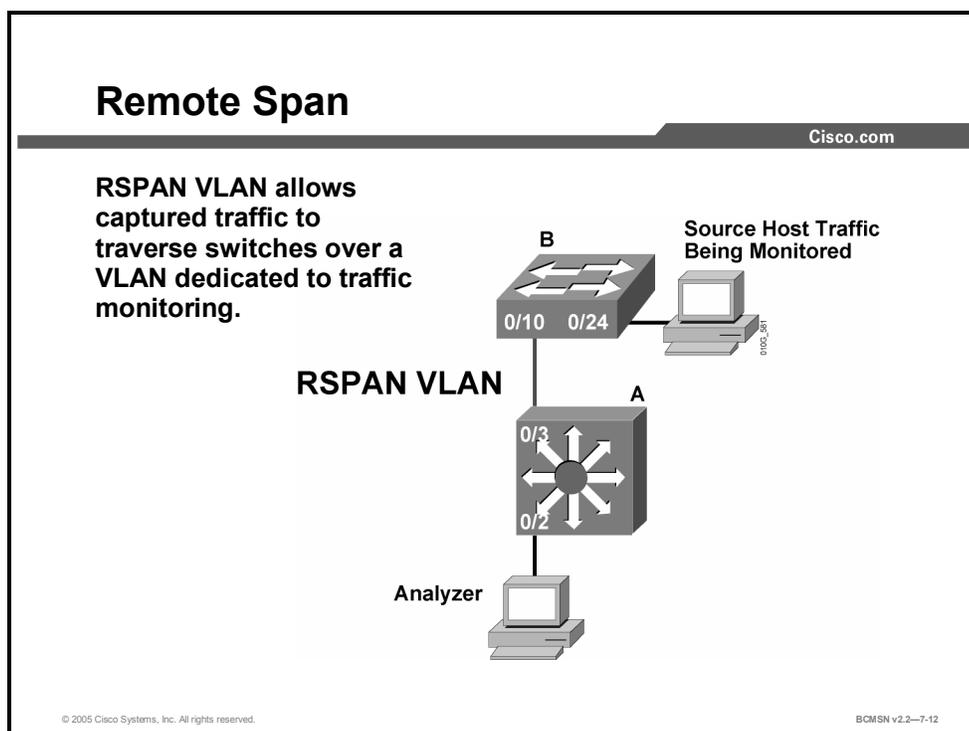
Resources

For additional information, refer to this resource:

http://cisco.com/en/US/products/hw/switches/ps700/products_tech_note09186a008015c612.shtml

Monitoring Performance with RSPAN

This subtopic describes RSPAN.



Remote SPAN (RSPAN) is a variation of SPAN. Rather than sending traffic directly to the traffic analyzer located on the same switch as the port being monitored, RSPAN sends traffic from a monitored port through an intermediate switch network to a traffic analyzer on another switch. RSPAN supports source ports, source VLANs, and destination ports on different switches. It provides remote monitoring of ports on multiple switches across the network, as shown in the figure. The traffic for each RSPAN session is carried over a user-specified RSPAN VLAN that is dedicated for that RSPAN session in all participating switches.

RSPAN consists of an RSPAN source session, an RSPAN VLAN, and an RSPAN destination session. The RSPAN source session must be configured separately from the destination session, given that the two are on different network devices. To configure an RSPAN source session on one network device, associate a set of source ports and VLANs with an RSPAN VLAN. To configure an RSPAN destination session on another device, associate the destination port with the RSPAN VLAN. The intermediate switches need only have the RSPAN VLAN carried over source-to-destination switch links.

RSPAN Guidelines

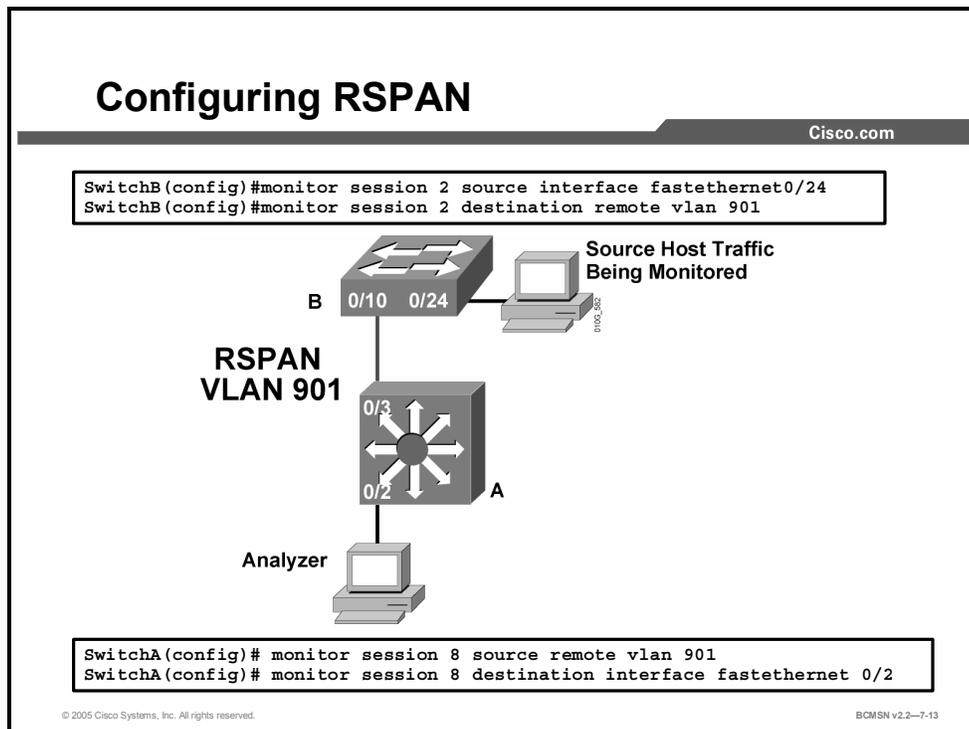
In addition to the guidelines and restrictions that apply to SPAN, the following guidelines apply to RSPAN:

- Networks impose no limit on the number of RSPAN VLANs that the networks carry.
- Intermediate switches might impose limits on the number of RSPAN VLANs that they can support, based on their capacity.
- The RSPAN VLANs must be configured on all source, intermediate, and destination network switches.

- RSPAN VLANs can be used only for RSPAN traffic.
- Access ports must not be assigned to RSPAN VLANs.
- Any ports in an RSPAN VLAN, except those selected to carry RSPAN traffic, should not be configured.
- MAC address learning is disabled on the RSPAN VLAN.
- RSPAN source ports and destination ports must be on different network devices.
- RSPAN VLANs cannot be configured as sources in VLAN SPAN sessions.
- Any VLAN can be configured as an RSPAN VLAN.

Configuring RSPAN

This subtopic provides a sample RSPAN configuration.



This configuration example would be appropriate only for switch models that support the monitor session command. In the example, VLAN 901 has been created as an RSPAN VLAN on both switches.

```
Switch(config)#vlan 901
Switch(config-vlan)# remote-span
```

This example shows how to configure RSPAN on the source switch.

```
SwitchB(config)#monitor session 2 source interface fastethernet0/24
SwitchB(config)#monitor session 2 destination remote vlan 901
```

This example shows how to configure RSPAN on the destination switch.

```
SwitchA(config)# monitor session 8 source remote vlan 901
SwitchA(config)# monitor session 8 destination interface fastethernet 0/2
```

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **CDP packets can expose some network information.**
- **Authentication information and data carried in Telnet sessions are in clear text.**
- **VTY ACLs should be used to limit Telnet access to switch devices.**
- **SSH provides a more secure option for Telnet.**
- **Best practices should be followed when placing switches into service.**
- **Gathering the aggregate network traffic of a switched network requires configuration of SPAN facilities.**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—7-14

Module Summary

This topic summarizes the key points discussed in this module.

Module Summary

Cisco.com

- **Key switch security issues should be identified on a switched network and proper measures taken to mitigate known attacks.**
- **VLAN trunk links should be secured to defend against VLAN hopping attacks.**
- **DHCP snooping, port security, and dynamic ARP inspection are used to protect the network against spoofing attacks.**
- **Implement AAA services to ensure secure authentication, authorization, and accounting and require port authentication using 802.1X.**
- **When placed into service switches should be configured according to best practices to secure the switch device and its protocols from attacks that can be launched through a switch.**

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

Network security is of the highest importance in all organizations that use network technology. Cisco networks provide multiple layers of secure protection against malicious attacks and unauthorized access. Protection is provided from device-level access control to network-wide mitigation of security threats. Security services such as PVLANS, DHCP snooping, and AAA provide an effective array of countermeasures to address the wide variety of network attacks.

References

For additional information, refer to these resources:

- Cisco Systems, Inc., *Catalyst 4500 Series Software Configuration Guide, 7.4: Configuring Port Security*:
http://cisco.com/en/US/products/hw/switches/ps663/products_configuration_guide_chapter09186a00801162fc.html
- Cisco Systems, Inc., *Catalyst 4500 Series Switch IOS Software Configuration Guide, 12.1(19)EW: Configuring Port Security*:
http://cisco.com/en/US/products/hw/switches/ps4324/products_configuration_guide_chapter09186a008019d0de.html
- Cisco Systems, Inc., *SAFE Layer 2 Security In-depth Version 2*:
http://www.cisco.com/en/US/netsol/ns340/ns394/ns171/ns128/networking_solutions_white_paper09186a008014870f.shtml
- Cisco Systems, Inc., *Packet Magazine*, First Quarter, 2003, “Layer2: The Weakest Link”:
http://cisco.com/en/US/partner/about/ac123/ac114/ac173/ac222/about_cisco_packet_feature09186a0080142deb.html

- Cisco Systems, Inc., *VLAN Security White Paper*
http://cisco.com/en/US/partner/products/hw/switches/ps708/products_white_paper09186a008013159f.shtml
- Cisco Systems, Inc., *Configuring Private VLANs (4500 series)*
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/12_1_13/config/pvlans.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 is a procedure used to mitigate VLAN hopping attacks? (Source: Mitigating VLAN Attacks)
- A) turning off DTP on all ports
 - B) turning off CDP on unused ports
 - C) turning on VTP pruning on trunk ports
 - D) turning on MST instead of allowing PVST+
- Q2) To configure dynamic ARP inspection properly, you should _____. (Source: Mitigating Spoof Attacks)
- A) configure all access switch ports as trusted and all trunk ports as trusted
 - B) configure all access switch ports as untrusted and all trunk ports as untrusted
 - C) configure all access switch ports as untrusted and all trunk ports as trusted
 - D) configure all access switch ports as trusted and all trunk ports as untrusted
- Q3) A mechanism that provides isolation between ports within the same VLAN without needing separate IP address space is a _____. (Source: Mitigating VLAN Attacks)
- A) VACL
 - B) PVST+ router
 - C) intra-VLAN router
 - D) private VLAN
- Q4) Cisco network devices support all of the following authentication mechanisms except _____. (Source: Defending Network Switches)
- A) line password
 - B) local password
 - C) local database
 - D) enable password
- Q5) Which type of access control list is applied to VTY lines to control Telnet access to a switch? (Source: Defending Network Switches)
- A) extended IP VACLs
 - B) extended IP access control lists
 - C) standard IP access control lists
 - D) standard IP VACLs
- Q6) What is the nature of a MAC flooding attack? (Source: Securing Switches from Attack)
- A) flooding a network with MAC address resolution attempts
 - B) flooding a network with enough frames so that no other devices can communicate
 - C) flooding a switch CAM table, forcing the switch to flood frames
 - D) flooding a switch with enough frames to consume all CPU resources

Module Self-Check Answer Key

Q1) A

Q2) C

Q3) D

Q4) B

Q5) C

Q6) C

Configuring Campus Switches to Support Voice and Video Applications

Overview

Campus networks carry a variety of data with diverse purposes and impacts on resources. When voice, video, and data application are effectively delivered over a single campus infrastructure, return on investment (ROI) yields are very high for equipment investments. Proper design and configuration efforts will ensure that voice, video, and data traffic efficiently coexist on a single campus infrastructure.

Module Objectives

Upon completing this module, you will be able to configure campus switches to optimize traffic flow when voice, video, and data applications traverse a single converged network. This ability includes being able to meet these objectives:

- Configure an access switch port to support the attachment of a Cisco phone with appropriate trust and voice VLAN configuration.
- Identify which IP multicast features should be enabled on switches at various layers to forward and restrict multicast traffic as needed.

Accommodating Voice Traffic on Campus Switches

Overview

IP telephony services are often provided over the campus infrastructure. To have data and voice application traffic harmoniously coexist, mechanisms must be set in place to differentiate traffic and to offer priority processing to delay sensitive voice traffic. Quality of service (QoS) policies mark and qualify traffic as it traverses the campus switch blocks. Specific VLANs keep voice traffic separate from other data to ensure that it is carried through the network with special handling and with minimal delay. Specific design and implementation considerations should be made at all campus switches supporting Voice over IP (VoIP).

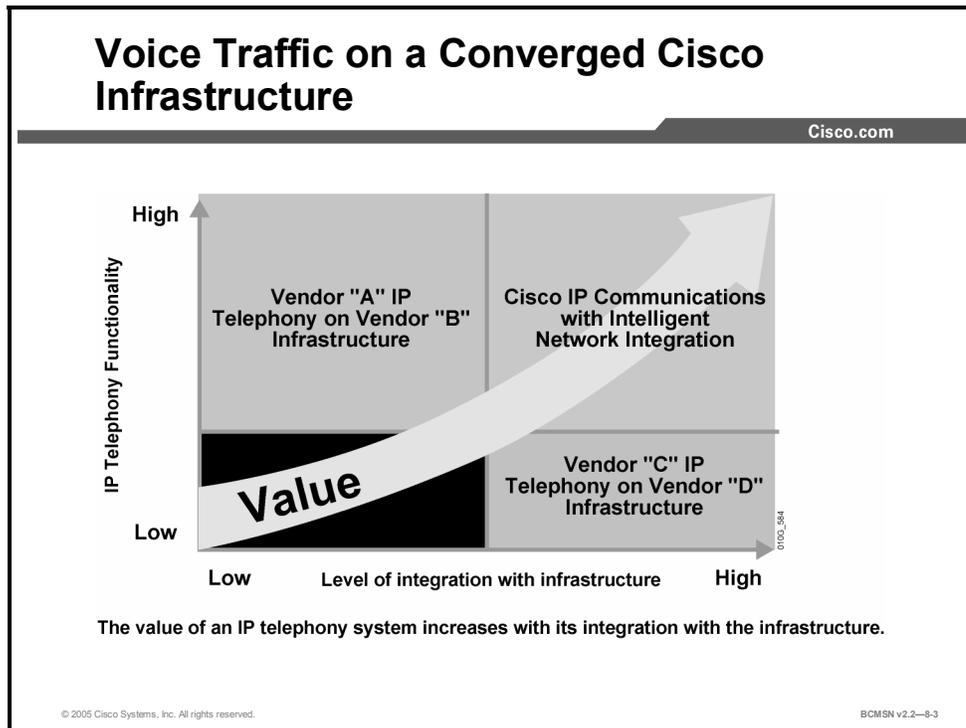
Objectives

Upon completing this lesson, you will be able to configure IP telephony support on an access switch to provide priority processing for voice traffic. This ability includes being able to meet these objectives:

- List the benefits of carrying voice traffic over a Cisco infrastructure
- Define a voice VLAN
- Identify network features that should be implemented in the Building Access and Building Distribution submodules to support voice traffic
- List network design considerations required to support voice traffic and devices
- Define the purpose of QoS in the campus network
- Describe why QoS is needful when voice traffic is present on network devices
- Describe a QoS trust boundary
- Categorize QoS classification and marking capabilities
- Identify basic commands to be considered when voice traffic will traverse a switch
- Configure an access switch for the attachment of a Cisco IP Phone

Voice Traffic on a Cisco Infrastructure

This topic describes how voice traffic impacts a switched campus network.



Cisco's converged end-to-end network solution offers the strengths of the Cisco data networking components such as routers, switches, and firewalls, which have infrastructure security and reliability as a foundation. An IP telephony solution can then be implemented over that network.

