



Cisco CCNP Switching Exam Certification Guide

Tim Boyles and Dave Hucaby, CCIE #4594



Cisco Press
201 W 103rd Street
Indianapolis, IN 46290

Cisco CCNP Switching Exam Certification Guide

Tim Boyles and David Hucaby

Copyright © 2001 Cisco Systems, Inc.

Cisco Press logo is a trademark of Cisco Systems, Inc.

Published by:

Cisco Press

201 West 103rd Street

Indianapolis, IN 46290 USA

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

Printed in the United States of America 1 2 3 4 5 6 7 8 9 0 03 02 01 00

1st Printing November 2000

Library of Congress Cataloging-in-Publication Number: 00-105170

ISBN: 1-58720-000-7

Warning and Disclaimer

This book is designed to provide information about the Cisco CCNP Switching Exam #640-504. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

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

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

Trademark Acknowledgments

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

Feedback Information

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

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at cisco-press@mcp.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher	John Wait
Editor-In-Chief	John Kane
Cisco Systems Program Manager	Bob Anstey
Executive Editor	Brett Bartow
Acquisitions Editor	Amy Lewis
Managing Editor	Patrick Kanouse
Development Editor	Christopher Cleveland
Copy Editor	Chuck Gose
Technical Editors	Stephen Daleo, Anthony Kwan, Chris Paggen, Casimir Sammanasu
Team Coordinator	Tammi Ross
Book Designer	Gina Rexrode
Cover Designer	Louisa Klucznik
Compositor	Octal Publishing, Inc.
Proofreaders	Dayna Isley
	Sarah Cisco
	Shannon Martin
	Larry Sweazy
Indexer	

About the Authors

Tim Boyles is the Director of Network Architecture for @Link Networks, a national CLEC which specializes in broadband data and communications solutions for small- and medium-sized businesses. Prior to that he worked as a Senior Consultant at Lucent Networkcare, formerly known as INS, where he was responsible for the design and implementation of large switch-based networks as well as multiple service provider projects. Tim has been in the networking business for 16 years with multiple vendor certifications, including CCNP. He holds an engineering undergraduate degree from the University of Missouri-Rolla and an MBA from California State University. Tim is a co-author of the *CLSC Exam Certification Guide*.

David Hucaby, CCIE #4594, is a Lead Network Engineer for the University of Kentucky, where he designs, implements, and maintains campus networks using Cisco products. Prior to his current position, David was a senior network consultant, where he provided design and implementation consulting, focusing on Cisco-based VPN and IP telephony solutions. David has a B.S. and M.S. in Electrical Engineering from the University of Kentucky.

About the Technical Reviewers

Stephen Daleo, CCNP, is a Certified Cisco Systems Instructor (CCSI) and a consultant with Mentor Technologies (formerly Chesapeake Computer Consultants, Inc.). Stephen has been teaching the recommended courses for Cisco Career Certifications since 1996, including ICND, BSCN, BCMSN, BCRAN, and CIT. Previous to joining Mentor Technologies, Stephen worked as a Network Systems Analyst for the North Broward Hospital District, where he designed and implemented their Metropolitan WAN consisting of four major sites and ten smaller remote sites. Stephen has a B.S. in Computer Science from Florida International University and an M.S. in Computer Technology from Barry University. Stephen is currently pursuing his CCIE certification.

Anthony Kwan, CCNP, CCDP, has worked in the Internetworking arena for over eight years and holds more than 14 Internetworking certifications. His networking expertise focuses on LAN/WAN design and troubleshooting, as well as voice, video, and VPN integration.

Christophe Paggen, CCIE #2659, joined Cisco Systems, Inc., in 1996, where he currently is a Network Design Engineer in the Advanced Network Solutions group. His primary focus is the redesign, optimization, and performance tuning of large-scale IP and multiprotocol enterprise networks, with a specialization in campus, local-area, and metropolitan-area networks. He holds a B.S. in Computer Science from IESSL (Liege, Belgium) and an M.S. in Economics from Université de Mons (Belgium).

Casimir Sammanasu is a Program Manager with Cisco Systems, Inc., and holds an M.S. Computer Science degree from DePaul University, Chicago, and an MBA degree from the University of Dallas. Casimir has developed LAN switching courses at Cisco in the past and is presently responsible for Cisco IOS curriculum that includes advanced technologies such as QoS, Multicast, Security, and VPN.

Dedications

Tim Boyles—Glory and thanks be to God for giving me the talent and for sustaining me when the going gets tough. To my wife, René, for putting up with the late nights and weekends. To my children, Andrew and Alyssa, for allowing me to take some time out of their schedule to finish the project. (Although they think it's pretty cool to see their old man in print!)

In memory of my daughter Ashley, who sees all things from the heavens.

“The heavens declare the glory of God; the skies proclaim the work of his hands.”—Psalms 19:1

Dave Hucaby—First, my thanks to Jesus Christ, my Lord and my best, best friend. Networking is great, but the abundant life you give is too wonderful! Thanks to my wife and best friend, Marci, for her love and support in everything I do. I'm also grateful to her for encouraging me to return for the second day of the CCIE lab, when I was ready to pack up and go home early. I'm glad I listened to her! Thanks to my girls—Lauren for encouraging me to play with her and forget stressful things, and Kara for waiting to be born until the book was nearly done. Thankfully, God enabled me to write late at night, while everybody else slept. Although this impacted our family time very little, a tired daddy is just not as much fun.

Lastly, I would like to thank my parents for their support; I'm especially grateful to my dad for sharing with me his love of engineering and his skills at technical writing.

Acknowledgments

Tim Boyles:

Chris Cleveland, Development Editor, who persevered to make this project all that it could be. Thanks for sorting out all the issues!

Brett Bartow, Executive Editor for keeping the project going among all the twists and turns. Thanks for steering the ship!

Dave Hucaby, for listening to all my late-night rants and being a great co-author to work with!

Howard Jones, for pinch-hitting on some last minute editing.

All of the technical editors that contributed to the success of this book. Thanks for keeping me honest with the material and all your diligence to make this a quality product. Thanks to, Chris Paggen, Steven Daleo, Casimir Samanasu, and Anthony Kwan. I couldn't have done it without you!

Dave Hucaby: Working with Chris Cleveland, Brett Bartow, and Amy Lewis, all with Cisco Press, has been great! These folks have been very patient with a new author and have gone extra miles to keep me focused on the task at hand. I've long been an avid fan and reader of Cisco Press books and am grateful for the opportunity to co-author one myself. Thanks to Tim Boyles for sharing the load and giving me advice along the way. Nathain Ingram, my Christian brother, deserves my thanks for being a steady source of encouragement and a great friend. Thanks to Eddie Lawrence for helping me work out some Catalyst switch logistics. Finally, I would like to thank the technical reviewers for making this a more accurate book. As well, I'm grateful to Kennedy Clark and Kevin Hamilton for writing the *real* switching book, *Cisco LAN Switching*. The more I'm exposed to other networking folks, the more I realize how little I know.