The power of this approach is that each new application—such as video, Web, or telephony—represents just another media type to traverse the same infrastructure medium rather than requiring the creation of a different communication medium for each media type. Intelligent devices are automatically given rights and priorities, and the applications themselves can intelligently communicate with the infrastructure to meet the constantly changing needs of the system as specified by the organization. A network infrastructure where IP Telephony requirements are addressed alongside standard network operations is a hallmark of the Cisco IPT converged solution.

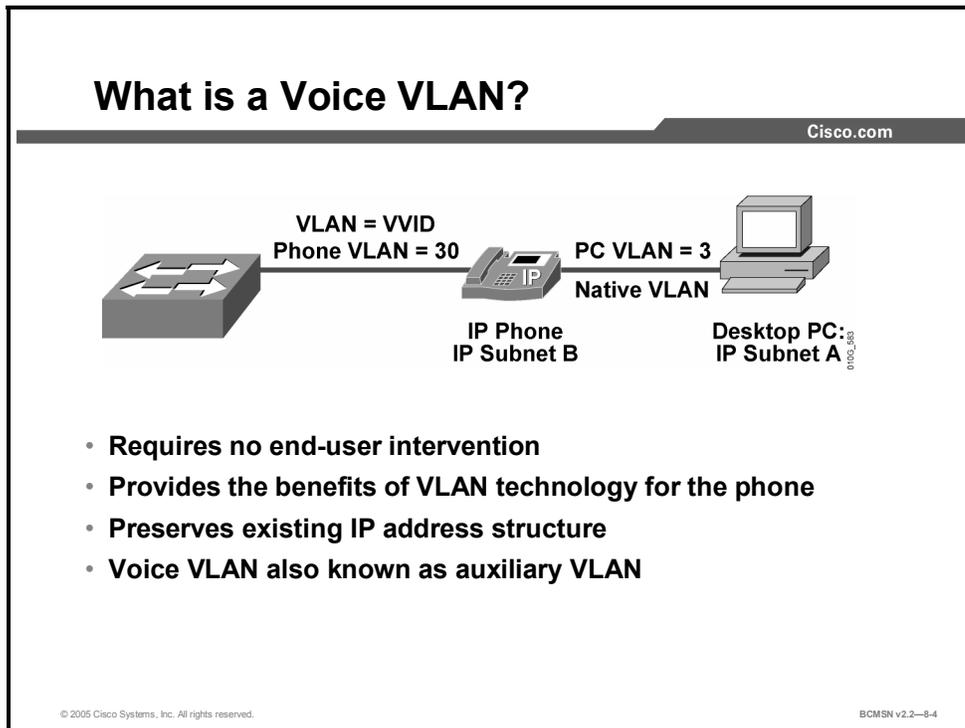
Benefits of IP Telephony on a Cisco Infrastructure

Cisco IP phones are able to use the Ethernet switches in the network as the "voice call switch matrix." Calls are managed differently, and the inherent time slot and bandwidth limitations of traditional time-division multiplexing (TDM) architectures are removed. Switching of a call is done only between the devices required to switch the call: the IP phones, voice gateways, and Ethernet switches. Calls do not have to be routed back to a traditional TDM switching matrix to complete the call.

Cisco IP phones are also able to receive call processing capability directly from the Cisco IOS software running on the access router for remote or small office locations. The tight integration

What Is a Voice VLAN?

This topic describes the use of voice VLANs in a switched IP telephony network.

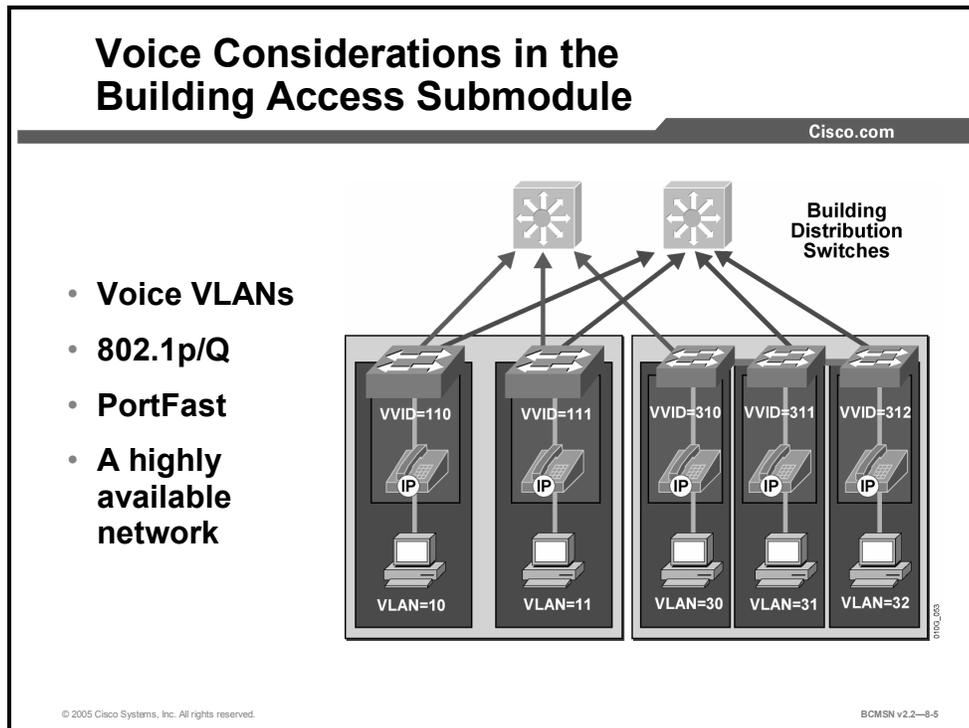


Some Cisco Catalyst switches offer a “voice VLAN” feature. The voice VLAN, also known as an auxiliary VLAN, provides automatic VLAN association for IP phones. Since the phones, and therefore the phone traffic, are associated with a specific VLAN, the phone traffic will be on different IP subnets even though voice and data coexist on the same physical infrastructure. A voice VLAN will have a specific VLAN ID, which is referred to as the voice VLAN ID (VVID).

When a phone is connected to the switch, the switch sends necessary voice VLAN information to the IP phone, placing it into the voice VLAN without end-user intervention. Placing phone traffic onto a distinct VLAN allows it to be segmented from the data traffic; this facilitates better network management and troubleshooting. Additionally, QoS or security policies can be enforced specifically for the traffic traversing the phone VLANs without affecting the data traffic. If the phone is moved, upon connection, it will again be associated with an appropriate voice VLAN even though its physical location may have changed.

Voice Considerations in Campus Submodules

This topic describes voice traffic considerations in the campus network.



Deploying IP telephony in the enterprise campus requires the implementation of various features particular to each submodule.

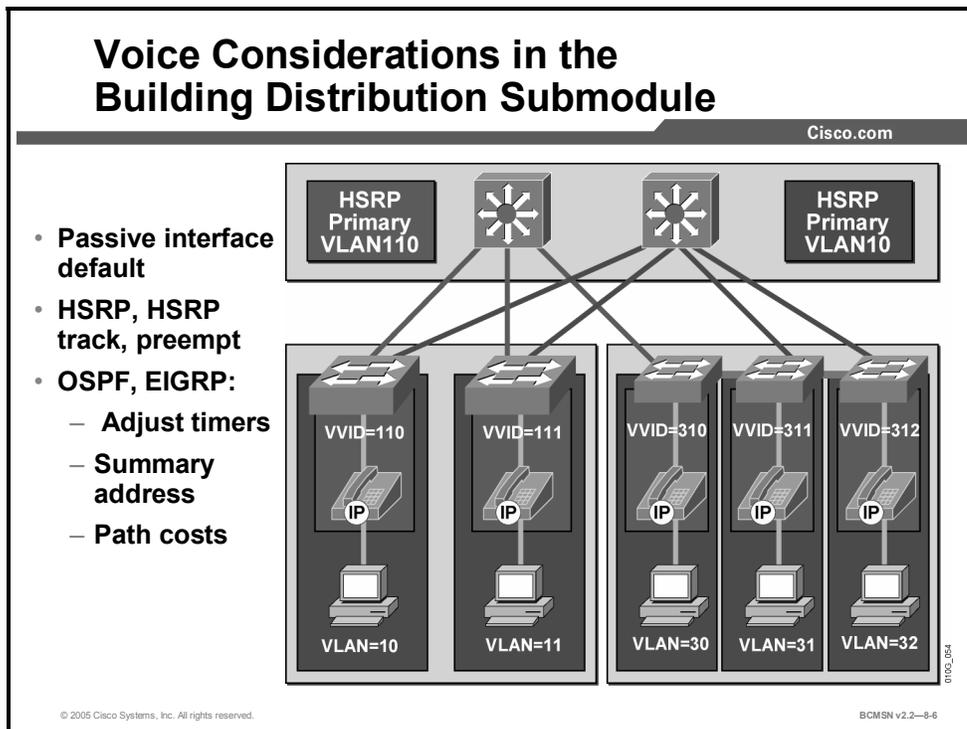
Building Access Submodule

Within the Building Access submodule, these features support IP telephony:

- Voice VLANs
- 802.1p/Q
- Hardware support for multiple output queues
- Hardware support for inline power to IP phones
- PortFast
- Root guard
- Unidirectional Link Detection (UDLD)
- UplinkFast

Building Distribution Submodule

This subtopic describes voice VLANs in the Building Distribution submodule.



Within the Building Distribution submodule, these features support IP telephony:

- Passive interfaces
- Layer 3 redundancy with Hot Standby Router Protocol (HSRP), HSRP track, and HSRP preempt
- Open Shortest Path First (OSPF) or Enhanced Interior Gateway Routing Protocol (EIGRP) routing with adjusted timers, summary addresses, and path costs

Network Design Considerations for Voice

This topic describes the network and device design considerations to support voice traffic.

Network Design for Voice Traffic and Devices: General Considerations

Cisco.com

- **What features are required for each campus network device?**
- **Can the physical plant support IP telephony?**
- **How will the phones be powered?**
- **Is adequate bandwidth available?**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—8-7

IP telephony requires that a network provide sufficient bandwidth and quick convergence after failures or changes. Most IP telephony installations are built on an existing network infrastructure, which requires configuration with priority given to voice traffic.

General Design Considerations

To determine if an infrastructure can support the addition of voice, evaluate these considerations.

- **Features required for each device in the campus network:** IP phones require power, and most enterprises put IP telephony applications on a separate VLAN with priority handling.
- **Physical plant capable of supporting IP telephony:** The wiring and cabling plant must be adequate for IP telephony needs. At a minimum, Category 5 cabling is required, and consideration should be given to increasing wall jacks and switch ports to support phone and PC connections.
- **Provision switches with inline power to support IP phones:** Deploy inline power to the IP Phones through a Catalyst Inline Power Patch Panel or through the individual switch ports using inline power modules in the Catalyst . This may increase the power requirements of the switch itself.
- **Network bandwidth adequate for data, voice, and call control traffic:** Along with data traffic, consider both voice and call control traffic loads. There should be no steady-state congestion or latency over the LAN links. This is critical for voice operations in the LAN.

Note Plan to work with a voice specialist to complete traffic engineering analysis for the network.

Bandwidth Provisioning

This subtopic describes bandwidth provisioning in a voice network.

Network Design for Voice Traffic and Devices: Bandwidth Provisioning

Cisco.com

- **Consider voice call control and call traffic.**
- **Consider multicast and data traffic.**
- **Do not exceed 75 percent of the total available bandwidth for each link.**
- **Provision for voice bearer traffic:**
 - **Voice bearer traffic (bps) = (packet payload + all headers in bits) * (packet rate per second)**

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

Properly provisioning the network bandwidth is a major component of designing a successful IP telephony network. The required bandwidth can be calculated by adding the bandwidth requirements for each major application, including voice, video, and data. This sum represents the minimum bandwidth requirement for any given link, and it should not exceed approximately 75 percent of the total available bandwidth for the link.

From a traffic standpoint, an IP telephony call consists of two traffic types:

- **Voice carrier stream:** This consists of Real-Time Transport Protocol (RTP) packets that contain the actual voice samples.
- **Call control signaling:** This consists of packets belonging to one of several protocols—those used to set up, to maintain, to tear down, or to redirect a call, depending upon call endpoints. Examples are H.323 or Media Gateway Control Protocol (MGCP).

A Voice over IP (VoIP) packet consists of the voice payload, IP header, User Datagram Protocol (UDP) header, RTP header, and Layer 2 link header. Coder-decoder (codec) type (G.711, G.729, and so forth) is configurable by device. However, G.729 does not support fax or modem traffic. The IP header is 20 bytes, the UDP header is 8 bytes, and the RTP header is 12 bytes. The link header varies in size according to the Layer 2 media used; Ethernet requires 14 bytes of header. The voice payload size and the packetization period are device dependent.

To calculate the bandwidth that voice streams consume, use this formula:

(Packet payload + all headers in bits) * Packet rate per second—for example, 50 packets per second (pps) when using a 20-ms packet period.

Power Considerations

This subtopic discusses power considerations in a voice network.

Network Design for Voice Traffic and Devices: Power Considerations

Cisco.com

- **Inline power or power patch panel for IP Phones**
- **UPS and generator backup**
- **UPS systems with autorestart capability**
- **UPS system monitoring**
- **A 4-hour service response contract for UPS system problems**
- **Recommended equipment operating temperatures maintained 24/7**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—8-9

Accurate calculations of power requirements are critical for an effective IP telephony solution. Power can be supplied to the IP phones directly from Catalyst switches with inline power capabilities or by inserting a Catalyst Inline Power Patch Panel. In addition to IP phones, failover power and total load must be considered for all devices in the IP telephony availability definition, including Building Distribution and Campus Backbone submodules, gateways, Cisco CallManager, and other servers and devices. Power calculations, therefore, must be network based rather than device based.

Providing highly available power protection requires an uninterruptible power supply (UPS) with a minimum battery life to support 1 hour and a 4-hour response for power system failures, or a generator with an onsite service contract. This solution must include UPS or generator backup for all devices associated with the IP telephony network. In addition, consider UPS systems that have autorestart capability and a service contract for 4-hour support response.

Recommendations for IP telephony high-availability power and environment:

- UPS and generator backup
- UPS systems with autorestart capability
- UPS system monitoring
- A 4-hour service response contract for UPS system problems
- Recommended equipment operating temperatures maintained 24/7

Intelligent Network Services

This subtopic discusses intelligent services needed on a voice network.

Network Design for Voice Traffic and Devices: Intelligent Services

Cisco.com

- **High availability should provide redundancy and failover for critical components.**
- **Provide security features considered on all enterprise network devices.**
- **Configure QoS for voice traffic:**
 - **Determine how voice will be classified.**
 - **Determine which queuing method to use for voice.**

© 2005 Cisco Systems, Inc. All rights reserved.

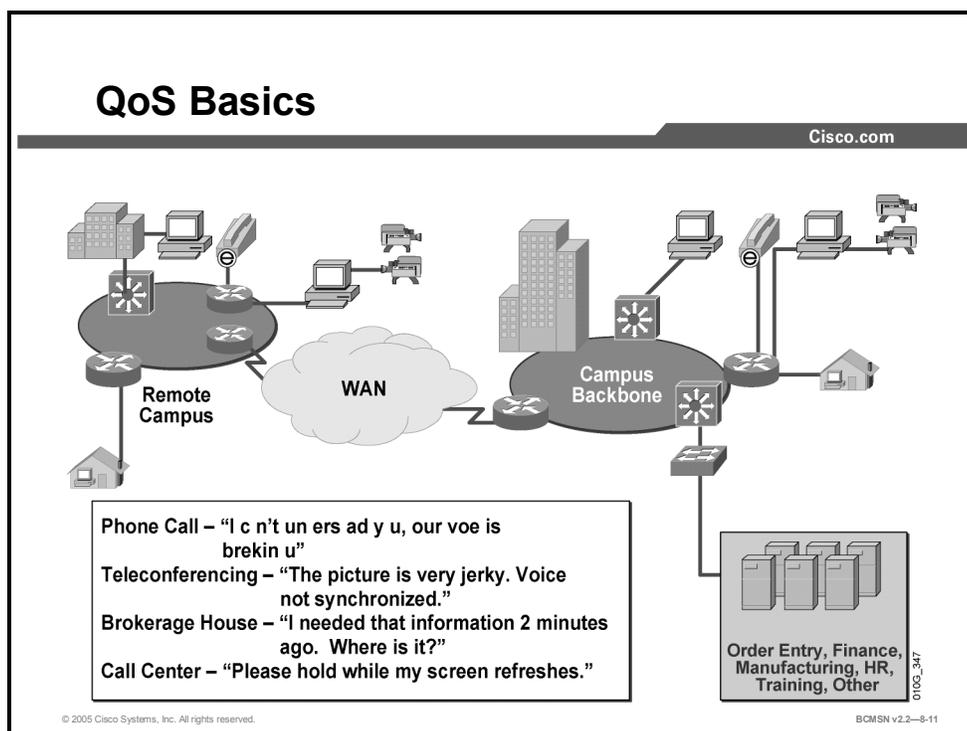
BCMSN v2.2—8-10

Network management, high availability, security, QoS, and intelligent network services must expand to incorporate voice-specific attributes.

- **Network management:** The merging of network management tasks associated with both voice and data networks is one of the key benefits of using a converged network as opposed to a voice-only network. However, it is still necessary to understand the traditional voice-only management concepts to relate the features available in that technology to the converged network management techniques.
- **High availability:** As with any network capability, plan redundancy for critical voice network components such as Cisco CallManager and the associated gateway and infrastructure devices.
- **Security:** The subject of securing voice communications has received more visibility recently, as network convergence becomes an accepted design model. With the advent of IP telephony traffic traversing the LAN infrastructure, the potential exists for malicious attacks on call-processing components and telephony applications. As with all network devices, there should be a predefined security policy for all devices, applications, and users associated with the voice network that is appropriate for the level of caution required. Consider security measures for voice call-processing platforms, applications, and telephony traffic.
- **QoS:** The goal of QoS is to provide critical applications with a higher priority for service so that they are the least likely to be delayed or to be dropped in times of congestion. When a network becomes congested, some traffic will be delayed or lost. Voice traffic has strict requirements concerning delay and delay variation (also known as “jitter”), and, compared to most data traffic, voice traffic is relatively intolerant of loss. To establish priority processing for voice traffic, you can employ a wide range of IP QoS features, such as classification, queuing, congestion detection, traffic shaping, and compression.

QoS Basics

This topic describes the features and attributes of QoS.



Network managers must be prepared for increasing amounts of traffic, requiring more bandwidth than is currently available. This is especially important when dealing with voice traffic. Almost any network can take advantage of QoS for optimum efficiency, whether it is a small corporate network, an Internet service provider, or an enterprise network. QoS is the application of features and functionality required to actively manage and satisfy the networking requirements of applications sensitive to loss, delay, and delay variation (jitter). QoS allows preference to be given to critical application flows for the available bandwidth. QoS tools enable manageability and predictable service for a variety of networked applications and traffic types in a complex network.

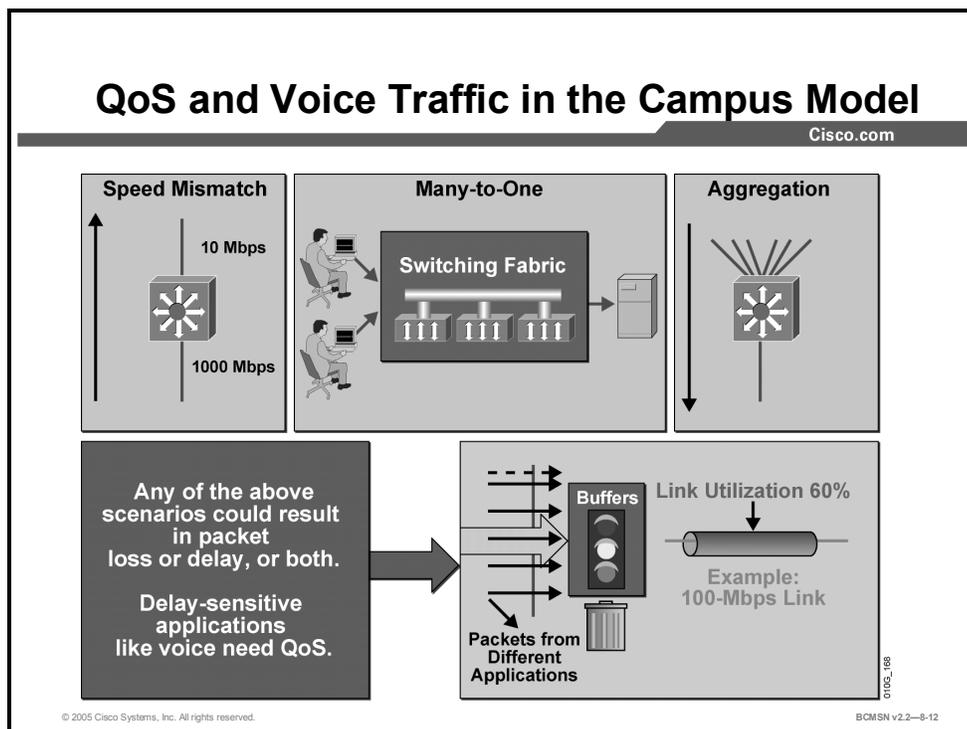
The Cisco IOS implementation of QoS software provides these benefits.

- **Priority access to resources:** QoS allows administrators to control which traffic is allowed to access specific network resources such as bandwidth, equipment, and WAN links. Critical traffic may take possession of a resource because the QoS implementation drops low-priority frames.
- **Efficient management of network resources:** If network management and accounting tools indicate that specific traffic is experiencing latency, jitter, and packet loss, then QoS tools can be used to adjust how that traffic is handled.
- **Tailored services:** The control provided by QoS enables Internet service providers to offer carefully tailored grades of service to their customers. For example, a service provider can offer one service level agreement (SLA) to a customer website that receives 3000 to 4000 hits per day and another to a site that receives only 200 to 300 hits per day.
- **Coexistence of mission-critical applications:** QoS technologies ensure that mission-critical business applications receive priority access to network resources while providing

adequate processing for applications that are not delay sensitive. Multimedia and voice applications tolerate little latency and require priority access to resources. Other delay-tolerant traffic traversing the same link, such as Simple Mail Transfer Protocol (SMTP) over TCP, can still be adequately serviced.

QoS and Voice Traffic in the Campus Module

This topic describes how QoS is applied for voice traffic in the campus module.



Regardless of the speed of individual switches or links, speed mismatches, many-to-one switching fabrics, and aggregation may cause a device to experience congestion which can result in latency. If congestion occurs and congestion management features are not in place, then some packets will be dropped, causing retransmissions that inevitably increase overall network load. QoS can mitigate latency caused by congestion on campus devices.

QoS is implemented by classifying and marking traffic at one device while allowing other devices to prioritize or to queue the traffic according to those marks applied to individual frames or packets. The table lists the campus devices involved in QoS marking or prioritizing.

QoS Application in the Campus Network

Campus Device	QoS Application
Access Layer	Initial point at which traffic enters the network. Traffic can be marked (or remarked) at Layers 2 and 3 by the access switch as it enters the network or "trusted" that it is entering the network with an appropriate tag.
Distribution Layer	Marks of traffic inbound from the access layer can be trusted or reset depending on the ability of the access layer switches. Priority access into the core is provided based on Layer 3 QoS tags.
Core	No traffic marking occurs at the core. Layer 2 or 3 QoS tags are trusted from distribution layer switches and used to prioritize and to queue the traffic as it traverses the core.

Network Availability Problem Areas

This subtopic describes network availability issues of concern in voice traffic.

Network Availability Problem Areas

Cisco.com

Delay (Latency)

Delay Variation (Jitter)

Packet Loss

Network administrators need a way to manage problem areas on an application basis.

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2-8-13

An enterprise network may experience any of these network availability problems.

Delay: Delay (or latency) is the amount of time that it takes a packet to reach the receiving endpoint after being transmitted from the sending endpoint. This time period is termed the “end-to-end delay” and can be broken into two areas: fixed network delay and variable network delay. Fixed network delay includes encoding and decoding time (for voice and video) as well as the amount of time required for the electrical and optical pulses to traverse the media en route to their destination. Variable network delay generally refers to network conditions, such as congestion, that may affect the overall time required for transit. In data networks, for example, these types of delay occur:

- **Packetization delay:** The amount of time that it takes to segment data (if necessary), sample and encode signals (if necessary), process data, and turn the data into packets
- **Serialization delay:** The amount of time that it takes to place the bits of a packet, encapsulated in a frame, onto the physical media
- **Propagation delay:** The amount of time that it takes to transmit the bits of a frame across the physical wire
- **Processing delay:** The amount of time that it takes for a network device to take the frame from an input interface, place it into a receive queue, and then place it into the output queue of the output interface
- **Queuing delay:** The amount of time that a packet resides in the output queue of an interface

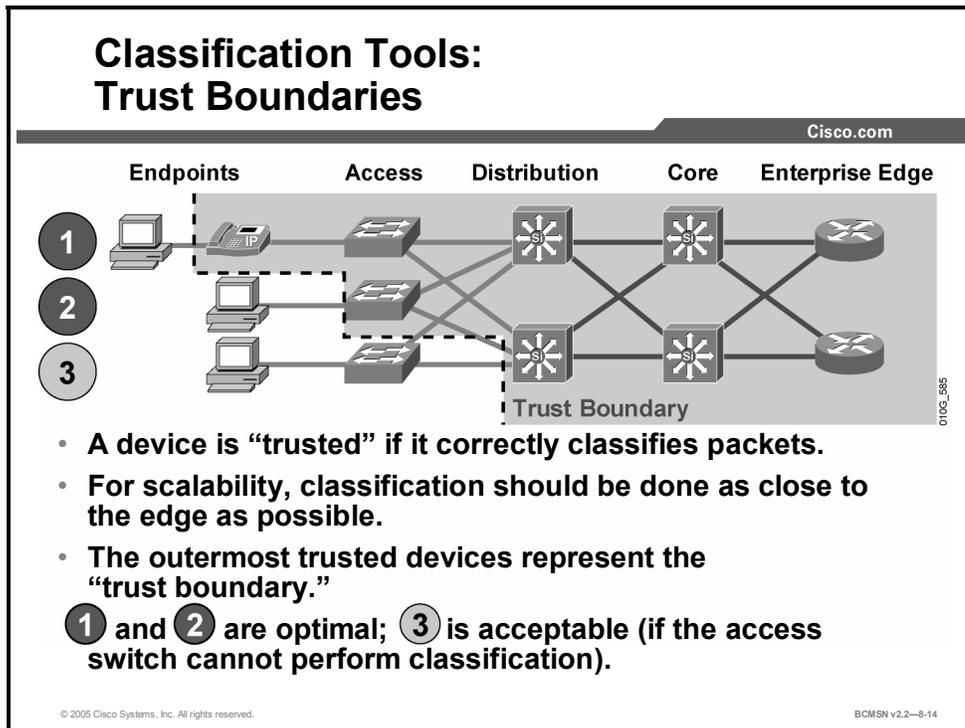
- **Delay variation:** Delay variation (or jitter) is the difference in the end-to-end delay between packets. For example, if one packet requires 100 ms to traverse the network from the source endpoint to the destination endpoint, and the following packet requires 125 ms to make the same trip, then the delay variation is calculated as 25 ms.

Each end station and Cisco network device in a voice or video conversation has a jitter buffer. Jitter buffers are used to smooth out changes in arrival times of data packets containing voice and video. A jitter buffer is dynamic and can adjust for changes in arrival times of packets. If you have instantaneous changes in arrival times of packets that are outside the capabilities of a jitter buffer to compensate, you will have one of these situations:

- A jitter buffer underrun, when arrival times between packets containing voice or video increase to the point where the jitter buffer has been exhausted and contains no packets to process the signal for the next piece of voice or video.
 - A jitter buffer overrun, when arrival times between packets containing voice or video decrease to the point where the jitter buffer cannot dynamically resize itself quickly enough to accommodate. When an overrun occurs, packets are dropped and voice quality is degraded.
- **Packet loss:** Packet loss is a measurement of packets transmitted and received compared to the total number that were transmitted. Loss is expressed as the percentage of packets that were dropped. Tail drops occur when the output queue is full. These are the most common drops that can occur when a link is congested. Other types of drops (input, ignore, overrun, no buffer) are not as common but may require a hardware upgrade because they are usually a result of network device congestion.

QoS Trust Boundaries

This topic describes QoS trust boundaries.



In a campus QoS implementation, boundaries are defined where the existing QoS values attached to frames and to packets are to be accepted or altered. These “trust boundaries” are established by configuring trust levels on the ports of key peripheral network devices where QoS policies will be enforced as traffic makes its way into the network. At these boundaries, traffic will be allowed to retain its original QoS marking or have new marking ascribed as a result of policies associated with its entry point into the network.

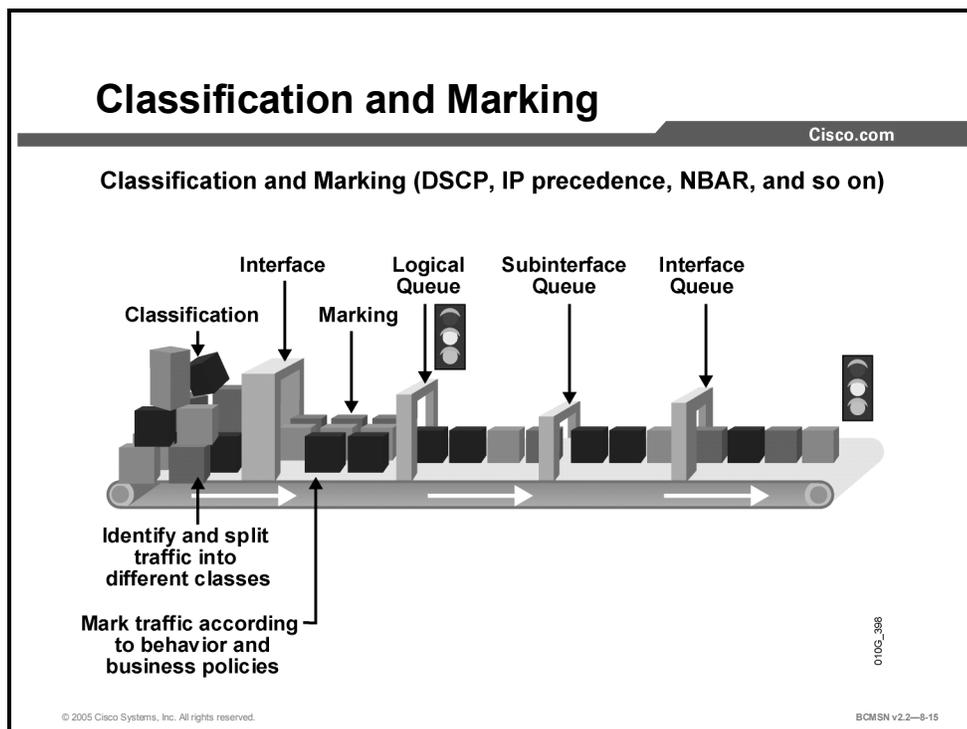
Trust boundaries establish a border for traffic entering the campus network. As traffic traverses the switches of the campus network, it is handled and prioritized according to the marks received or trusted when the traffic originally entered the network at the trust boundary.

At the trust boundary device, QoS values are trusted if they are considered to accurately represent the type of traffic and precedence processing the traffic should receive as it enters the campus network. If untrusted, the traffic will be marked with a new QoS value appropriate for the policy in place at the point where the traffic entered the campus network. Ideally, the trust boundary exists at the first switch receiving traffic from a device or IP phone. It is also acceptable to establish the trust boundary as all the traffic from an access switch enters a Building Distribution layer port.