Contents at a Glance

	Introduction	xxiii
Chapter 1	All About the Cisco Certified Network Professional and Design Professional Certification	3
Chapter 2	Campus Network Design Models	15
Chapter 3	Basic Switch and Port Configuration	65
Chapter 4	VLANs and Trunking	97
Chapter 5	Redundant Switch Links	145
Chapter 6	Trunking with ATM LANE	203
Chapter 7	InterVLAN Routing	241
Chapter 8	Multilayer Switching	265
Chapter 9	Overview of Hot Standby Routing Protocol	301
Chapter 10	Multicasts	333
Chapter 11	Configuring Multicast Networks	369
Chapter 12	Controlling Access in the Campus Environment	393
Chapter 13	Monitoring and Troubleshooting	425
Chapter 14	Scenarios for Final Preparation	463
Appendix A	Answers to the “Do I Know This Already?” Quizzes and Q&A Sections	477
Index		529

Contents

	Introduction	xxiii
	Goals and Methods	xxiii
	Who Should Read This Book?	xxiii
	Strategies for Exam Preparation	xxiv
	How This Book Is Organized	xxiv
	Approach	xxvi
	Icons Used in This Book	xxviii
	Command Syntax Conventions	xxix
Chapter 1	All About the Cisco Certified Network Professional and Design Professional Certification	3
	Overview of Cisco Certifications	4
	Exams Required for Certification	5
	Other Cisco Certifications	6
	What's on the Switching Exam?	6
	Topics on the Exam	7
	Recommended Training Path for CCNP and CCDP	8
	How to Use This Book to Pass the Exam	9
	I've Taken BCMSN—Now What?	11
	I've Taken CLSC—Now What?	11
	I've Learned Switching From Experience, But I Will Not Be Taking the BCMSN Course—Now What?	12
	Conclusion	13
Chapter 2	Campus Network Design Models	15
	How to Best Use This Chapter	15
	“Do I Know This Already?” Quiz	16
	<i>Foundation Topics</i>	20
	Switching Functionality	20
	Layer 2 Switching	20
	Layer 3 Routing	21
	Layer 3 Switching	22

Layer 4 Switching	22
Multilayer Switching (MLS)	23
Campus Network Models	23
Shared Network Model	24
LAN Segmentation Model	25
Network Traffic Models	28
Predictable Network Model	30
Hierarchical Network Design	30
Access Layer	31
Distribution Layer	31
The Core Layer	32
Cisco Products in the Hierarchical Design	32
Access Layer Switches	33
Distribution Layer Switches	34
Core Layer Switches	36
Product Summary	37
Modular Network Design	39
The Switch Block	40
Sizing a Switch Block	41
The Core Block	43
Collapsed Core	44
Dual Core	45
Core Size in a Campus Network	46
Core Scalability	47
Layer 3 Core	48
Foundation Summary	49
Q&A	53
Scenarios	57
Scenario 2-1: Small Campus Network Design	57
Scenario 2-2: Medium Campus Network Design	57
Scenario 2-3: Large Enterprise Campus Network Design	57
Scenario Answers	59
Scenario 2-1 Answers: Small Campus Network Design	59
Scenario 2-2 Answers: Medium Campus Network Design	60
Scenario 2-3 Answers: Large Enterprise Campus Network Design	61

Chapter 3	Basic Switch and Port Configuration	65
	How to Best Use This Chapter	65
	“Do I Know This Already?” Quiz	66
	<i>Foundation Topics</i>	70
	Desktop Connectivity with Ethernet	70
	Ethernet	70
	Fast Ethernet	71
	Full-Duplex Fast Ethernet	72
	Gigabit Ethernet	73
	Desktop Connectivity with Token Ring	74
	Token Ring Bridging	75
	Connecting Switches	77
	Console Port Cables/Connectors	77
	Ethernet Port Cables/Connectors	77
	Gigabit Ethernet Port Cables/Connectors	78
	Token Ring Port Cables/Connectors	79
	Switch Management	80
	Identifying the Switch	80
	Setting the Hostname/System Name on an IOS-Based Switch	80
	Setting the Hostname/System Name on a CLI-Based Switch	80
	Passwords and User Access	81
	Setting Login Passwords on an IOS-Based Switch	81
	Setting Login Passwords on a CLI-Based Switch	81
	Remote Access	82
	Enabling Remote Access on an IOS-Based Switch	82
	Enabling Remote Access on a CLI-Based Switch	82
	Communicating Between Switches	83
	Cisco Discovery Protocol	83
	Switch Clustering and Stacking	85
	Switch Port Configuration	86
	Identifying Ports	86
	Assigning a Port Description on an IOS-Based Switch	86
	Assigning a Port Description on a CLI-Based Switch	86
	Port Speed	86
	Assigning Port Speed on an IOS-Based Switch	87
	Assigning Port Speed on a CLI-Based Switch	87
	Ethernet Port Mode	87
	Assigning the Ethernet Link Mode on an IOS-Based Switch	87
	Assigning the Ethernet Link Mode on a CLI-Based Switch	87

Token Ring Port Mode	88
Assigning the Token Ring Link Mode on a CLI-Based Switch	88
Foundation Summary	89
Q&A	92
Chapter 4 VLANs and Trunking	97
How to Best Use This Chapter	97
“Do I Know This Already?” Quiz	98
Foundation Topics	102
Virtual LANs	102
VLAN Membership	103
Static VLANs	103
Dynamic VLANs	105
Extent of VLANs	105
End-to-End VLANs	106
Local VLANs	106
VLAN Trunks	106
VLAN Frame Identification	108
Inter-Switch Link Protocol	109
IEEE 802.1Q Protocol	109
LAN Emulation (LANE)	111
IEEE 802.10	111
Dynamic Trunking Protocol	111
VLAN Trunk Configuration	111
VLAN Trunk Configuration on an IOS-Based Switch	112
VLAN Trunk Configuration on a CLI-Based Switch	112
VLAN Trunking Protocol	114
VTP Domains	114
VTP Modes	115
VTP Advertisements	115
VTP Configuration	119
Configuring a VTP Management Domain	119
Configuring a VTP Management Domain on an IOS-Based Switch	119
Configuring a VTP Management Domain on a CLI-Based Switch	119
Configuring the VTP Mode	119
Configuring the VTP Mode on an IOS-Based Switch	120
Configuring the VTP Mode on a CLI-Based Switch	120
Configuring the VTP Version	120

Configuring the VTP Version on an IOS-Based Switch	121	
Configuring the VTP Version on a CLI-Based Switch	122	
VTP Status	122	
VTP Pruning	123	
Enabling VTP Pruning on an IOS-Based Switch	125	
Enabling VTP Pruning on a CLI-Based Switch	125	
Token Ring VLANs	126	
TrBRF	127	
TrCRF	128	
TrCRF Redundancy	130	
VTP and Token Ring VLANs	130	
Duplicate Ring Protocol (DRiP)	131	
<i>Foundation Summary</i>	132	
<i>Q&A</i>	136	
<i>Scenarios</i>	140	
Scenario 4-1	140	
Scenario 4-2	141	
<i>Scenarios Answers</i>	142	
Scenario Answers 4-1	142	
Scenario Answers 4-2	142	
Chapter 5	Redundant Switch Links	145
	How to Best Use This Chapter	145
	“Do I Know This Already?” Quiz	146
	<i>Foundation Topics</i>	150
	Switch Port Aggregation with EtherChannel	150
	Bundling Ports with EtherChannel	150
	Distributing Traffic in EtherChannel	151
	Port Aggregation Protocol (PAgP)	153
	EtherChannel Configuration	154
	EtherChannel Configuration on a CLI-Based Switch	155
	EtherChannel Configuration on an IOS-Based Switch	155
	Displaying EtherChannel Configuration	155
	Spanning-Tree Protocol	156
	Bridging Loops	156
	Preventing Loops with Spanning-Tree Protocol	159