Note Best practices suggest classifying and marking traffic as close to the traffic source as possible.

QoS Traffic Classification and Marking

This topic describes traffic classification and marking for QoS.



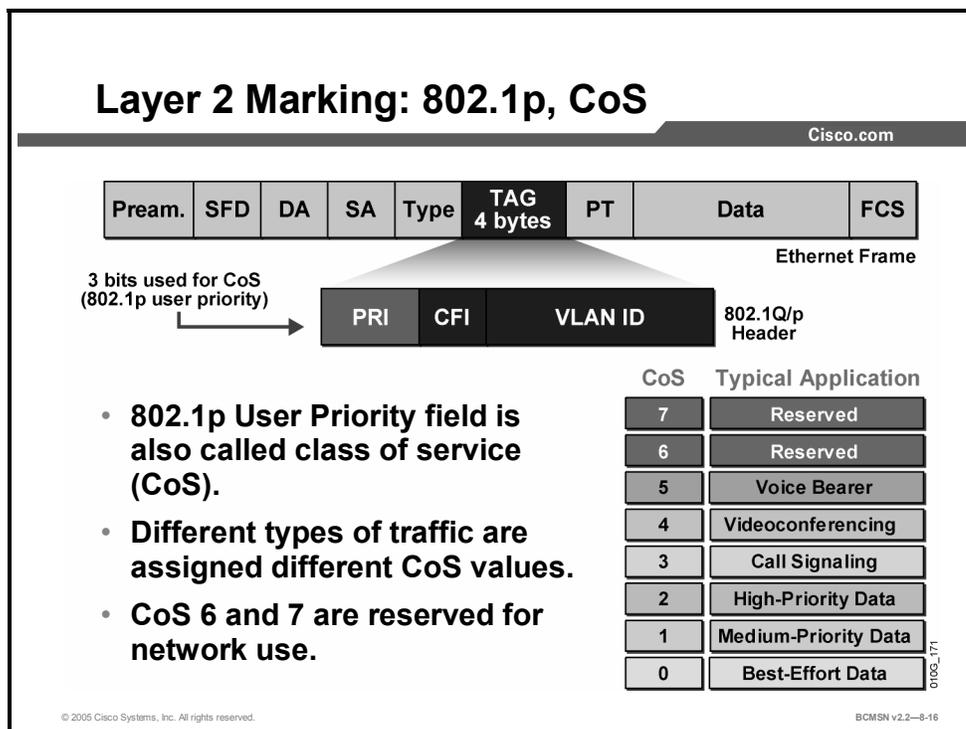
Classification and marking is the process of identifying traffic for proper prioritization as that traffic traverses the campus network. Traffic is classified by examining information at various layers of the Open Systems Interconnection (OSI) model. All traffic classified in a certain manner will receive an associated mark or QoS value. IP traffic can be classified according to any values configurable in an access control list (ACL) or any of the following criteria:

- **Layer 2 parameters:** MAC address, Multiprotocol Label Switching (MPLS), ATM cell loss priority (CLP) bit, Frame Relay discard eligible (DE) bit, ingress interface
- **Layer 3 parameters:** IP precedence, differentiated services code point (DSCP), QoS group, IP address, ingress interface
- **Layer 4 parameters:** TCP or UDP ports, ingress interface
- **Layer 7 parameters:** Application signatures, ingress interface

All traffic classified or grouped according to these criteria will be marked according to that classification. QoS marks or values establish priority levels or priority classes of service for network traffic as it is processed by each switch. Once traffic is marked with a QoS value, then QoS policies on switches and interfaces will handle traffic according to the values contained in individual frames and packets. As a result of classification and marking, traffic will be prioritized accordingly at each switch to ensure that delay-sensitive traffic receives priority processing as the switch manages congestion, delay, and bandwidth allocation.

Layer 2 QoS Marking

This subtopic describes how QoS values are carried in the Layer 2 header.



QoS Layer 2 classification occurs by examining information in the Ethernet or 802.1Q header such as destination MAC address or VLAN ID. QoS Layer 2 marking occurs in the Priority field of the 802.1Q header. LAN Layer 2 headers have no means of carrying a QoS value, so 802.1Q encapsulation is required if Layer 2 QoS marking is to occur. The Priority field is 3 bits long and is also known as the 802.1p User Priority or Class of Service (CoS) value.

This 3-bit field hosts CoS values ranging from 1 to 7, 1 being associated with delay tolerant traffic such as TCP/IP. Voice traffic, which by nature is not delay tolerant, receives higher default CoS values, such as 3 for Call Signaling. A CoS value of 5 is given to Voice Bearer traffic, which is the phone conversation itself in which voice quality is impaired if any packets are dropped or delayed.

As a result of Layer 2 classification and marking, the following QoS operations can occur:

- **Input queue scheduling:** When a frame enters a port, it can be assigned to one of a number of port-based queues prior to being scheduled for switching to an egress port. Typically, multiple queues are used where traffic requires different service levels.
- **Policing:** Policing is the process of inspecting a frame to see if it has exceeded a predefined rate of traffic within a certain time frame that is typically a fixed number internal to the switch. If a frame is determined to be in excess of the predefined rate limit, it can either be dropped, or the CoS value can be marked down.
- **Output queue scheduling:** The switch will place the frame into an appropriate outbound (egress) queue for switching. The switch will perform buffer management on this queue by ensuring that the buffer does not overflow.

Layer 3 QoS Marking

This subtopic describes QoS information carried in Layer 3 headers.

Layer 3 Marking: IP Precedence, DSCP

Cisco.com

Version Length	ToS Byte	Len	ID	Offset	TTL	Proto	FCS	IP SA	IP DA	Data
----------------	----------	-----	----	--------	-----	-------	-----	-------	-------	------

IPv4 Packet

7	6	5	4	3	2	1	0
IP Precedence			Unused				
DiffServ Code Point (DSCP)						Flow Ctrl	

Standard IPv4

DiffServ Extensions

- IPv4
 - Three most significant bits of ToS byte are called IP precedence.
 - Other bits are unused.
- DiffServ
 - Six most significant bits of ToS byte are called DiffServ Code Point (DSCP)
 - DSCP is backward compatible with IP precedence.
 - Remaining two bits are used for flow control.

© 2005 Cisco Systems, Inc. All rights reserved.
BCMSN v2.2—8-17

QoS Layer 3 classification results from the examination of header values such as destination IP address or protocol. QoS Layer 3 marking occurs in the Type of Service (ToS) byte in the IP header. The first 3 bits of the ToS byte are occupied by IP Precedence, which correlates to the 3 CoS Bits carried in the Layer 2 header.

The ToS byte can also be used for Differentiated Services Code Point (DSCP) marking. DSCP allows prioritization hop by hop as packets are processed on each switch and interface. The ToS bits are used by DSCP values as shown in the table. The first 3 DSCP bits, correlating to Precedence and CoS, identify the DSCP Class of Service for the packet.

ToS Byte:	P2	P1	P0	T3	T2	T1	T0	Zero
DS Byte:	DS5	DS4	DS3	DS2	DS1	DS0	ECN1	ECN0
	(Class Selector)			(Drop Precedence)				

The next 3 DS bits establish a drop precedence for the packet. Packets with a high DSCP drop precedence value will be dropped before those with a low value if a device or a queue becomes overloaded and must drop packets. Voice traffic will be marked with a low DSCP drop precedence value to minimize voice packet drop.

Each 6-bit DSCP value is also given a DSCP Codepoint name. DSCP classes 1-4 are Assured Forwarding (AF) classes. Therefore, if the DSCP class value is 3 and the drop precedence is 1, the DSCP Codepoint would be AF31.

Basic Switch Commands to Support Attachment of a Cisco IP Phone

This topic describes Catalyst switch commands associated with attachment of a Cisco IP phone.

About Basic Switch Commands to Support Attachment of a Cisco IP Phone

Cisco.com

Configure Voice VLAN

- **switchport voice vlan 110**

Configure Trust and CoS options

- **mls qos trust cos**
- **mls qos trust device cisco-phone**
- **mls qos extend trust**
- **switchport priority extend cos *cos_value***

Verify Configuration

- **show interfaces fa 0/4 switchport**
- **show mls qos interface fa 0/4**

© 2005 Cisco Systems, Inc. All rights reserved.

BCMSN v2.2—8-18

These commands are used to configure and verify two basic required functions on a switch port connected to an IP phone with a PC connected to that phone.

Switchport Voice over IP Commands

Command	Description
Switch(config-if)# switchport voice vlan vlan-id	Instructs the Cisco IP phone to forward all voice traffic through the specified VLAN. By default, the Cisco IP phone forwards the voice traffic with an 802.1Q priority of 5.
Switch(config-if)# mls qos trust cos Switch(config-if)# mls qos trust device cisco-phone	Sets the switch port to trust the CoS values received on this port only if CDP discovers that a Cisco phone is attached. Effectively the trust boundary has been moved to the phone.
Switch(config-if)# mls qos extend trust	Sets the IP phone access port to trust the CoS value of frames received from the PC or attached device.
Switch(config-if)# switchport priority extend cos cos_value	Sets the IP phone access port to override the CoS value of frames received from the PC or attached device and replace them with the given value.
Switch# show interfaces interface-id switchport	Displays IP telephony support configured on the interface.
Switch# show mls qos interface interface-id	Displays QoS support configured on the interface.

How to Configure a Switch for Attachment of a Cisco IP Phone

This topic describes how to configure and verify fundamental VoIP features on a Catalyst switch.

How to Configure a Switch for Attachment of a Cisco IP Phone

Cisco.com

- Voice traffic tagged for voice VLAN
- Data VLAN traffic from PC can be
 - Untrusted
 - Trusted
 - Set to a specific value

The diagram illustrates a network configuration for a Cisco IP Phone. A Catalyst switch (2950, 2955, or 3550) is connected to a 3-Port Switch via Fa 0/4 (VLAN 110). The 3-Port Switch has three ports: P1, P2, and P3. P2 is connected to a Cisco IP Phone 7960, which contains a Phone ASIC. P3 is connected to a PC. The PC is connected to the 3-Port Switch via VID 10. The Catalyst switch is connected to the 3-Port Switch via Fa 0/4 (VLAN 110).

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—8-19

These commands are used to configure and to verify basic features used to manage voice traffic on Catalyst switch ports.

Step	Description
1.	Enable voice VLAN on a switch port and associate a VLAN ID. <code>Switch(config-if)# switchport voice vlan vlan-id</code>
2.	Trust the CoS value of frames as they arrive at the switch port. <code>Switch(config-if)# mls qos trust cos</code>
3.	Make this trust conditional on a Cisco IP phone being attached. <code>Switch(config-if)# mls qos trust device cisco-phone</code> Or Set the CoS value to frames coming from the PC attached to the IP phone. <code>Switch(config-if)# switchport priority extend cos cos_value</code>
3.	Display voice parameters configured on the interface. <code>Switch# show interfaces interface-id switchport</code>
4.	Display QoS parameters configured on the interface. <code>Switch# show mls qos interface interface-id</code>

Example

```
Switch(config)# interface fastethernet 0/4
Switch(config-if)# switchport voice vlan 110
Switch(config-if)# mls qos trust cos
Switch(config-if)# mls qos trust device cisco-phone
Switch(config-if)# ctrl-Z
Switch# show interfaces fastethernet 0/4
```

```
Switch# show mls qos interface fastethernet 0/4
FastEthernet0/4
trust state: trust cos
trust mode: trust cos
COS override: dis
default COS: 0
pass-through: none
trust device: cisco-phone
```

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Voice traffic is optimally carried over a Cisco infrastructure.**
- **Voice VLANs are used to facilitate VoIP communication.**
- **Voice traffic should be given distinct consideration in the Access and Distribution submodules.**
- **Follow specific principles for effective network design when implementing VoIP.**
- **QoS is used to prioritize traffic on the campus network.**
- **It is needful to consider QoS when voice traffic is present in a campus network.**
- **Configuration is necessary to establish network trust boundaries.**
- **Classification and marking can be based on received 802.1p CoS bits.**
- **Basic switch commands are considered when voice traffic will traverse a switch.**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—8-20

Configuring IP Multicast

Overview

IP multicast delivers video streams to multiple end users with minimal consumption of campus network resources. Implementing IP multicast requires understanding of IP multicast address structure as well as routing protocols and processes used to carry multicast traffic between campus switches. Three notable IP multicast protocols and processes are Protocol Independent Multicast (PIM), Internet Group Management Protocol (IGMP) and Reverse Path Forwarding (RPF). PIM and IGMP ensure that multicast traffic streams are forward onto Layer 3 segments where receivers exist and are removed from segments with no receivers. On access switches, flooding of multicast packets is minimized using IGMP snooping. Reverse Path Forwarding maintains a loop free topology as multicast streams are forwarded to an IP multicast enabled router.

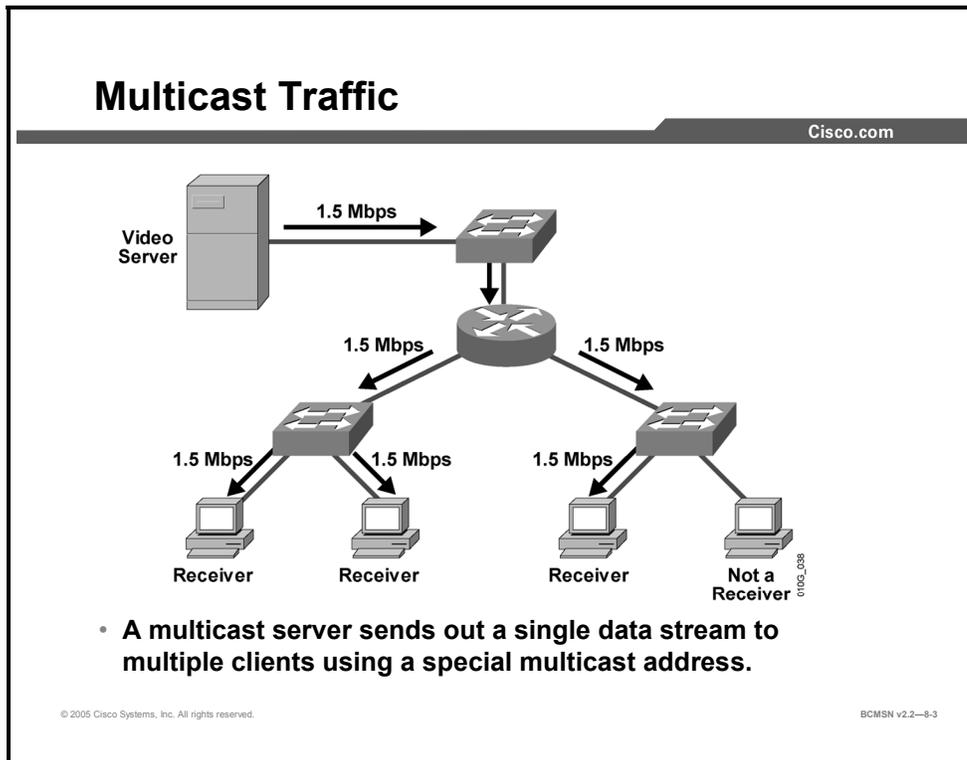
Objectives

Upon completing this lesson, you will be able to identify what IP multicast features to enable to optimize the forwarding of multicast packets across network switches and routers. This ability includes being able to meet these objectives:

- Describe IP multicast
- Describe IP multicast address conventions in IPv4
- Define RPF
- Define PIM
- Describe the three modes of PIM
- Define IGMP
- Describe the process used on Layer 2 switches to snoop IGMP information
- Identify commands used to configure a switch to forward IP multicast traffic
- Configure and verify multicast forwarding in an MLS network

IP Multicast

This topic describes IP multicast operations.



Multimedia applications offer the integration of sound, graphics, animation, text, and video, which are all delivered over IP and an existing campus infrastructure. These types of applications have become an effective means of corporate communication; however, sending combined media over a campus data network brings with it the potential for high bandwidth consumption. IP multicast is an efficient means of delivering media to many hosts over a single IP flow.

Multimedia traffic can work its way through the network in one of several ways.

Unicast: An application sends one copy of each packet to every client unicast address requiring the multimedia flow. If the group is large, the same information is carried multiple times from sender to receiver. Unicast transmission has significant scaling restrictions.

Broadcast: An application sends one copy of each packet to a destination broadcast address. The network interface card of all hosts seeing the broadcast packets must process all broadcast packets. This is inefficient if only a small group in the network actually needs the data flow. Multimedia broadcast is rarely implemented.

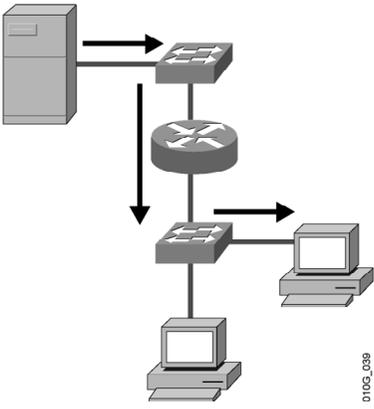
Multicast: A multimedia server sends one copy of each packet to a single destination IP address that can be received by many end stations if they choose to “listen” on that address. In the figure, the video server transmits a single video stream to a set of host devices listening to a specific multicast address. Only 1.5 Mbps of server-to-network bandwidth is utilized regardless of the number of receiving hosts.

Multicasting conserves bandwidth by replicating packets only onto segments where listening devices exist, allowing an arbitrary number of clients to subscribe to the multicast address. Multicast flows minimize processing required by network devices and nonlistening hosts.

IP Multicast Characteristics

Cisco.com

- **Transmits to a host group**
- **Delivers with “best effort” reliability**
- **Supports dynamic membership**
- **Supports diverse numbers and locations**
- **Supports membership in more than one group**



010C_039

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—8-4

IP multicast is the transmission of an IP data frame to a host group that is defined by a single IP multicast address. IP multicasting has these characteristics:

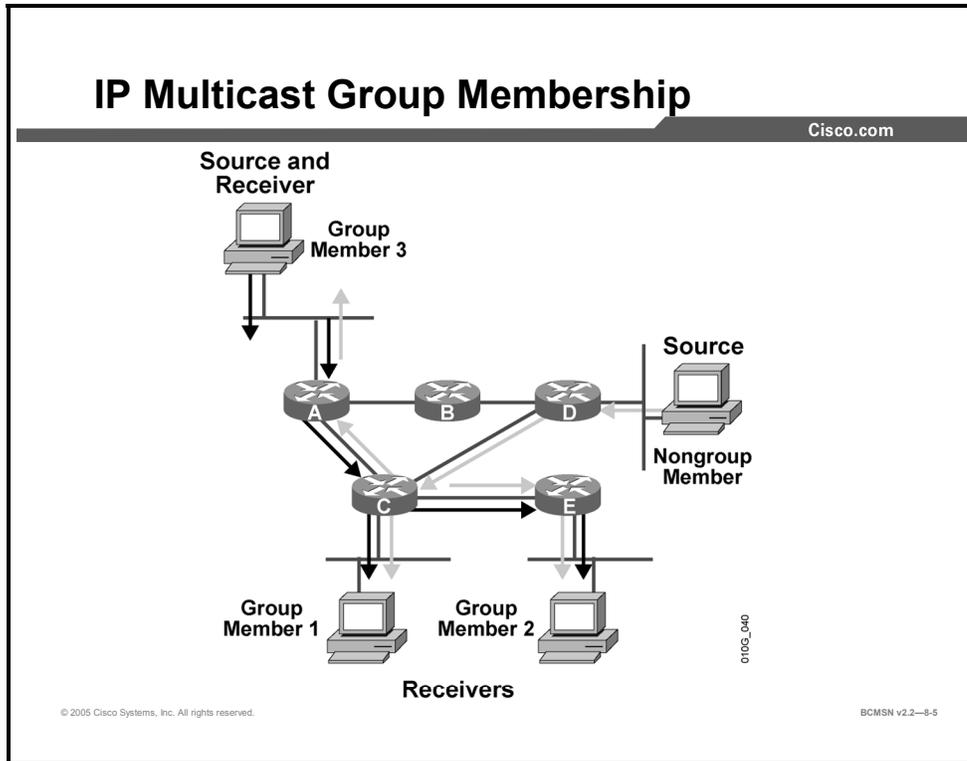
- Delivers a multicast datagram to a destination multicast address (also known as a multicast group) with the same best-effort reliability as a regular unicast IP datagram.
- Allows group members to join and leave dynamically.
- Supports all host groups regardless of the location or number of members.
- Supports the membership of a single host in one or more multicast groups.
- Can carry multiple data streams to a single group address.
- Can use a single group address for multiple host applications.
- Multicast server does not keep track of the number of recipients.

In IP multicasting, the variability in delivery time is limited to the differences in end-to-end network delay along the complete server-to-client path. In a unicast scenario, the server must sequence transmission of multiple copies of the data to multiple unicast hosts, so variability in delivery time is large, especially for large transmissions or large distribution lists.

Multicast traffic is handled at the transport layer using the User Datagram Protocol (UDP). Unlike the Transmission Control Protocol (TCP), UDP adds no reliability, flow control, or error recovery functions to IP. Because of the simplicity of UDP, data packet headers contain fewer bytes and consume less network overhead than TCP. Therefore, reliability in multicast is managed by observing receivers and monitoring the network's ability to deliver the multicast packets in sequence without delay.

IP Multicast Group Membership

This subtopic describes IP multicast group membership.



IP multicast relies on the concept of group members and a group address. The group address is a single IP multicast address that is the destination address of all packets sent from a source. Receiving devices join that group and listen for packets with the destination IP address of the group. In normal TCP/IP operations, packets arrive at a destination address by traversing routers that forward the packets hop by hop, based on the IP destination of the packet and entries in the routing table. Routers do not forward traffic in this manner to multicast destination addresses.

By default, multicast traffic is blocked at the Layer 3 devices, as is the case with broadcast traffic. Routers between a multicast source and its receiving hosts must be configured with a multicast protocol that will determine where the receivers exist so that the flow is sent only onto segments with receivers and is not replicated onto segments with no group members listening for that flow. Receivers join a group by registering with a local router. The routers between the source and receiver must be configured to deliver the multicast stream from the source to all segments hosting devices that have joined the group. Once a receiver joins a group, packets addressed to that group are placed onto the segment where the receivers are located, and multiple group members can receive a single flow. PIM and IGMP are multicast protocols that keep multicast traffic confined to portions of the network where group members are located.

A single device can be a source for one flow and be a group member or receiver for a second flow. In the figure, two flows are arriving at group members 1 and 2. One flow comes from the device on the top, which is a source for the stream. This is represented by the dark arrows. A second flow from the device on the right is represented by the pale arrows. For the first flow, the device on the top of the figure is a source, but for the second flow, it is a group member listening to the flow from the device on the right. Group members 1 and 2 are listeners or group members for both flow and therefore send no multicast packets.

IP Multicast Address Structure

This topic describes IP Multicast address structure.

Multicast IP Address Structure

Cisco.com

28 bits

Class D	1	1	1	0	Multicast Group ID
---------	---	---	---	---	--------------------

0100_041

- **A Class D address consists of 1110 as the high-order bits in the first octet, followed by a 28-bit group address.**
- **Class D addresses range from 224.0.0.0 through 239.255.255.255. The high-order bits in the first octet identify this 224-base address.**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—8-6

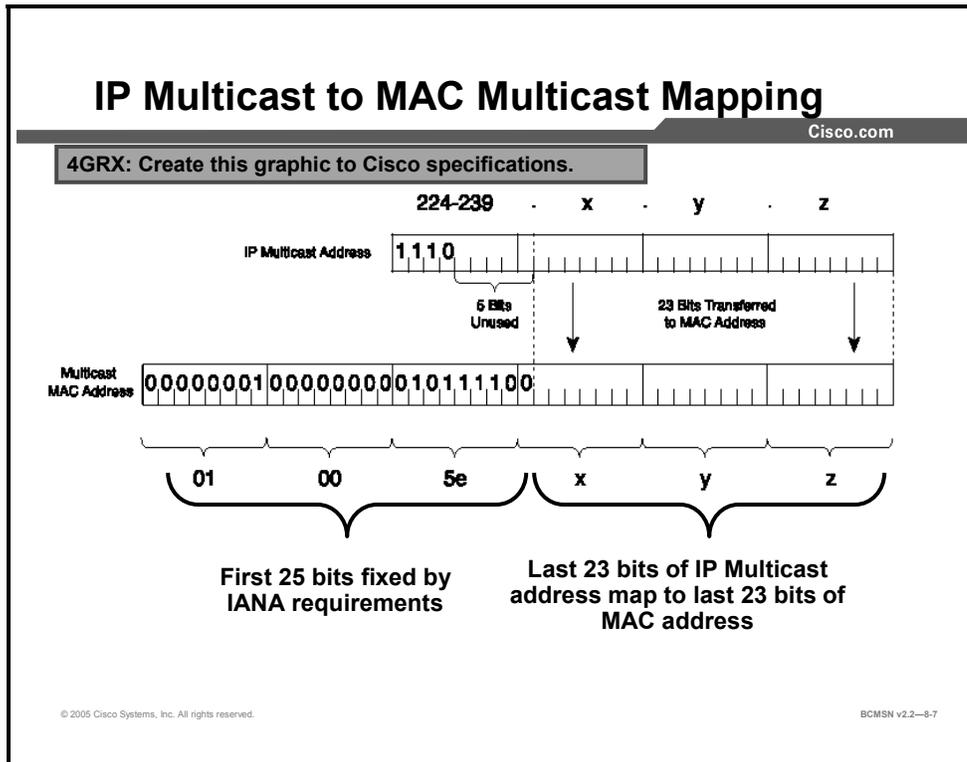
Multicast uses Class D IP address space. A Class D address consists of 1110 as the high-order bits in the first octet, followed by a 28-bit group address. The range of IP multicast addresses is divided into classes based on the high-order bits of a 32-bit IP address.

The remaining 28 bits of the IP address identify the multicast group ID. This multicast group ID is a single address typically written as decimal numbers in the range 224.0.0.0 through 239.255.255.255. The high-order bits in the first octet identify this 224-base address.

Multicast addresses may be dynamically or statically allocated. Dynamic multicast addressing provides applications with a group address on demand. Dynamic multicast addresses have a specific lifetime, and applications must request and use the address as needed. Static addresses are used at all times. As with IP addressing, there is the concept of private address space that may be used for local, organization-wide traffic and public or Internet-wide multicast addresses. There are also addresses reserved for specific protocols that require well-known addresses. The Internet Assigned Numbers Authority (IANA) manages the assignment of multicast addresses that are called permanent host groups and are similar in concept to well-known TCP and UDP port numbers.

IP Multicast to MAC Address Mapping

This subtopic describes how an IP multicast address and a MAC address correspond to one another.



Due to decisions taken early in the development of multicasting, only the MAC address range from 0100.5e00.0000 through 0100.5e7f.ffff is available for carrying multicast frames.

This makes the first 25 bits of the MAC address fixed and allows for the last 23 bits of the MAC address to correspond to the last 23 bits in the IP multicast group address.

Because the upper 5 bits of the IP multicast address are dropped in this mapping, the result is that two different IP multicast addresses may map to the same MAC multicast address. For example, 224.1.1.1 and 225.1.1.1 map to the same multicast MAC address. If one user subscribed to group A (as designated by 224.1.1.1) and the other users subscribed to group B (as designated by 225.1.1.1), they would both receive both A and B streams at Layer 2. At Layer 3, however, only the packets associated with the IP address of the selected multicast group would be viewable because the port ranges used within the address will be different between aliased streams. Network administrators should consider this when assigning IP multicast addresses.

IP Multicast Address Ranges

This topic describes how multicast address space is grouped.

IP Multicast Addresses	
Cisco.com	
Description	Range
Reserved link local address	224.0.0.0 to 224.0.0.255
Globally scoped addresses	224.0.1.0 to 238.255.255.255
Source specific multicast	232.0.0.0 to 232.255.255.255
GLOP addresses	233.0.0.0 to 233.255.255.255
Limited scope addresses	239.0.0.0 to 239.255.255.255

© 2005 Cisco Systems, Inc. All rights reserved. BCM5N v2.2-8-8

The IANA controls the assignment of IP multicast addresses. IANA has assigned the IPv4 Class D address space to be used for IP multicast. Therefore, all IP multicast group addresses fall in the range from 224.0.0.0 through 239.255.255.255. The Class D address is used for the destination IP address of multicast traffic for a specific group. The source address of a multicast datagram is the unicast address of the device sourcing the multicast flow to the destination multicast address.

Reserved Link-Local Addresses

The IANA has reserved addresses in the range 224.0.0.0 to 224.0.0.255 to be used by network protocols on a local network segment. A router should never forward packets with these addresses. Packets with link-local destination addresses are typically sent with a Time to Live (TTL) value of 1 and are not forwarded by a router. Network protocols use these addresses for automatic router discovery and to communicate important routing information. For example, Open Shortest Path First (OSPF) protocol uses the IP addresses 224.0.0.5 and 224.0.0.6 to exchange link-state information.

Address 224.0.0.1 identifies the all-hosts group. Every multicast-capable host must join this group. If a **ping** command is issued using this address, all multicast-capable hosts on the network must answer the ping request.

Address 224.0.0.2 identifies the all-routers group. Multicast routers must join that group on all multicast-capable interfaces.

Globally Scoped Addresses

Multicast addresses in the range from 224.0.1.0 through 238.255.255.255 are called “globally scoped addresses.” They can be used to multicast data between organizations and across the Internet. Some of these addresses have been registered with IANA; for example IP address 224.0.1.1 has been reserved for Network Time Protocol (NTP).

Source-Specific Multicast Addresses

Addresses in the 232.0.0.0 to 232.255.255.255 range are reserved for Source-Specific Multicast (SSM). SSM is an extension of Protocol Independent Multicast (PIM), which allows for an efficient data delivery mechanism in one-to-many communications.

GLOP Addresses

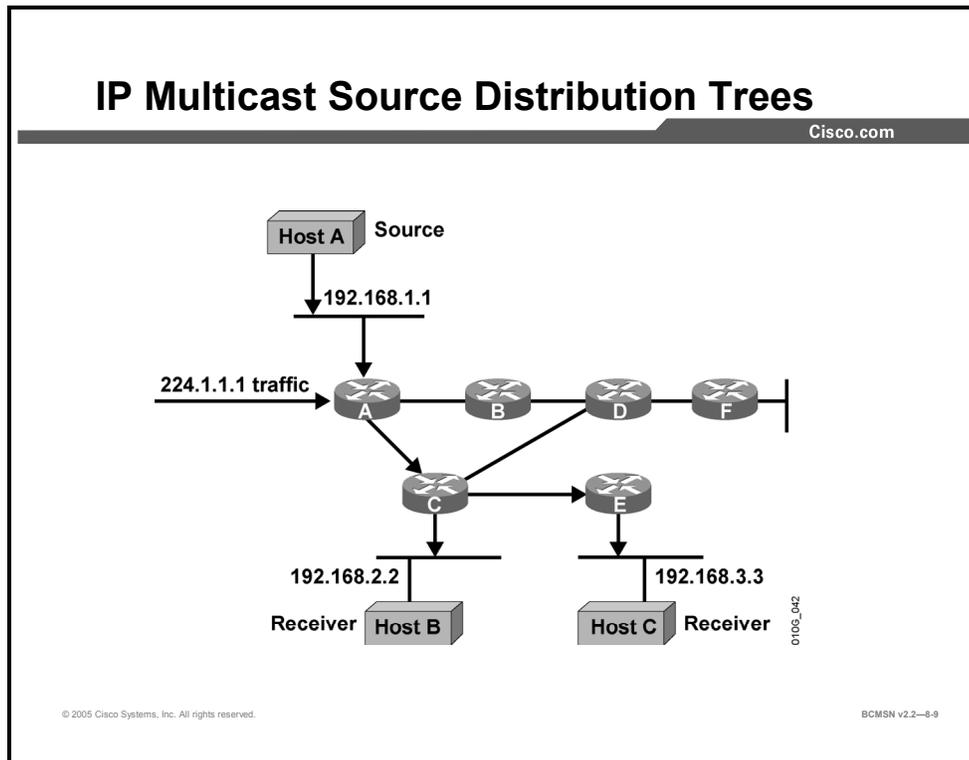
Multicast GLOP (a word, not an acronym) addresses in the range 233.0.0.0 to 233.255.255.255 can be used statically by organizations that have an Autonomous System (AS) number registered by a network registry and listed in the RWhois database. The second and third octets of the domain multicast address are represented by the AS number. For example, AS 62010 is written in hexadecimal format as “F23A.” This value is separated into two parts, F2 and 3A, and those numbers, converted to decimal, would be 242 and 58. This would yield a multicast GLOP address of 233.242.58.0/24. Multicast group addresses in that address space can be used by the organization with AS 62010 and routed throughout the Internet Multicast Backbone.