Spanning-Tree Communication: Bridge Protocol Data Units	160
Electing a Root Bridge	161
Electing Root Ports	163
Electing Designated Ports	165
STP States	168
STP Timers	170
Topology Changes	171
Spanning-Tree Design	172
Types of STP	172
Common Spanning Tree (CST)	172
Per-VLAN Spanning Tree (PVST)	172
Per-VLAN Spanning Tree Plus (PVST+)	173
STP Configuration	173
Root Bridge Placement	174
Root Bridge Configuration	178
Spanning-Tree Customization	179
Tuning the Root Path Cost	180
Tuning the Port ID	181
Viewing STP Status	182
Tuning Spanning-Tree Convergence	182
Modifying STP Timers	182
Redundant Link Convergence	184
Foundation Summary	188
Q&A	193
Scenarios	199
Scenario 5-1: Spanning-Tree Protocol Operation	199
Scenario Answers	200
Scenario 5-1 Answers: Spanning-Tree Protocol Operation	200
Chapter 6 Trunking with ATM LANE	203
How to Best Use This Chapter	203
“Do I Know This Already?” Quiz	204
Foundation Topics	208
ATM Review	208
Cells and SAR	209
ATM Model	210
Virtual Circuits	211
ATM Addressing	211

VPI/VCI Addresses	212
NSAP Addresses	212
Inherent ATM Protocols	213
LAN Emulation (LANE)	213
LANE Components	213
LANE Operation	216
Step 1: Contacting the LECS	216
Step 2: Contacting the LES	216
Step 3: Contacting the BUS	217
Step 4: Communicating Between LECs	217
Address Resolution	218
Address Resolution Scenario 1: Using IP ARP to Resolve MAC Addresses	218
Address Resolution Scenario 2: Using LE_ARP to Resolve NSAP Addresses	218
Design of LANE Components	219
LANE Component Placement	219
LANE Component Redundancy (SSRP)	220
LANE Configuration	220
Configuring the LES and BUS	223
Configuring the LECS	223
Configuring Each LEC	224
Viewing the LANE Configuration	224
Viewing Default NSAP Addresses	224
Viewing LES Status	225
Viewing BUS Status	225
Viewing the LECS Database	226
Viewing LEC Status	226
<i>Foundation Summary</i>	228
<i>Q&A</i>	231
<i>Scenarios</i>	236
Scenario 6-1	236
<i>Scenarios Answers</i>	238
Scenario 6-1 Answers	238
Chapter 7 InterVLAN Routing	241
How to Best Use This Chapter	241
“Do I Know This Already?” Quiz	242
<i>Foundation Topics</i>	245

InterVLAN Routing Background	245
InterVLAN Routing Design	245
Routing with Multiple Physical Links	246
Routing over Trunk Links	247
802.1Q and ISL Trunks	247
ATM LANE	248
Routing with an Integrated Router	249
InterVLAN Routing Configuration	250
Accessing the Route Processor	250
Establishing VLAN Connectivity	251
Establishing VLAN Connectivity with Physical Interfaces	251
Establishing VLAN Connectivity with Trunk Links	252
Establishing VLAN Connectivity with LANE	253
Establishing VLAN Connectivity with Integrated Routing Processors	254
Configure Routing Processes	254
Additional InterVLAN Routing Configurations	255
Foundation Summary	257
Q&A	259
Chapter 8 Multilayer Switching	265
How to Best Use This Chapter	265
“Do I Know This Already?” Quiz	266
Foundation Topics	269
Overview of Multilayer Switching	269
Multilayer Switching Components	270
MLS-RP Advertisements	271
Hello Messages	271
XTAGs	271
MLS Caching	272
Disabling MLS	274
Configuring Multilayer Switching	275
Displaying VTP Domain Information	277
Enabling MLS	278
VTP Domain Issues	279
MLS Management Interface	279
Verifying MLS-RP	280
Flow Masks	282
Output Lists	283

Input Access Lists	284
Configuring the MLS-SE	285
MLS Caching	285
Verifying MLS Configurations	287
External Router Support	288
Switch Inclusion Lists	289
Displaying MLS Cache Entries	289
Foundation Summary	291
Q&A	293
Scenarios	296
Scenario 8-1	296
Scenario 8-2	297
Scenarios Answers	298
Scenario 8-1 Answers	298
Router Configuration for Scenario 8-1	298
Switch Configuration for Scenario 8-1	298
Display for show mls include Command (Question 7)	299
Scenario 8-2 Answers	299
Chapter 9 Overview of Hot Standby Router Protocol	301
How to Best Use This Chapter	301
“Do I Know This Already?” Quiz	302
Foundation Topics	306
HSRP Overview	306
Issues with Traditional Methods	306
Default Gateways	306
Proxy ARP	307
Routing Information Protocol (RIP)	308
ICMP Router Discovery Protocol (IRDP)	308
Hot Standby Router Protocol	309
HSRP Group Members	310
Addressing HSRP Groups Across ISL Links	311
Multiple HSRP Groups	312
HSRP Operations	313
Active Router	313
Locating the Virtual Router MAC Address	313
Active and Standby Router Behavior	314

Anatomy of an HSRP Message	315
HSRP States	316
Configuring HSRP	317
Configuring an HSRP Standby Interface	317
Configuring HSRP Standby Priority	318
Configuring HSRP Standby Preempt	319
Configuring the Hello Message Timers	319
Understanding HSRP Interface Tracking	320
Configuring HSRP Tracking	322
HSRP Status	323
Troubleshooting HSRP	323
<i>Q&A</i>	<i>325</i>
<i>Scenarios</i>	<i>329</i>
Scenario 9-1	329
<i>Scenario Answers</i>	<i>330</i>
Scenario 9-1 Answers	330
Chapter 10 Multicasts	333
How to Best Use This Chapter	334
“Do I Know This Already?” Quiz	335
<i>Foundation Topics</i>	<i>338</i>
Multicast Overview	338
Unicast Traffic	338
Broadcast Traffic	340
Multicast Traffic	341
Characteristics of Multicast Traffic	342
Multicast Addressing	343
Multicast Address Structure	343
Mapping IP Multicast Addresses to Ethernet	344
Managing Multicast Traffic	345
Subscribing and Maintaining Groups	346
IGMP Version 1	347
Joining a Group Using IGMP Version 1	347
General Queries Using IGMP Version 1	348
Membership Queries Using IGMP Version 1	348
Leaving a Group Using IGMP Version 1	348
IGMP Version 2	349

Joining a Group Using IGMP v2	350
Querier Election Using IGMPv2	350
Maintaining a Group Using IGMPv2	352
Leaving a Group Using IGMPv2	352
Switching Multicast Traffic Using CGMP	353
Routing Multicast Traffic	354
Distribution Trees	355
Source-Specific Distribution Trees	355
Shared Distribution Trees	356
Scope of Delivery	357
Multicast Routing Protocols	358
Dense Mode Routing Protocols	358
DVMRP	359
MOSPF	359
PIMDM	360
Sparse Mode Routing Protocols	360
CBT	361
PIMSM	361
<i>Foundation Summary</i>	362
<i>Q&A</i>	364
Chapter 11 Configuring Multicast Networks	369
How to Best Use This Chapter	369
“Do I Know This Already?” Quiz	370
<i>Foundation Topics</i>	373
Planning for Multicast Services in a Network	373
Configuring IP Multicast	373
Enabling IP Multicast Routing	374
Enabling PIM on an Interface	374
Enabling PIM in Dense Mode	375
Enabling PIM in Sparse Mode	375
Enabling PIM in Sparse-Dense Mode	376
Verifying PIM Configuration	376
Selecting a Designated Router	376
Displaying PIM Neighbors	376
Configuring a Rendezvous Point	377
Auto-RP	378
Configuring Time-To-Live	381
Debugging Multicast	381