Limited Scope Addresses

Like private IP address space that is used within the boundaries of a single organization, “limited” or “administratively scoped” addresses in the range 239.0.0.0 to 239.255.255.255 are constrained to a local group or organization. Companies, universities, or other organizations can use limited scope addresses to have local multicast applications that will not be forwarded over the Internet. Typically, routers are configured with filters to prevent multicast traffic in this address range from flowing outside of an AS. Within an autonomous system or domain, the limited scope address range can be further subdivided so that local multicast boundaries can be defined. This subdivision is called “address scoping” and allows for address reuse between smaller domains. These addresses are described in RFC 2365, *Administratively Scoped IP Multicast*.

What Is RPF?

This topic describes how IP multicast traffic traverses distribution trees.



Multicast-capable routers create distribution trees that control the path that IP multicast traffic takes through the network. Unlike typical IP traffic, multicast traffic is forwarded away from the source, rather than toward the receiver. Because it is the reverse of typical IP packet processing, the process is called Reverse Path Forwarding (RPF). Multicast-capable routers create distribution trees that control the path that IP multicast traffic takes through the network as it is forwarded from the source toward all receivers. Multicast distribution trees fall into the categories of source-based trees and shared trees and the type of tree is dependent upon the multicast protocol in use.

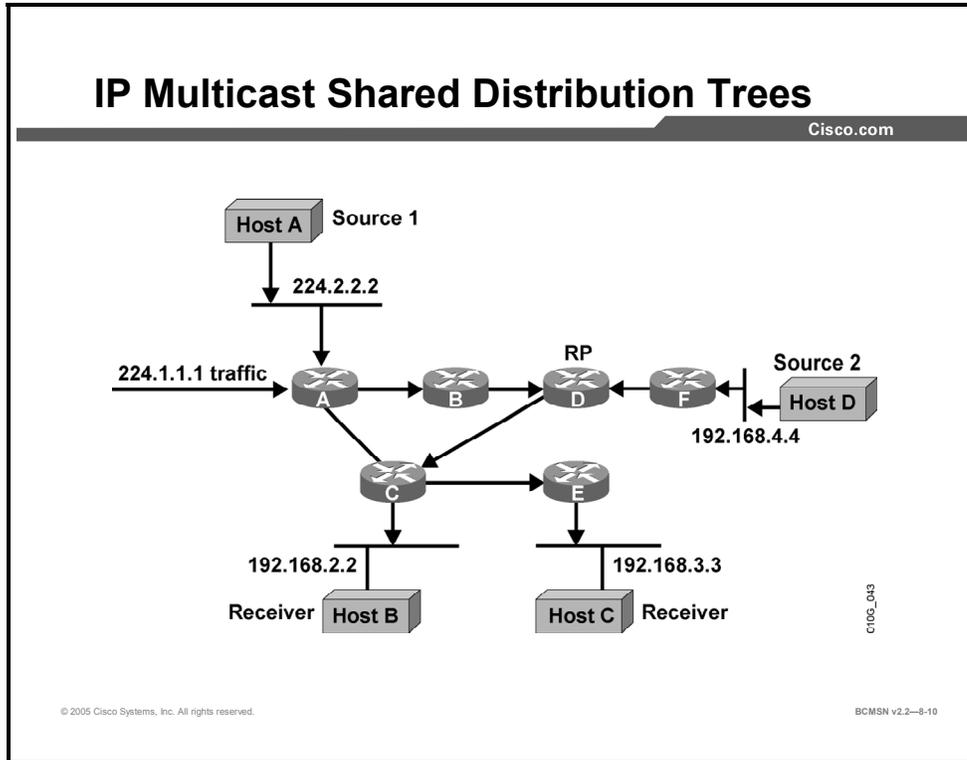
Source Distribution Trees

A source tree is the simplest form of a multicast distribution tree, with its root at the source and branches forming a tree through the network toward the receivers. This type of tree uses the shortest path through the network and is therefore also called a “shortest path tree” (SPT).

The figure shows an example of an SPT for group 224.1.1.1 rooted at the source (host A) and distributing to two receivers (hosts B and C).

An SPT is identified by a special notation of (S, G), pronounced “S comma G,” where S is the IP address of the source, and G is the multicast group address to which receivers belong. Using this notation, the SPT for the example shown in the figure would be (192.168.1.1, 224.1.1.1). The unicast IP address of the receivers is irrelevant.

The (S, G) notation implies that a separate SPT exists for each individual source sending to each group. For example, if host B is also sending traffic to group 224.1.1.1 and hosts A and C are receivers, a separate (S, G) SPT would exist with a notation of (192.168.2.2, 224.1.1.1).



Shared Distribution Trees

Unlike source trees that have their root at the source, shared trees use a single common root placed at a chosen point in the network. This shared root is called a “rendezvous point” (RP). The figure shows a shared unidirectional tree for group 224.2.2.2, with the root located at router D. The source traffic is sent toward the RP. The traffic is then forwarded from the RP to reach all the receivers. If the receiver is located between the source and the RP, it will be serviced directly.

In this example, multicast traffic from the sources (hosts A and D) travels to the RP (router D) and then down the tree to the two receivers (hosts B and C). Because all sources in the multicast group use a common shared tree, a notation written as (*, G), pronounced “star comma G,” represents the tree. In this case, “*” means all sources, and G represents the multicast group. Therefore, the shared tree shown in the figure would be written as (*, 224.2.2.2).

Source Trees Versus Shared Trees

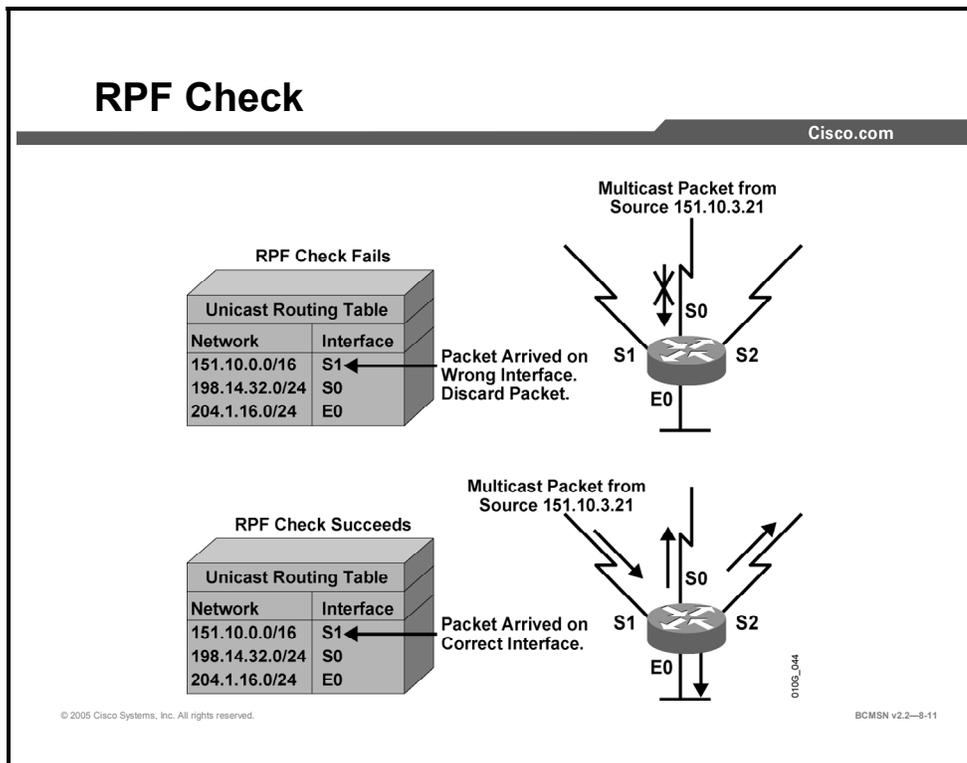
Members of multicast groups can join or leave at any time; therefore, the distribution trees must be dynamically updated. When all the active receivers on the Layer 3 segments that are associated with a particular branch stop requesting traffic for a multicast group, the routers prune that branch from the distribution tree, stopping traffic flow to that branch. If one receiver on that branch becomes active and requests the multicast traffic, the router will dynamically modify the distribution tree and will again start forwarding traffic.

Source trees have the advantage of creating the optimal path between the source and the receivers. This advantage guarantees the minimum amount of network latency for forwarding multicast traffic. However, this optimization comes at a cost: the routers must maintain path information for each source. In a network that has thousands of sources and thousands of groups, this overhead can quickly become a resource issue on the routers. Memory consumption from the size of the multicast routing table is a factor that network designers must take into consideration.

Shared trees have the advantage of requiring the minimum amount of network state information in each router. This advantage lowers the overall memory requirements for a network that allows shared trees only. The disadvantage of shared trees is that, under certain circumstances, the paths between the source and receivers might not be the optimal path, which might introduce some latency in packet delivery. For example, in the figure, the shortest path between host A (source 1) and host B (a receiver) would be router A and router C. Because router D is used as the RP for a shared tree, the traffic must traverse routers A, B, D, and then C. Network designers must carefully consider the placement of the RP when implementing a shared tree-only environment.

Reverse Path Forwarding Check

This subtopic describes the process of RFP check in multicast routing.



With unicast traffic, packets are routed from the source to the destination by considering the packet's destination IP address. The routing table looks up the destination address and forwards a single copy of the unicast packet out the correct interface in the direction of the destination.

In multicast forwarding, the source sends traffic to a group of hosts represented by a multicast group address. The multicast router determines which direction is upstream (toward the source) and which is downstream (toward the listeners). If there are multiple downstream paths, the router replicates the packet down all appropriate downstream paths.

RPF uses the existing unicast routing table to validate the network from where upstream multicast traffic should arrive. When a multicast packet arrives at a router, the router will perform an RPF check on the packet. If the check is successful, the packet is forwarded; otherwise it will be dropped. This RPF check helps to guarantee that the distribution tree is loop free.

For packets flowing down a source tree, the RPF check procedure follows this sequence:

- Step 1** Router looks up the source address in the unicast routing table to determine if the packet has arrived on the interface located on the reverse path back to the source.
- Step 2** If a packet has arrived on the interface leading back to the source, the RPF check is successful and the packet will be forwarded.
- Step 3** If the RPF check in Step 2 fails, the packet is quietly dropped.

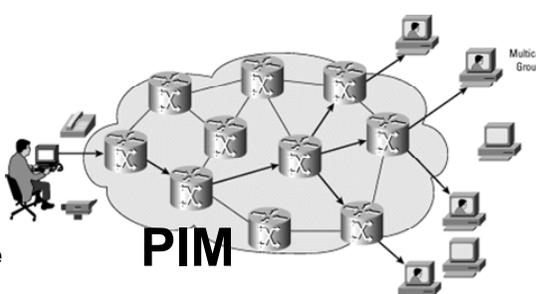
At the top of the figure, the RPF check fails. A multicast packet from source 151.10.3.21 is received on interface S0. A check of the unicast route table shows that S1 is the interface where

What Is PIM?

This topic describes Protocol Independent Multicast.

What is PIM?

Cisco.com



- Configured on all routers between source and receiver
- Conserves network resources while making flow available to appropriate receivers
- Traffic flow dependent upon PIM implementation

4GRX: Create this graphic to Cisco specifications. Found on link:
http://cisco.com/en/US/tech/tk828/tech_brief09186a00800a4415.html

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—8-12

PIM is a multicast routing protocol that makes packet-forwarding decisions independent of standard or unicast IP routing protocols such as Enhanced Interior Gateway Routing Protocol (EIGRP). At the same time, PIM can leverage unicast routing tables to perform multicast forwarding functions. For example, PIM uses the unicast routing table to perform the RPF check function. Unlike unicast routing protocols, PIM does not send and receive routing updates between routers.

PIM has three forwarding modes.

PIM Dense Mode (PIM-DM): This mode uses a push model to flood multicast traffic to every corner of the network. Routers located throughout the tree can then prune the flow if they receive no requests for the particular multicast group flow. This method would be efficient in certain deployments in which there are typically active receivers on every subnet in the network.

PIM Sparse Mode (PIM-SM): This mode uses a pull model to deliver multicast traffic. Only network segments with active receivers explicitly requesting the flow will receive the traffic for a multicast group. PIM-SM distributes information about active sources by forwarding data packets on the shared tree. Because PIM-SM uses shared trees, the source forwards its flow to a rendezvous point (RP), which then sends it to any requesting routers.

PIM Sparse-Dense Mode: This mode allows individual groups to be run in either sparse or dense mode, depending on whether RP information is available for that group. If the router gleans RP information for a particular group, it will be treated as sparse mode; otherwise that group will be treated as dense mode.

PIM Versions 1 and 2

PIM version 1 was Cisco proprietary. In addition to being an IEEE standard, version 2 includes the following improvements:

- A single, active RP exists per multicast group, with multiple backup RPs. This compares to multiple active RPs for the same group in PIM version 1.
- A bootstrap router (BSR) provides a fault tolerant, automated RP discovery and distribution mechanism. Thus, routers dynamically learn the group-to-RP mappings.
- Sparse mode and dense mode are properties of a group, as opposed to an interface. PIM Join and Prune messages have more flexible encodings for multiple address families.
- Registered messages to an RP indicate whether they were sent by a border router or by a designated router.
- PIM packets are no longer inside IGMP packets; they are stand-alone packets.

References

For additional information, refer to this resource:

IP Multicast: http://cisco.com/en/US/tech/tk828/tech_brief09186a00800a4415.html

PIM Modes

This topic describes PIM modes.

Protocol-Independent Multicast Dense Mode

Cisco.com

Prune Message

- **Protocol-independent means that the protocol is not dependent on a specific unicast routing protocol.**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—8-13

PIM is IP routing protocol independent and can leverage whichever unicast routing protocols are used to populate the unicast routing table. PIM dense mode operation is based on the assumption that the multicast group members are densely distributed throughout the network and that bandwidth is plentiful, meaning that almost all hosts on the network belong to the group. PIM dense mode multicast routing protocol relies on periodic flooding of the network with multicast traffic to set up and maintain the distribution tree.

PIM dense mode works best when there are numerous members belonging to each multimedia group. PIM floods the multimedia packet out to all routers in the network and then prunes routers that do not support members of that particular multicast group.

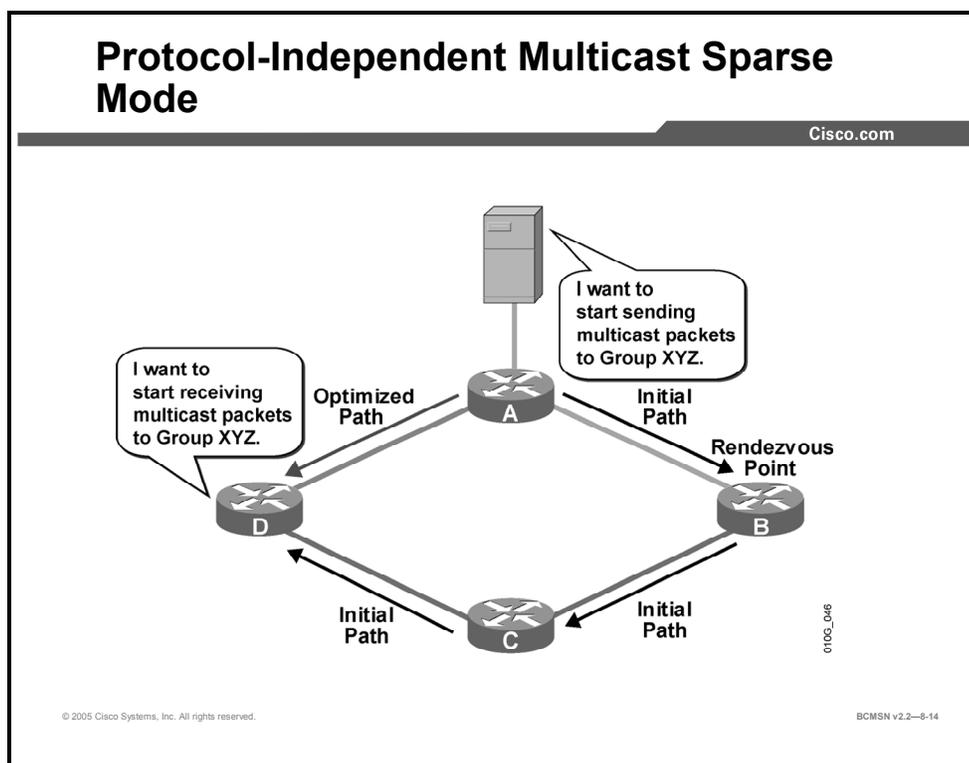
PIM dense mode is most useful under the following circumstances:

- Senders and receivers are in close proximity to one another.
- There are few senders and many receivers.
- The volume of multicast traffic is high.

The stream of multicast traffic is constant.

PIM Sparse Mode

This subtopic describes PIM sparse mode operations.



The second approach to multicast routing is based on the assumption that the multicast group members are sparsely distributed throughout the network and bandwidth is not necessarily widely available.

It is important to note that sparse mode does not imply that the group has few members, just that they are widely dispersed. In this case, flooding would unnecessarily waste network resources. Sparse mode multicast routing protocols rely on more selective techniques to set up and maintain multicast trees. Sparse mode protocols begin with an empty distribution tree and add branches only as the result of explicit requests to join the distribution.

Sparse mode PIM is optimized for environments where there are many multipoint data streams. Sparse multicast is most useful when

- There are few receivers in a group.
- The type of traffic is intermittent.

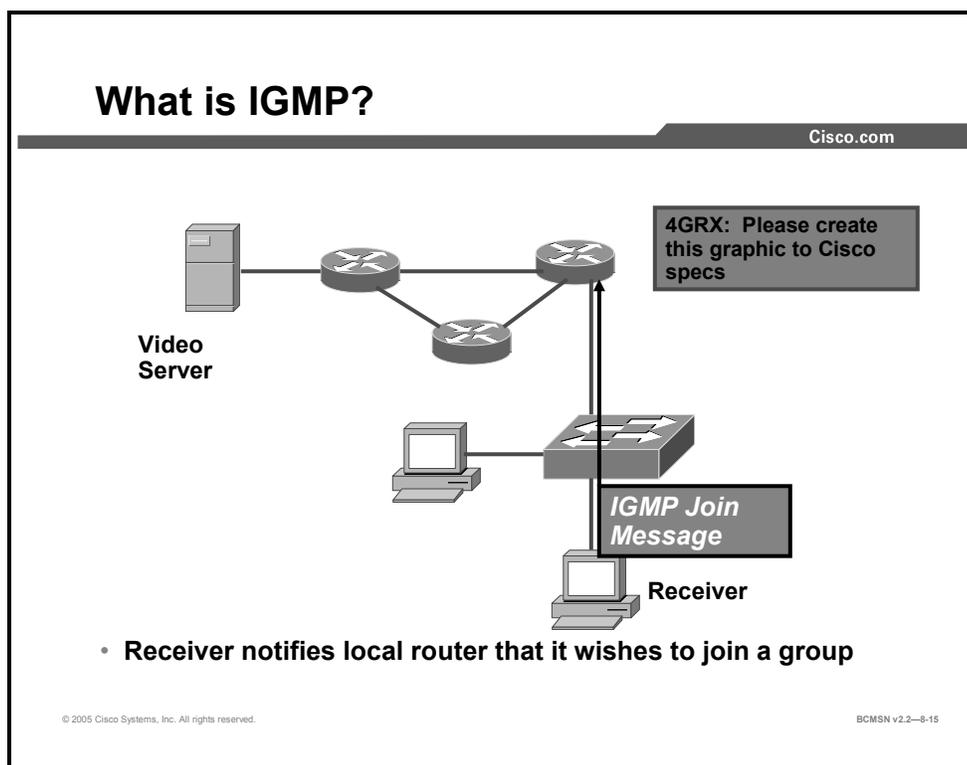
In sparse mode, each data stream goes to a relatively small number of segments in the campus network. Instead of flooding the network to determine the status of multicast members, sparse mode PIM defines a rendezvous point. When a source begins to generate a flow, it is directed to a rendezvous point. When a router determines that it has receivers out its interfaces, it registers with the rendezvous point. The routers in the path will optimize the path automatically to remove any unnecessary hops. Sparse mode PIM assumes that no hosts want the multicast traffic unless they specifically request it.

PIM is able to simultaneously support dense mode for some multicast groups and sparse mode for others. Cisco has implemented an alternative to choosing just dense mode or just sparse

mode on a router interface. PIM sparse-dense mode allows the network to determine which IP multicast groups should use sparse mode and which groups should use dense mode. PIM sparse mode and sparse-dense mode require the use of a rendezvous point.

What Is IGMP?

This topic discusses the IGMP multicast protocols that manage client joins and leaves.



IGMP is used to register individual hosts with a multicast group. The host sends a join message to a local router multicast address. If the router is running a multicast routing protocol, it will accept the join and then forward the multicast stream for that group onto the segment where the registering host is present. IGMP messages are IP datagrams with a protocol value of 2 and a destination address 224.0.0.2 and a TTL of 1.

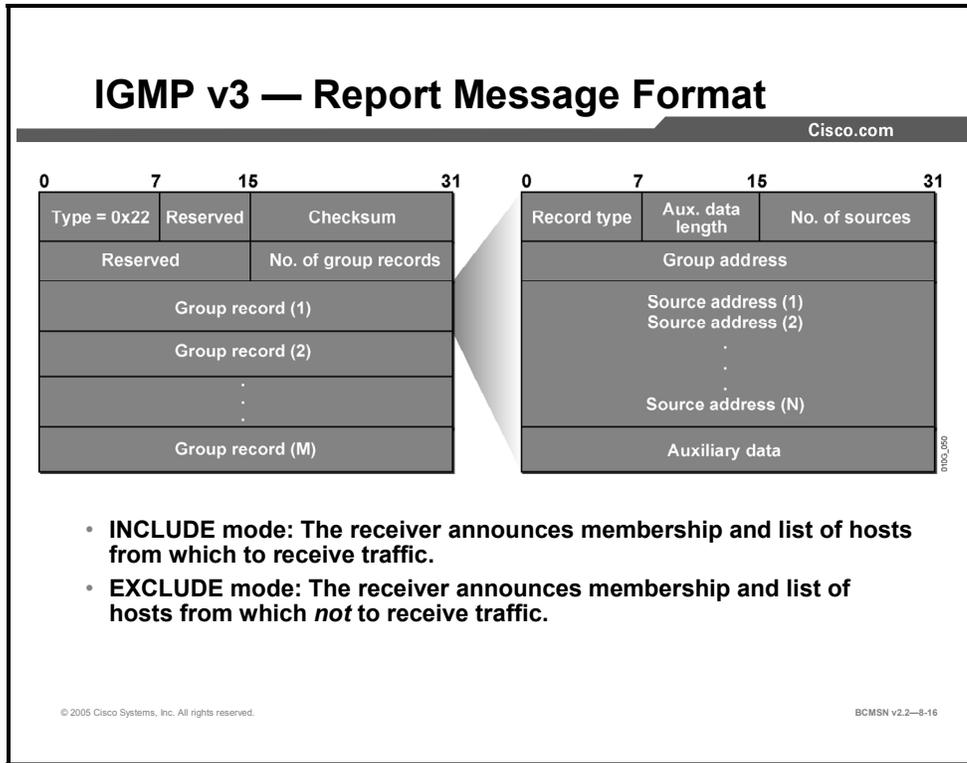
In addition to listening to IGMP join messages, multicast routers also periodically send out queries to discover which groups are active or inactive on a particular subnet. Any end station that is part of the multicast group receives this IGMP query and responds with a host membership report for each group to which it belongs. This is sent to all hosts 224.0.0.1 with a TTL of 1.

As of this writing, version 3 is the most current iteration of IGMP and is covered in more detail. Previous versions had attributes and limitations as listed in the table.

IGMP Version	Attributes and Limitations
v1	Single group membership per IGMP message. No group-specific query, no leave group message.
v2	Single group membership per IGMP message. Could not specify source host from which the multicast flow is desired.
v3lite	Cisco proprietary for SSM applications on hosts without v3 support.

IGMP Message Format

This subtopic describes the content of the IGMP frame.



IGMP version 3, the next step in the evolution of IGMP, adds support for source filtering, multiple group memberships, joins, and leaves. This enables a multicast receiving host to indicate to the router the groups from which it wants to receive multicast traffic, as well as the source unicast addresses from which this traffic is expected. This membership information enables Cisco IOS software to forward traffic from only those sources requested by the receiver. IGMP v 3 supports report and query messages that have different packet structures as shown.

IGMP v 3 Report Message

This table describes the fields in the IGMP v 3 report message sent from the host to the router.

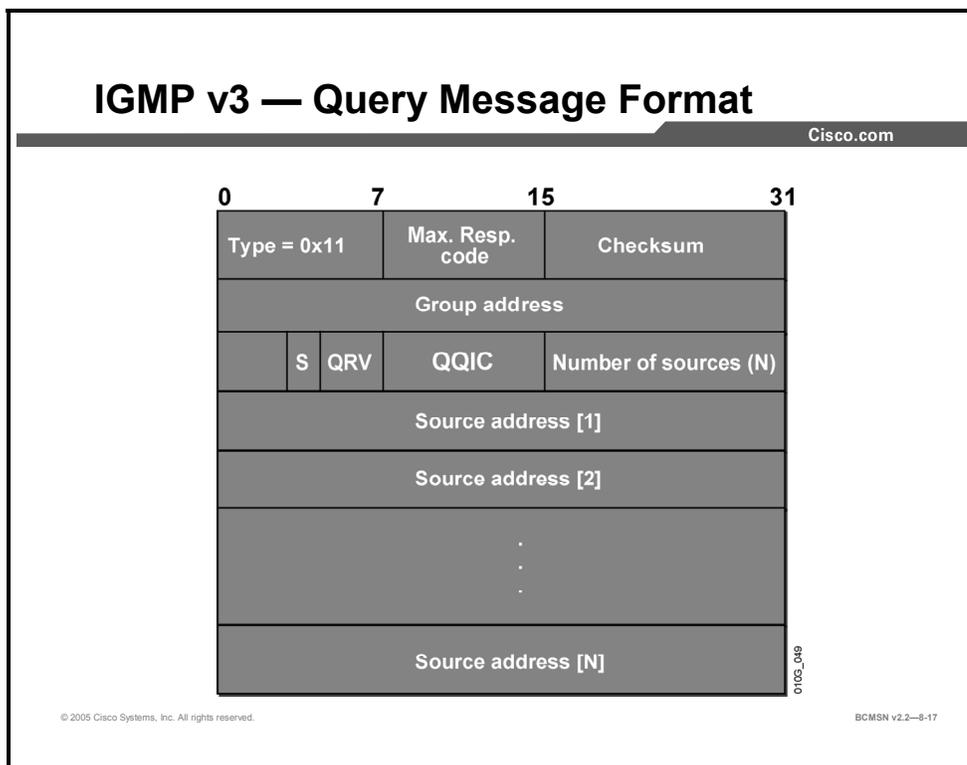
Field	Definition
No. of group records [M]	Number of group records present in the report
Group record [1...M]	Block of fields containing information regarding the sender membership with a single multicast group
Record type	The group record type (e.g., INCLUDE, EXCLUDE)
No. of sources [N]	Number of sources present in the record
Source address [1...N]	Address of the source(s)

With IGMP v 3, receivers signal membership to a multicast host group in these two modes:

- **INCLUDE mode:** The receiver announces membership to a host group and provides a list of source addresses (the INCLUDE list) from which it does want to receive traffic.

IGMP v 3 Query Message

This subtopic describes the IGMP v 3 query message.



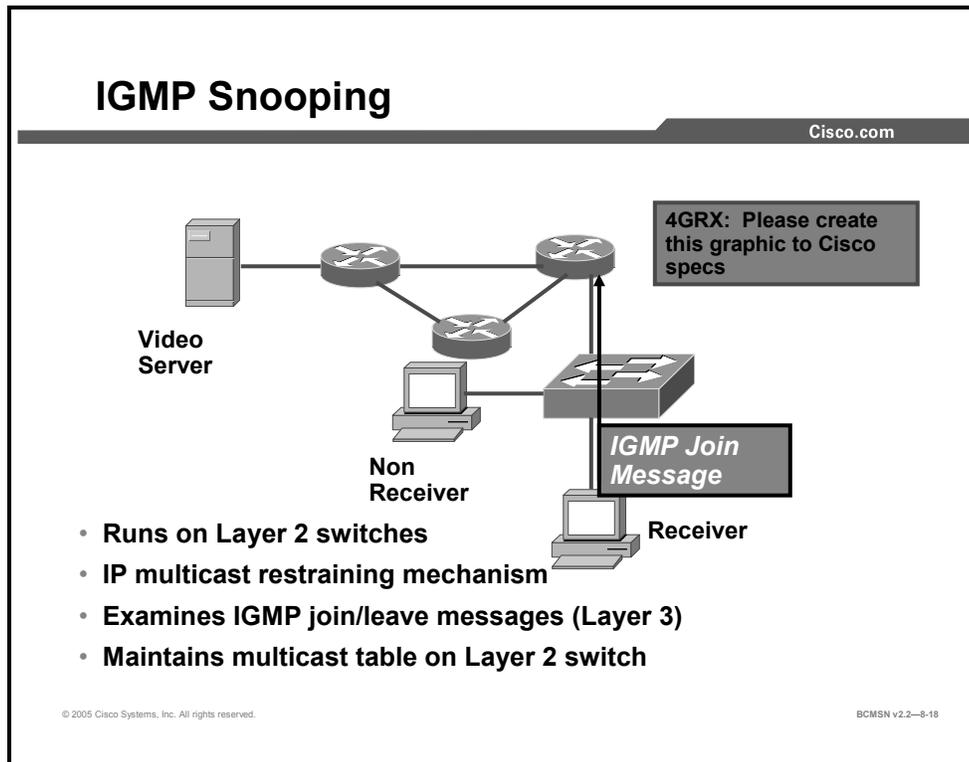
The IGMP query message sent from the multicast router to the all-hosts address 224.0.0.1 has a different format than the report or join messages have.

Fields in the IGMP v 3 Query Message

Field	Definition
Type = 0x11	IGMP query.
Max resp. code	Maximum response code (in seconds). This field specifies the maximum time allowed before sending a responding report.
Group address	Multicast group address. This address is 0.0.0.0 for general queries.
S	S flag. This flag indicates that processing by routers is being suppressed.
QRV	Querier Robustness Value. This value affects timers and the number of retries.
QQIC	Query Interval Code of the Querier (in seconds). This field specifies the query interval used by the querier.
No. of sources [N]	Number of sources present in the query. This number is nonzero for a group-and-source query.
Source address [1...N]	Address of the source(s)

Describing the IGMP Snooping Process

This topic describes how a switch can snoop IGMP messages.