Configuring Internet Group Management Protocol (IGMP)	382
Configuring Cisco Group Management Protocol (CGMP)	383
Configuring CGMP Leave	384
Foundation Summary	385
Q&A	386
Scenarios	389
Scenario 11-1	389
Scenarios Answers	390
Scenario 11-1 Answers	390
Chapter 12 Controlling Access in the Campus Environment	393
How to Best Use This Chapter	393
“Do I Know This Already?” Quiz	394
Foundation Topics	398
Access Policies	398
Managing Network Devices	400
Physical Access	400
Passwords	400
Privilege Levels	402
Virtual Terminal Access	404
Access Layer Policy	406
Access Layer Port Security	407
Configuring Port Security at the Access Layer	407
Enabling and Verifying Port Security Using the set CLI on set Command-Based Switches	407
Enabling and Verifying Port Security on Cisco IOS Command-Based Switches	408
Distribution Layer Policy	408
Filtering Traffic at the Distribution Layer	409
IP Standard Access List Overview	410
IP Extended Access List Overview	411
Controlling Routing Update Traffic	413
Configuring Route Filtering	413
IP Route Filtering	414
Core Layer Policy	415

Foundation Summary 416**Q&A 417****Scenarios 420**

Scenario 12-1 420

Scenario 12-2 421

Scenarios Answers 422

Scenario 12-1 Answers 422

Scenario 12-2 Answers 422

Chapter 13 Monitoring and Troubleshooting 425

How to Best Use This Chapter 425

“Do I Know This Already?” Quiz 426

Foundation Topics 430

Monitoring Cisco Switches 430

Out-of-Band Management 430

Console Port Connection 430

Serial Line Internet Protocol (SLIP) 432

In-Band Management 433

SNMP 434

Telnet Client Access 438

Cisco Discovery Protocol (CDP) 439

Embedded Remote Monitoring 440

Switched Port Analyzer 441

CiscoWorks 2000 442

General Troubleshooting Model 444

Troubleshooting with show Commands 446

Physical Layer Troubleshooting 447

Troubleshooting Ethernet 448

Network Testing 449

Traceroute 450

Network Test Equipment 451

Volt-Ohm Meters, Digital Multimeters, and Cable Testers 452

TDRs and OTDRs 452

Breakout Boxes, Fox Boxes, and BERTs/BLERTs 453

Network Monitors 453

Network Analyzers 453

	<i>Foundation Summary</i>	454
	<i>Q&A</i>	456
	<i>Scenarios</i>	459
	Scenario 13-1	459
	Scenario 13-2	459
	<i>Scenarios Answers</i>	460
	Scenario 13-1 Answers	460
	Scenario 13-2 Answers	460
Chapter 14	Scenarios for Final Preparation	463
	Scenario 14-1	463
	Scenario 14-2	465
	Scenario 14-3	467
	<i>Scenarios Answers</i>	<i>469</i>
	Scenario 14-1 Answers	469
	Scenario 14-2 Answers	471
	Scenario 14-3 Answers	472
Appendix A	Answers to the “Do I Know This Already?” Quizzes and Q&A Sections	477
Index		529



INTRODUCTION

Professional certifications have been an important part of the computing industry for many years and will continue to become more important. Many reasons exist for these certifications, but the most popularly cited reason is that of credibility. All other considerations held equal, the certified employee/consultant/job candidate is considered more valuable than one who is not.

Goals and Methods

The most important and somewhat obvious goal of this book is to help you pass the Switching exam (#640-504). In fact, if the primary objective of this book was different, then the book's title would be misleading; however, the methods used in this book to help you pass the CCNP Switching exam are designed to also make you much more knowledgeable about how to do your job. While this book and the accompanying CD together have more than enough questions to help you prepare for the actual exam, the method in which they are used is not to simply make you memorize as many questions and answers as you possibly can.

The key approach used in this book is to help you discover the exam topics that you need to review in more depth, to help you fully understand and remember those details, and to help you prove to yourself that you have retained your knowledge of those topics. So this book does not try to help you pass by memorization but helps you truly learn and understand the topics. The Switching exam is just one of the foundation topics in the CCNP certification and the knowledge contained within is vitally important to consider yourself a truly skilled routing/switching engineer or specialist. This book would do you a disservice if it didn't attempt to help you learn the material. To that end, the book will help you pass the Switching exam by using the following methods:

- Helping you discover which test topics you have not mastered
- Providing explanations and information to fill in your knowledge gaps
- Supplying exercises and scenarios that enhance your ability to recall and deduce the answers to test questions
- Providing practice exercises on the topics and the testing process via test questions on the CD

Who Should Read This Book?

This book is not designed to be a general networking topics book, although it can be used for that purpose. This book is intended to tremendously increase your chances of passing the CCNP Switching exam. Although other objectives can be achieved from using this book, the book is written with one goal in mind: to help you pass the exam.

So why should you want to pass the CCNP Switching exam? Because it's one of the milestones towards getting the CCNP certification; no small feat in itself. What would getting the CCNP mean to you? A

raise, a promotion, recognition? How about to enhance your resume? To demonstrate that you are serious about continuing the learning process and that you're not content to rest on your laurels. To please your reseller-employer, who needs more certified employees for a higher discount from Cisco. Or one of many other reasons.

Strategies for Exam Preparation

The strategy you use for CCNP Switching might be slightly different than strategies used by other readers, mainly based on the skills, knowledge, and experience you already have obtained. For instance, if you have attended the BCMSN course, then you might take a different approach than someone who learned switching via on-the-job training. Chapter 1, "All About the Cisco Certified Network Professional and Design Professional Certification," includes a strategy that should closely match your background.

Regardless of the strategy you use or the background you have, the book is designed to help you get to the point where you can pass the exam with the least amount of time required. For instance, there is no need for you to practice or read about IP addressing and subnetting if you fully understand it already. However, many people like to make sure that they truly know a topic and thus read over material that they already know. Several book features will help you gain the confidence that you need to be convinced that you know some material already and to also help you know what topics you need to study more.

How This Book Is Organized

Although this book could be read cover-to-cover, it is designed to be flexible and allow you to easily move between chapters and sections of chapters to cover just the material that you need more work with. Chapter 1 provides an overview of the CCNP and CCDP certifications and offers some strategies for how to prepare for the exams. Chapters 2 through 13 are the core chapters and can be covered in any order. If you do intend to read them all, the order in the book is an excellent sequence to use. Chapter 14, "Scenarios for Final Preparation," provides many scenarios that will help you review and refine your knowledge, without giving you a false sense of preparedness that you would get with simply reviewing a set of multiple-choice questions.

The core chapters, Chapters 2 through 13, cover the following topics:

- **Chapter 2, "Campus Network Design Models"**—The role of switches as they apply to the OSI model are discussed here, as well as the design of the campus network using switches and routers. A three layer hierarchical model is also discussed in addition to the various Cisco products used in such a design.
- **Chapter 3, "Basic Switch and Port Configuration"**—This chapter covers the Ethernet, Fast Ethernet, and Gigabit Ethernet network media technologies; the use of Token Ring LAN media in switched networks; the physical cabling and connectivity used with Catalyst switches; basic Catalyst switch configuration and administration as well as techniques for interswitch communication; and the switch commands that can be used to configure a LAN port for use.
- **Chapter 4, "VLANs and Trunking"**—This chapter presents the process of defining common workgroups within a group of switches. Switch configuration for VLANs is covered, along with the method of identifying and transporting VLANs on various types of links. VLAN administration and management is presented through the configuration of the VLAN Trunking Protocol (VTP).

- **Chapter 5, “Redundant Switch Links”**—This chapter presents technologies that can be used in a campus network to provide higher reliability. Redundancy between switches, fault tolerance and recovery, and timely access are all techniques that are discussed. Each of these makes use of redundant links between switches and switch blocks.
- **Chapter 6, “Trunking with ATM LANE”**—This chapter presents a review of ATM and focuses on the use of LANE technology for trunking. While ATM is a very complex technology, it is presented only briefly to set the foundation for a more detailed discussion of LANE.
- **Chapter 7, “InterVLAN Routing”**—This chapter discusses routing between VLANs to provide complete connectivity across the switched network. Several design methodologies are presented, along with Cisco Catalyst and router configuration procedures for interVLAN routing.
- **Chapter 8, “Multilayer Switching”**—This chapter is an overview of multilayer switching (MLS), as well as how to configuring MLS on different devices that make up the switch block. Also covered are flow masks.
- **Chapter 9, “Overview of Hot Standby Routing Protocol”**—This chapter covers the use of HSRP in a campus environment, specifically how to implement redundant architectures and provide load sharing and backup capabilities to today’s enterprise networks.
- **Chapter 10, “Multicasts”**—This chapter discusses the definition of multicasts, multicast protocols, multicast networking on routers and switches, and different multicast routing protocols.
- **Chapter 11, “Configuring Multicast Networks”**—This chapter describes how to configure basic multicast networks. A more complete description of IP multicast routing commands used in this chapter is found on Cisco CCO in the documentation section. This information builds on that covered in Chapter 10.
- **Chapter 12, “Controlling Access in the Campus Environment”**—This chapter covers the definition of access policies, as well as basic security configurations of routers and switches. Also discussed are the different layers of the switch block and what policies should cover at each layer.
- **Chapter 13, “Monitoring and Troubleshooting”**—This chapter discusses a general model for troubleshooting, in addition to methods of monitoring and troubleshooting and the commands associated with each.

Additional scenarios in Chapter 14 provide a method of final preparation with more questions and exercises. Example test questions and the testing engine on the CD allow simulated exams for final practice.

Each of these chapters uses several features to help you make best use of your time in that chapter. The features are as follows:

- **“Do I Know This Already?” Quizzes and Quizlets**—Each chapter begins with a quiz that helps you determine the amount of time you need to spend studying that chapter. The quiz is broken into subdivisions, called “quizlets,” that correspond to a section of the chapter. Following the directions at the beginning of each chapter, the “Do I Know This Already?” quiz will direct you to study all or particular parts of the chapter.

- **Foundation Topics**—This is the core section of each chapter that explains the protocols, concepts, and configuration for the topics in the chapter.
- **Foundation Summary**—Near the end of each chapter, a summary collects the most important tables and figures from the chapter. The “Foundation Summary” section is designed to help you review the key concepts in the chapter if you score well on the “Do I Know This Already?” quiz, and they are excellent tools for last-minute review.
- **Q&A**—These end-of-the-chapter questions focus on recall, covering topics in the “Foundation Topics” section by using several types of questions. And because the “Do I Know This Already?” quiz questions can help increase your recall as well, they are restated in the “Q&A” sections. Restating these questions, along with new questions, provides a larger set of practice questions for when you finish a chapter and for final review when your exam date is approaching.
- **Scenarios**—Located at the end of most chapters, the scenarios allow a much more in-depth examination of a network implementation. Rather than posing a simple question asking for a single fact, the scenarios let you design and build networks (at least on paper) without the clues inherent in a multiple-choice quiz format.
- **CD-based practice exam**—The companion CD contains a large number of questions not included in the text of the book. You can answer these questions by using the simulated exam feature or by using the topical review feature. This is the best tool for helping you prepare for the test-taking process.

Approach

Retention and recall are the two features of human memory most closely related to performance on tests. This exam preparation guide focuses on increasing both retention and recall of the topics on the exam. The other human characteristic involved in successfully passing the exam is intelligence; this book does not address that issue!

Adult retention is typically less than that of children. For example, it is common for four-year-olds to pick up basic language skills in a new country faster than their parents. Children retain facts as an end unto itself; adults typically either need a stronger reason to remember a fact or must have a reason to think about that fact several times to retain it in memory. For these reasons, a student who attends a typical Cisco course and retains 50 percent of the material is actually quite an amazing student.

Memory recall is based on connectors to the information that needs to be recalled—the greater the number of connectors to a piece of information, the better chance and better speed of recall. For example, if the exam asks what VTP stands for, you automatically add information to the question. You know the topic is switching because of the nature of the test. You might recall the term “VTP domain,” which implies that this is a type of switch domain. You might also remember that we’re talking about VLANs. Having read the answer “VLAN Trunking Protocol,” then you might even have the infamous “aha” experience, in which you are then sure that your answer is correct (and possibly a brightly lit bulb is hovering over your head). All these added facts and assumptions are the connectors that eventually lead your brain to the fact that needs to be recalled. Of course, recall and retention work together. If you do not retain the knowledge, it will be difficult to recall it.

This book is designed with features to help you increase retention and recall. It does this in the following ways:

- By providing succinct and complete methods of helping you decide what you recall easily and what you do not recall at all.
- By giving references to the exact passages in the book that review those concepts you did not recall so that you can quickly be reminded about a fact or concept. Repeating information that connects to another concept helps retention, and describing the same concept in several ways throughout a chapter increases the number of connectors to the same pieces of information.
- By including exercise questions that supply fewer connectors than multiple-choice questions. This helps you exercise recall and avoids giving you a false sense of confidence, as an exercise with only multiple-choice questions might do. For example, fill-in-the-blank questions require you to have better recall than multiple-choice questions.
- By pulling the entire breadth of subject matter together. A separate chapter (Chapter 14) contains scenarios and several related questions that cover every topic on the exam and gives you the chance to prove that you have gained mastery over the subject matter. This reduces the connectors implied by questions residing in a particular chapter and requires you to exercise other connectors to remember the details.
- Finally, accompanying this book is a CD-ROM that has exam-like, multiple-choice questions. These are useful for you to practice taking the exam and to get accustomed to the time restrictions imposed during the exam.


Icons Used in This Book



Router



Bridge



Hub



DSU/CSU



Catalyst switch



Multilayer switch




ATM switch




ISDN switch



Communication server



Gateway



Access server




PC



PC with software



Sun Workstation



Mac




Terminal



File server



Web server



CiscoWorks Workstation



Printer



Laptop



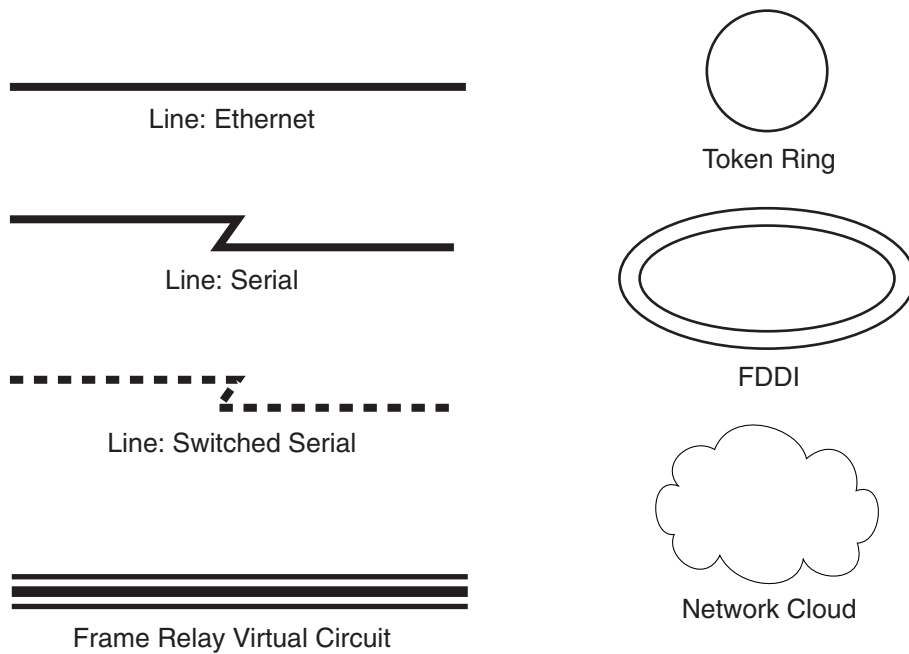
IBM mainframe



Front End Processor



Cluster Controller



Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets [] indicate optional elements.
- Braces { } indicate a required choice.
- Braces within brackets [{ }] indicate a required choice within an optional element.
- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- *Italics* indicate arguments for which you supply actual values.