The default behavior for a Layer 2 switch is to forward multicast traffic to every port in the VLAN on which the traffic was received. Therefore, a switch between a requesting host and a multicast router will forward a multicast flow intended for a single host out all switch ports on the same VLAN as the receiving host. IGMP snooping is an IP multicast constraining mechanism for switches. It examines IGMP frames so that multicast traffic is not forwarded out all VLAN ports but only those over which hosts sent IGMP messages toward the router.

IGMP snooping runs on a Layer 2 switch. The switch snoops the content of the IGMP join and leave messages sent between the hosts and the router. When the switch sees an IGMP report from a host to join a particular multicast group, the switch creates a content addressable memory (CAM) table entry associating the port number where that message was seen to the Layer 2 multicast address for the group that the host joined. When the frames of the multicast flow arrive at the switch with the destination multicast MAC address, they are forwarded down only those ports where the IGMP messages were snooped and associated CAM table entries were created. When the switch snoops the IGMP leave group message from a host, the switch removes the table entry.

IP Multicast Configuration Commands

This topic describes the commands used to configure and verify IP multicast.

IP Multicast Configuration Commands

Cisco.com

Configure

- **ip multicast-routing**
- **ip pim sparse-dense mode**
- **ip pim rp-address 192.168.68.3**

Verify

- **show ip pim interface**
- **show ip mroute**

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2—8-19

These commands are used to configure IP multicast on a router or switch.

Multicast Configuration and Verification Commands

Command	Description
Switch(config)# ip multicast-routing	Enables multicast routing on the Layer 3 device.
Switch(config-if)# ip pim pim-mode	Specifies if PIM will operate in sparse, dense or sparse-dense mode as it communicates with other multicast routers on this interface.
Switch(config)# ip pim rp-address ip-address	Specifies the address of the RP that this router will communicate with for PIM sparse mode operation.
Switch# show ip pim interface	Displays information about PIM on an interface basis. Numerous arguments are available to display specific information.
Switch# show ip mroute [hostname / group_number]	Displays multicast routing table entries on a source, group basis. Numerous arguments are available to display specific information.

How to Enable IP Multicast

This topic describes how to execute commands to enable and verify multicast routing.

Enabling IP Multicast

Cisco.com

```
Switch(config)#ip multicast-routing
```

- Globally enables IP multicast routing

```
Switch(config-if)#ip pim [sparse-mode | dense-mode | sparse-dense-mode]
```

- Configures PIM on a specific interface

```
Switch(config)#ip rp-address ip-address [acl-number] [override]
```

- Configure the IP address of the RP

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—8-20

By default, a Layer 3 device will isolate multicast traffic to the segment on which it was generated, not forwarding it across the router to other network segments. Enabling IP multicast routing allows a Layer 3 device to forward multicast packets based upon the configuration of the multicast routing protocol. The general steps to enabling and verifying multicast routing are outlined here.

1. Enable IP Multicast Routing

A single command is used from global configuration to enable multicast routing:

```
Switch(config)#ip multicast-routing
```

2. Enable a Multicast Routing Protocol

The multicast routing protocol, which is PIM on a campus network, establishes the rules by which multicast traffic will be forwarded onto various network segments by the Layer 3 device. An interface can be configured to operate in PIM dense mode, sparse mode, or sparse-dense mode. The mode determines how the Layer 3 device populates its multicast routing table and how it forwards multicast packets received from directly connected segments. Enabling PIM on an interface also enables IGMP operation on that interface.

At interface configuration mode, configure the PIM mode of operation for the interface:

```
Switch(config-if)#ip pim dense-mode  
Switch(config-if)#ip pim sparse-mode  
Switch(config-if)#ip pim sparse-dense-mode
```

When the switch populates the multicast routing table, dense mode interfaces are always added to the table. Sparse mode interfaces are added to the table only when periodic join messages are received from downstream routers, or when there is a directly connected member on the interface. If configured for sparse or sparse-dense mode, multicast sparse mode operation will occur if an RP is known for the group. If so, the packets are encapsulated and sent toward the RP. When no RP is known, the packet is flooded in a dense mode fashion. By default, multicast routing is disabled on an interface.

3. Configure the RP for Sparse Mode Operation

The simplest way to implement PIM sparse mode is to statically configure the RP address on each router required to forward multicast traffic. The static method should only be used on small networks as the administrative tasks of maintaining these static RP addresses soon becomes too burdensome. Routers with group members on directly connected interfaces use the RP address to send PIM register messages up the tree. On routers in the tree that are closest to the RP, the RP address is used to send PIM join and prune messages to the RP, and to inform it about group membership down the tree from that point. The RP address must be configured on all routers anywhere in the tree that will carry traffic from a source to its member (including the RP router).

To configure the address of the RP, enter this command in global configuration mode:

```
Switch(config)#ip pim rp-address ip-address [access-list-number]
[override]
```

4. Verify IP Multicast Operations

These commands are available to verify and monitor IP multicast operations such as operational statistics, resource utilization, multicast database content, troubleshooting information, node reachability and multicast routing paths. Some of the commands are further described in this lesson.

Multicast Routing Statistics

Command	Description
<code>ping [group-name group-address]</code>	Sends an Internet Control Message Protocol (ICMP) echo request to a multicast group address
<code>show ip mroute [hostname group_number]</code>	Displays the contents of the IP multicast routing table
<code>show ip pim interface [type number] [count]</code>	Displays information about interfaces configured for PIM
<code>show ip interface</code>	Displays PIM information for all interfaces

5. Verify PIM

This subtopic describes how to verify PIM configuration and operation on an interface.

Verifying PIM

Cisco.com

```
Switch#show ip pim interface
```

- Displays information about interfaces configured for PIM

```
Switch#show ip pim interface
```

Address	Interface	Mode	Neighbor Count	Query Interval	DR
198.92.37.6	Ethernet0	Dense	2	30	198.92.37.33
198.92.36.129	Ethernet1	Dense	2	30	198.92.36.131
10.1.37.2	FastEthernet 1/1	Dense	1	30	0.0.0.0

© 2005 Cisco Systems, Inc. All rights reserved. BCMSN v2.2--8-21

This is sample output from the **show ip pim interface** command:

```
Switch#show ip pim interface
```

```
Address          Interface          Mode      Neighbor  Query
designated router                               Count     Interval

198.92.37.6      Ethernet0          Dense     2         30
198.92.37.33
198.92.36.129    Ethernet1          Dense     2         30
198.92.36.131
```

This is sample output from the **show ip pim interface** command with a count:

```
Switch#show ip pim interface count
```

```
Address          Interface          FS  Mpackets In/Out
171.69.121.35    Ethernet0          *   548305239/13744856
171.69.121.35    Serial0.33         *   8256/67052912
198.92.12.73     Serial0.1719      *   219444/862191
```

The following is sample output from the **show ip pim interface** command with a count when IP multicast is enabled. The example lists the PIM interfaces that are fast switched and process switched, and lists the packet counts for these. **H** indicates multicast-enabled interfaces.

```
Switch#show ip pim interface count
```

```
States: FS - Fast Switched, H - Hardware Switched
Address          Interface          FS  Mpackets In/Out
```

```

192.1.10.2      Vlan10      * H 40886/0
192.1.11.2      Vlan11      * H 0/40554
192.1.12.2      Vlan12      * H 0/40554
192.1.23.2      Vlan23      * 0/0
192.1.24.2      Vlan24      * 0/0

```

6. Verifying Multicast Routing and Clearing the Routing Table

This subtopic describes how to verify IP multicast with the **show ip mroute** command.

Verifying Multicast Routing

Cisco.com

```
Switch#show ip mroute [hostname / group_number]
```

- Displays the contents of the IP multicast routing table

```
Switch#show ip mroute

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
R - RP-bit set, F - Register flag, T - SPT-bit set
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.0.255.3), uptime 5:29:15, RP is 198.92.37.2, flags: SC
Incoming interface: FastEthernet 1/1, RPF neighbor 10.3.35.1, Dvmrp
Outgoing interface list:
  Ethernet0, Forward/Sparse, 5:29:15/0:02:57

(198.92.46.1, 224.0.255.3), uptime 5:29:15, expires 0:02:59, flags: C
Incoming interface: FastEthernet 1/1, RPF neighbor 10.3.35.1
Outgoing interface list:
  Ethernet0, Forward/Sparse, 5:29:15/0:02:57
```

© 2005 Cisco Systems, Inc. All rights reserved.
BCMSN v2.2-8-22

This example shows the output from several iterations of the **show ip mroute** command:

Example: show ip mroute for Sparse Mode

This is sample output from the **show ip mroute** command for a router operating in sparse mode:

```
Switch#show ip mroute

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
R - RP-bit set, F - Register flag, T - SPT-bit set
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.0.255.3), uptime 5:29:15, RP is 198.92.37.2, flags: SC
Incoming interface: FastEthernet 1/1, RPF neighbor 10.3.35.1, Dvmrp
Outgoing interface list:
  Ethernet0, Forward/Sparse, 5:29:15/0:02:57

(198.92.46.1, 224.0.255.3), uptime 5:29:15, expires 0:02:59, flags: C
Incoming interface: FastEthernet 1/1, RPF neighbor 10.3.35.1
Outgoing interface list:
  Ethernet0, Forward/Sparse, 5:29:15/0:02:57
```

Note Output interface timers are not updated for hardware-forwarded packets. Entry timers are updated approximately every 5 seconds.

Example: show ip mroute summary

This is sample output from the **show ip mroute** command with the **summary** keyword:

```
Switch#show ip mroute summary

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned
      R - RP-bit set, F - Register flag, T - SPT-bit set, J - Join
SPT
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.255.255.255), 2d16h/00:02:30, RP 171.69.10.13, flags: SJPC

(*, 224.2.127.253), 00:58:18/00:02:00, RP 171.69.10.13, flags: SJC

(*, 224.1.127.255), 00:58:21/00:02:03, RP 171.69.10.13, flags: SJC

(*, 224.2.127.254), 2d16h/00:00:00, RP 171.69.10.13, flags: SJCL
(128.9.160.67/32, 224.2.127.254), 00:02:46/00:00:12, flags: CLJT
(129.48.244.217/32, 224.2.127.254), 00:02:15/00:00:40, flags: CLJT
(130.207.8.33/32, 224.2.127.254), 00:00:25/00:02:32, flags: CLJT
(131.243.2.62/32, 224.2.127.254), 00:00:51/00:02:03, flags: CLJT
(140.173.8.3/32, 224.2.127.254), 00:00:26/00:02:33, flags: CLJT
(171.69.60.189/32, 224.2.127.254), 00:03:47/00:00:46, flags: CLJT
```

Example: show ip mroute active

This is sample output from the **show ip mroute** command with the **active** keyword:

```
Switch#show ip mroute active

Active IP Multicast Sources - sending >= 4 kbps

Group: 224.2.127.254, (sdr.cisco.com)
  Source: 146.137.28.69 (mbone.ipd.anl.gov)
  Rate: 1 pps/4 kbps(1sec), 4 kbps(last 1 secs), 4 kbps(life avg)

Group: 224.2.201.241, ACM 97
  Source: 130.129.52.160 (webcast3-e1.acm97.interop.net)
  Rate: 9 pps/93 kbps(1sec), 145 kbps(last 20 secs), 85 kbps(life
avg)

Group: 224.2.207.215, ACM 97
  Source: 130.129.52.160 (webcast3-e1.acm97.interop.net)
  Rate: 3 pps/31 kbps(1sec), 63 kbps(last 19 secs), 65 kbps(life
avg)
```

Example: show ip mroute count

This is sample output from the **show ip mroute** command with the **count** keyword:

```
Switch#show ip mroute count

IP Multicast Statistics - Group count: 8, Average sources per group:
9.87
Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second

Group: 224.255.255.255, Source count: 0, Group pkt count: 0
  RP-tree: 0/0/0/0
```

```

Group: 224.2.127.253, Source count: 0, Group pkt count: 0
  RP-tree: 0/0/0/0

Group: 224.1.127.255, Source count: 0, Group pkt count: 0
  RP-tree: 0/0/0/0

Group: 224.2.127.254, Source count: 9, Group pkt count: 14
  RP-tree: 0/0/0/0
  Source: 128.2.6.9/32, 2/0/796/0
  Source: 128.32.131.87/32, 1/0/616/0
  Source: 128.125.51.58/32, 1/0/412/0
  Source: 130.207.8.33/32, 1/0/936/0
  Source: 131.243.2.62/32, 1/0/750/0
  Source: 140.173.8.3/32, 1/0/660/0
  Source: 146.137.28.69/32, 1/0/584/0
  Source: 171.69.60.189/32, 4/0/447/0
  Source: 204.162.119.8/32, 2/0/834/0

Group: 224.0.1.40, Source count: 1, Group pkt count: 3606
  RP-tree: 0/0/0/0
  Source: 171.69.214.50/32, 3606/0/48/0, RPF Failed: 1203

Group: 224.2.201.241, Source count: 36, Group pkt count: 54152
  RP-tree: 7/0/108/0
  Source: 13.242.36.83/32, 99/0/123/0
  Source: 36.29.1.3/32, 71/0/110/0
  Source: 128.9.160.96/32, 505/1/106/0
  Source: 128.32.163.170/32, 661/1/88/0
  Source: 128.115.31.26/32, 192/0/118/0
  Source: 128.146.111.45/32, 500/0/87/0
  Source: 128.183.33.134/32, 248/0/119/0
  Source: 128.195.7.62/32, 527/0/118/0

```

Clearing the Multicast Routing Table

After configuration changes are made, the IP multicast tables may need to be cleared before accurate table information will display. This is particularly true when changes are made at several contiguous routers. Use the following EXEC mode command to clear the tables:

```
Switch#clear ip mroute
```

After executing this command, use the **show ip mroute** command to display new multicast routing table information built after the clear command was executed.

Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **IP Multicast differs significantly from standard IP traffic**
- **IP Multicast has specific address conventions in IPv4**
- **Specific IP multicast ranges have specific purposes**
- **Reverse Path Forwarding keeps the Multicast network loop-free**
- **PIM allows routing of Multicast traffic**
- **PIM has three modes of operation**
- **IGMP is used between workstations and routers**
- **IGMP messages can be snooped by switches to increase the efficiency of multicast frame forwarding on L2 switch**
- **Specific commands are used to route Multicast traffic**

© 2005 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—8-23

References

For additional information, refer to this resource:

- Cisco Systems, Inc., *IP Multicast, Internet Protocol IP Multicast Technology*.
http://cisco.com/en/US/tech/tk828/tech_brief09186a00800a4415.html

Module Summary

This topic summarizes the key points discussed in this module.

Summary

Cisco.com

- **A campus with IP telephony requires specific switch configuration parameters to give voice traffic priority processing**
- **Sending IP multicast across a campus requires Distribution and Access switches to be configured to send traffic where needed and block it where there are no consumers**

© 2004 Cisco Systems, Inc. All rights reserved.BCMSN v2.2—8-1

Campus networks must be designed to carry traffic that serves many purposes and has various impacts on network resources. Return on investment (ROI) escalates when voice, video, and data applications are delivered over a single Campus infrastructure. Proper design and configuration considerations must be made to ensure that voice, video, and data traffic are all being effectively carried over the Campus infrastructure.

References

For additional information, refer to this resource:

- Cisco Systems, Inc., *IP Multicast, Internet Protocol IP Multicast Technology*.
http://cisco.com/en/US/tech/tk828/tech_brief09186a00800a4415.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 types of multicast addresses are to be used only within a local group or organization? (Source: Configuring IP Multicast)
- A) GLOP addresses
 - B) Limited scope addresses
 - C) Globally scoped addresses
 - D) Source-specific multicast addresses
- Q2) Other than data, what types of network traffic will the converged Campus infrastructure be required to support? (Choose two.) (Source: Accommodating Voice Traffic on Campus Switches)
- A) Voice
 - B) Video
 - C) ATM
 - D) All of the above
- Q3) To separate the voice and data traffic, what is commonly used? (Source: Accommodating Voice Traffic on Campus Switches)
- A) VLAN
 - B) Data VLAN
 - C) Voice VLAN
 - D) None of the above

Module Self-Check Answer Key

Q1) B

Q2) A, B

Q3) C