All About the Cisco Certified Network Professional and Design Professional Certification

The Cisco Certified Network Professional (CCNP) and Cisco Certified Design Professional (CCDP) certifications are the second level of Cisco certifications and are becoming popular as more incentives become available to both certification holders and their employers. Cisco has designed both the CCNP and CCDP certifications as evidence that an individual has completed a rigorous path of testing in the network arena.

The CCNP and CCDP tracks require the candidate to be comfortable in advanced routing techniques, switching techniques, and dialup or RAS technology. In addition to those subjects, the CCNP must be able to, without a book, configure and troubleshoot a routed and switched network. The CCDP must demonstrate the skill to digest a vast quantity of user requirements and prepare a scalable design that fits the customer needs and requirements.

The CCNP is a more “hands on” certification that involves testing for a variety of routing and switching skills. Both configuration and troubleshooting are covered. Short of the CCIE, this certification is currently one of the most sought after.

The CCDP track focuses on designing scalable networks using routing and switching technologies. Testing involves the same battery of exams and subject matter as the CCNP track, with the exception of the Support exam. Instead, the CCDP track ends with the Cisco Internetwork Design exam. Because this certification focuses on the design aspects of internetworking, it is more suited for the pre-sales engineer or the network architect.

Because both the CCNP and CCDP maintain the same base set of requirements (except the final exam in the series), it is not surprising to find that a CCNP can produce a solid, scalable network design, while a CCDP can configure a router in a multiprotocol, routing, and switching environment. The key issue is the focus of the candidate in the business world.

Neither CCNP nor CCDP certification is a “one test and you pass” type of exam. Instead, each requires a series of either two or four exams. The exams are individually difficult because of the depth of understanding needed for each area of concentration. As well, each certification requires the Cisco Certified Network Associate (CCNA) certification as a prerequisite.

The focus of this book is the preparation and passing of the Cisco CCNP Switching Exam because this exam (or its content) is required for either CCNP or CCDP.

Overview of Cisco Certifications

Cisco's main motivation behind the current certification program is to provide a means of measuring the skills of people working for Cisco Resellers and Certified Partners. Cisco fulfills only a small portion of its orders via direct sale from Cisco; most times, a Cisco reseller is involved.

Cisco has not attempted to become the only source for consulting and implementation services for network deployment using Cisco products. In 1996–97 Cisco embarked on a channel program whereby business partners would be the eyes and ears to the smaller and midsize businesses that Cisco could not form a peer relationship with. Instead, Cisco partners of all sizes were carrying the Cisco flag into these smaller companies. With that business model, there was a great need to certify the skill levels of the partner companies.

The Cisco Certified Internetworking Expert (CCIE) program was Cisco's first cut at certifications. Introduced in 1994, the CCIE was designed to be one of the most respected, difficult-to-achieve certifications. To certify, a candidate had to pass a written test offered at Sylvan Prometric and then a two-day hands-on lab administered by Cisco.

Certifying resellers and services partners by using the number of employed CCIEs as the gauge worked well originally, partly because Cisco had significantly fewer partners than it does today. Cisco was using the number of CCIEs on staff as part of the criteria in determining the level of partner status for the company, which in turn dictated the discount received by the reseller when buying from Cisco. The number of resellers was growing and with Cisco's commitment to the lower tier market and smaller size business, it needed to have smaller integrators sized appropriately.

The CCIE certification fell short of the goal of helping to certify resellers and other partners as the number of partners increased to include some smaller integrators that were satisfying the medium and small business markets. Many smaller resellers that provided turnkey solutions for small businesses were not able to attain any degree of discount because of their size. Cisco, however, needed their skills to continue to capture the small business market, which is one of the largest markets in the internetworking arena today.

Cisco needed certifications that were less rigorous than the CCIE, which would allow Cisco more granularity in judging the skills on staff at a partner company. Therefore, Cisco created several additional certifications: the CCNA, CCDA, CCNP, and CCDP.

Two categories of certifications were developed: one to certify implementation skills and the other to certify design skills. Resellers working in a pre-sales environment need more design skills, whereas services companies require more implementation skills. So the CCNA and CCNP provide implementation-oriented certifications; whereas, the CCDA and CCDP provide design-oriented certifications.

Rather than just one level of certification besides CCIE, Cisco created two additional levels: an Associate level and a Professional level. CCNA is the more basic, and CCNP is the intermediate level between CCNA and CCIE. Likewise, CCDA is more basic than CCDP.

Several of the certifications require additional certifications as a prerequisite. For instance, CCNP certification requires CCNA first. Also, CCDP requires both CCDA and CCNA certification. CCIE, however, does not require any other certification prior to the written and lab tests, mainly for historical reasons.

Cisco certifications have become a much needed commodity in the internetworking world. The CCNP and CCDP certifications are truly another win-win situation for you and for Cisco.

Exams Required for Certification

To certify for CCNP or CCDP, successful completion of a group of exams is required. The exams generally match the same topics that are covered in one of the official Cisco courses. Table 1-1 outlines the exams and the courses with which they are most closely matched.

Table 1-1 *Exam-to-Course Mappings*

Certification	Exam Number	Name	Course Most Closely Matching Exam Requirements
CCNA	640-507	CCNA exam	Interconnecting Cisco Network Devices (ICND)
CCDA	640-441	CCDP Exam	Designing Cisco Networks (DCN)
CCNP	640-503	Routing Exam	Building Scalable Cisco Networks (BSCN)
	640-504	Switching Exam	Building Cisco Multilayer Switched Networks (BCMSN)
	640-505	Remote Access Exam	Building Cisco Remote Access Networks (BCRAN)
	640-509*	Foundation Exam	BSCN, BCMSN, and BCRAN
	640-506	Support Exam	Cisco Internetwork Troubleshooting (CIT)
CCDP	640-503	Routing Exam	Building Scalable Cisco Networks (BSCN)
	640-504	Switching Exam	Building Cisco Multilayer Switched Networks (BCMSN)
	640-505	Remote Access Exam	Building Cisco Remote Access Networks (BCRAN)
	640-509*	Foundation Exam	BSCN, BCMSN, and BCRAN
	640-025	CID Exam	Cisco Internetwork Design (CID)

*Exam 640-509 meets the same requirements as passing these three exams: 640-503, 640-504, and 640-505.

Other Cisco Certifications

The certifications mentioned so far are oriented toward routing and LAN switching. Cisco has many other certifications, which are summarized in Table 1-2. Refer to Cisco's web site at www.cisco.com/warp/public/10/wwtraining/certprog/index.html for more information.

Table 1-2 *Additional Cisco Certifications*

Certification	Purpose, Prerequisites
CCNA-WAN	Basic certification for Cisco WAN switches.
CCNP-WAN	Intermediate certification for Cisco WAN switches. Requires CCNA-WAN.
CCDP-WAN	Design certification for Cisco WAN switches. Requires CCNP-WAN.
CCIE-WAN	Expert level certification for Cisco WAN switches. No prerequisite. Requires exam and lab.
CCIE-ISP Dial	CCIE level certification for Internet service provider (ISP) and dial network skills. No prerequisite. Requires exam and lab.
CCIE-SNA-IP	Expert level certification for Cisco products and features used for melding SNA and IP networks. No prerequisite. Requires exam and lab.
CCIE-Design	Expert level certification that covers design principles related to the access, distribution, and core layers of large internetworks. It also requires candidates to have a thorough understanding of Campus Design, Multiservice, SNA-IP, and Network Management related design issues.
CCNP and CCDP specializations	Several specialized certifications are available for CCNP and CCDP (routing/switching). See www.cisco.com/warp/public/10/wwtraining/certprog/special/course.html for more details.

What's on the Switching Exam?

As with other Cisco exams, the exact exam content is not publicly known. In fact, Cisco makes fairly general Switching Exam content available to the public at www.cisco.com/warp/public/10/wwtraining/certprog/testing/pdf/bcmsn.pdf

In addition to the general content listed, this book is structured to cover the content of the Building Cisco Multilayer Switched Networks (BCMSN) course. This content provides full coverage of switching topics that might be encountered in either the Switching Exam or real-world CCNP/CCDP workplace.

Topics on the Exam

The following list outlines the various topics that you will likely encounter on the exam. The topics represent a detailed list for areas of focus but are not intended as a list of test question topics. Each listed item may have subitems that will be tested on.

Table 1-3 lists the exam topics in the order that they are found within this book.

Table 1-3 *CCNP/CCDP Switching Exam Topics*

Chapter	Topics
Chapter 2, “Campus Network Design Models”	Switching Functionality, Campus Network Models, Hierarchical Network Design, Cisco Products in the Hierarchical Design, and Modular Network Design
Chapter 3, “Basic Switch and Port Configuration”	Desktop Connectivity with Ethernet, Desktop Connectivity with Token Ring, Connecting Switches, Switch Management, and Switch Port Configuration
Chapter 4, “VLANs and Trunking”	Virtual LANs, VLAN Trunks, VLAN Trunk Configuration, VLAN Trunking Protocol, VTP Configuration, and VTP Pruning
Chapter 5, “Redundant Switch Links”	Switch Port Aggregation, Spanning-Tree Protocol (STP), STP Configuration, STP Design and Tuning, and STP Convergence Tuning
Chapter 6, “Trunking with ATM LANE”	ATM, LANE Operation, and LANE Configuration
Chapter 7, “InterVLAN Routing”	InterVLAN Routing Design and interVLAN Routing Configuration
Chapter 8, “Multilayer Switching”	Multilayer Switching, Flow Masks, and Multilayer Switching Configuration
Chapter 9, “Overview of Hot Standby Routing Protocol”	Configuring HSRP Operations, HSRP Router Roles, and HSRP Preempt Status
Chapter 10, “Multicasts”	Multicast Methods and Characteristics, Multicast and Ethernet Addressing, IGMP, Multicast Technology on Routers and Switches, and Multicast Routing Protocols
Chapter 11, “Configuring Multicast Networks”	Multicast Planning, Multicast Configuration, Configuring IGMP, and Configuring CGMP
Chapter 12, “Controlling Access in the Campus Environment”	Access Policies, Managing Network Devices, Access Layer Policy, Distribution Layer Policy, and Core Layer Policy

continues

Table 1-3 CCNP/CCDP Switching Exam Topics (Continued)

Chapter 13, “Monitoring and Troubleshooting”	Monitoring Cisco Switches, Monitoring Commands, General Troubleshooting Model, Troubleshooting Commands, and Physical Layer Troubleshooting
Chapter 14, “Scenarios for Final Preparation”	Case studies involving all areas of switching technology

The exam itself is a computer-based exam with multiple choice, fill-in-the-blank, and list-in-order style questions. The fill-in-the-blank questions must be filled in using the *complete* syntax for the command, including dashes and the like. For the fill-in-the-blank questions, a tile button is given that can be used to list a large number of commands in alphabetical order. This setup is a real life saver if you can’t remember if there is a dash or an “s” at the end of a command. Knowing the syntax is key, though, because the list contains some bogus commands as well as the real ones.

As with most of the Cisco exams, you cannot “mark” and return to a question. This requires that you answer a question before moving along, even if it means guessing at an answer. Remember that a blank answer is incorrect.

The exam can be taken at any Sylvan Prometric testing center (1-800-829-NETS or <http://www.2test.com>).

Recommended Training Path for CCNP and CCDP

The recommended training path for the Cisco CCNP 2.0 and CCDP 2.0 professional level certifications is as follows:

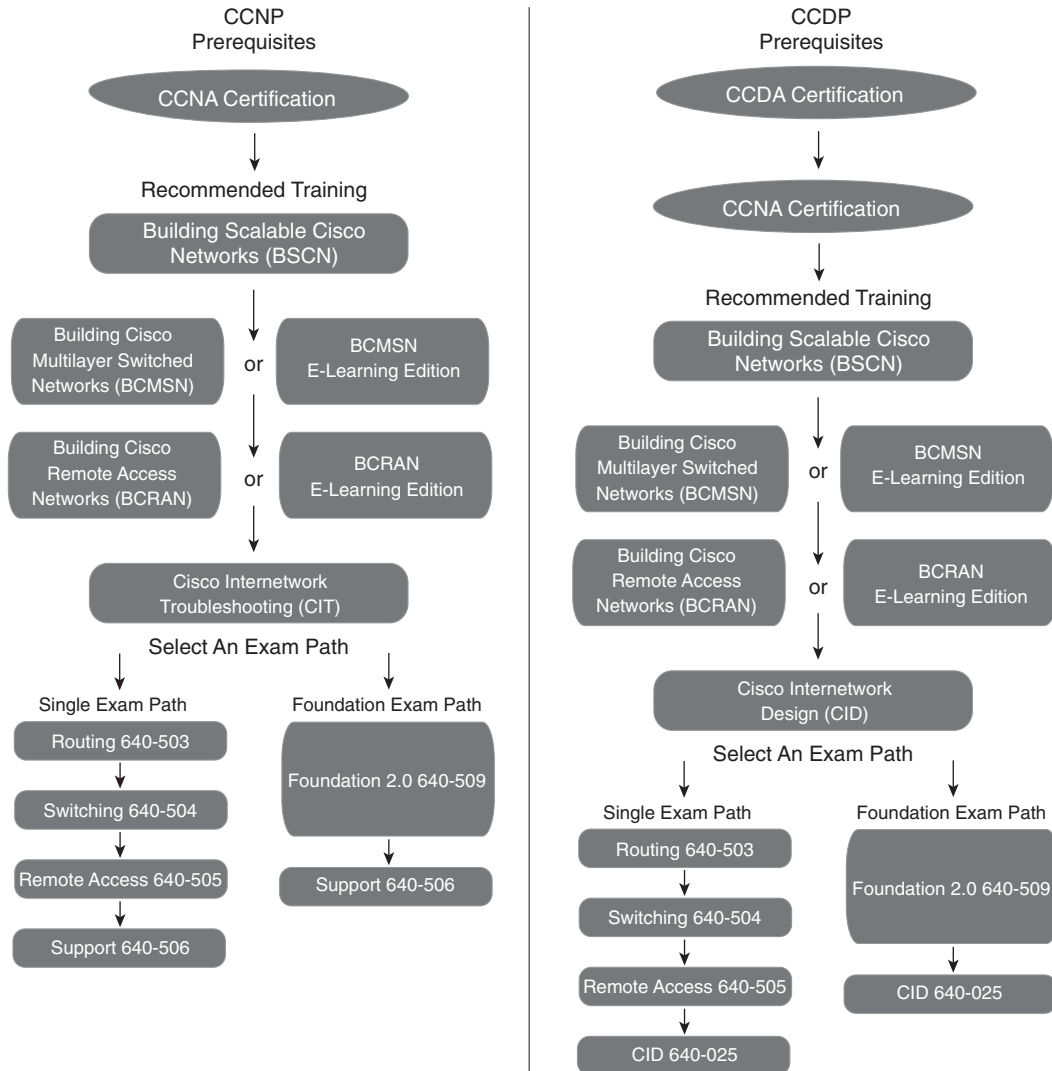
- **Building Scalable Cisco Networks (BSCN)**—Covers the advanced routing protocols and the scaling issues involved with a large routed network with multiple protocols.
- **Building Cisco Multilayer Switched Networks (BCMSN)**—Covers the switch infrastructure and the configuration in a large network environment.
- **Building Cisco Remote Access Networks (BCRAN)**—Covers the dialup and RAS issues involved in large scale remote access designs and implementations.

The CCNP then requires Cisco Internetworking Troubleshooting (CIT) as the final course. The CCDP requires Cisco Internetwork Design (CID) as the final course.

The recommended training courses will give you the basics to pass the exams for the CCNP or CCDP track. Cisco’s exams, however, will not necessarily correspond one-to-one with the curriculum of a given class. In essence, Cisco is not looking at the exams as a “fact-stuffing event” but rather as a gauge of how well you know and can use the technology.

Figure 1-1 illustrates the training track for CCNP and CCDP as of September 2000.

Figure 1-1 CCNP/CCDP 2.0 Training/Exam Track



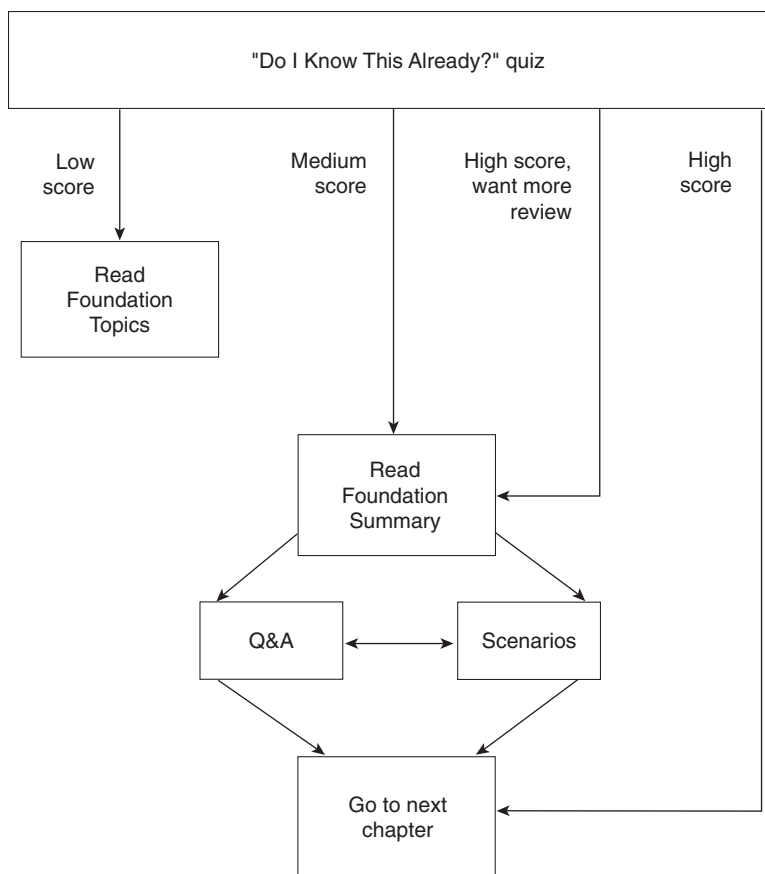
How to Use This Book to Pass the Exam

One way to use this book is to start at the beginning and read it cover to cover. Although that would help you prepare, most people would not take that much time, particularly if you already knew a lot about some of the topics in the book.

The rest of you might want to consider a different strategy on how to best use this book, depending on what training you have had. This book is designed to help you get the most out of the time you take to study.

At the beginning of each chapter, you are instructed on how to make the best use of your time reading that chapter, assuming that you are not going to read every detail. The instructions on how to use each chapter are outlined in Figure 1-2.

Figure 1-2 *How to Use Chapters 2 Through 13*



Each of these chapters begins with a quiz, which is broken into subdivisions called “quizlets.” If you get a high score, you might simply review the “Foundation Summary” section at the end of the chapter. If you score well on one quizlet but low on another, you are directed to the section of the chapter corresponding to the quizlet on which your score was low. If you score less than 50 percent on the overall quiz, you should read the whole chapter. Of course, these are simply guidelines.

After completing the core chapters (Chapters 2 through 13), you have several options for your next study activity. Chapter 14, “Scenarios for Final Preparation,” can be used to expand your thinking to more real-world examples. Network diagrams are presented, along with questions from a wide range of switching subjects covered in the core chapters.

If you want even more final preparation, you can go over the many practice questions located in each chapter and on the testing engine on the accompanying CD. All pre-chapter quizzes and chapter-ending questions, with answers, are in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections.” The CD includes testing software, as well as many additional questions in the format of the Switching Exam. The questions should be a valuable resource when performing final preparations.

I’ve Taken BCMSN—Now What?

For starters, you’ve taken the best path to prepare yourself. However, retaining and recalling most of the material covered in an instructor-led course is difficult after some time has passed. To reinforce what you’ve learned in the course, here are some strategies to follow:

- Use this book exactly as described in the opening pages of each of Chapters 2 through 13. Each chapter begins with a quiz that helps you assess the basic topics you need to study. The quiz also directs you to the appropriate sections of the chapter to increase your knowledge on specific topics rather than requiring you to read the entire chapter.
- Be certain to read the sections of this book that are not specifically covered in the BCMSN course. An example is Chapter 6, “Trunking with ATM LANE,” which is not part of the course. By looking through the list of chapters and topics, you should be able to pick out sections of the book that you are not so familiar with.
- Use Chapter 14, “Scenarios for Final Preparation,” as a “last pass” strategy. After reviewing all other material, quizzes, and “Q&A” sections at the end of each chapter, set aside time to go through the scenarios. Don’t expect to be able to answer all the scenario questions without having to refer to the chapters though. The scenarios were designed to make you think about a wide variety of topics and to provide some further structure for reviewing the book material.

I’ve Taken CLSC—Now What?

The current BCMSN class follows much of the material covered in the previous Cisco LAN Switching Configuration (CLSC) course. However, BCMSN has been massively reorganized and covers a good deal of additional subject matter. To fill in the gaps and provide a good study experience, here are some strategies to follow:

- Read and study through the *Building Cisco Multilayer Switched Networks* textbook from Cisco Press (ISBN 1-57870-093-0). This book closely follows the actual BCMSN course material and will give you a good review of the topics covered in the course.

- Read and study the chapters in this book that were not covered in CLSC. Some chapters to consider are
 - Chapter 2, “Campus Network Design Models”
 - Chapter 5, “Redundant Switch Links”
 - Chapter 6, “Trunking with ATM LANE”
 - Chapter 8, “Multilayer Switching”
- Use this book exactly as described in the opening pages of each of Chapters 2 through 13. Each chapter begins with a quiz that helps you assess the basic topics you need to study. The quiz also directs you to the appropriate sections of the chapter to increase your knowledge on specific topics rather than requiring you to read the entire chapter.
- Use Chapter 14, “Scenarios for Final Preparation,” as a “last pass” strategy. After reviewing all other material, quizzes, and “Q&A” sections at the end of each chapter, set aside time to go through the scenarios. Don’t expect to be able to answer all the scenario questions without having to refer to the chapters though. The scenarios were designed to make you think about a wide variety of topics and to provide some further structure for reviewing the book material.

I’ve Learned Switching From Experience, But I Will Not Be Taking the BCMSN Course—Now What?

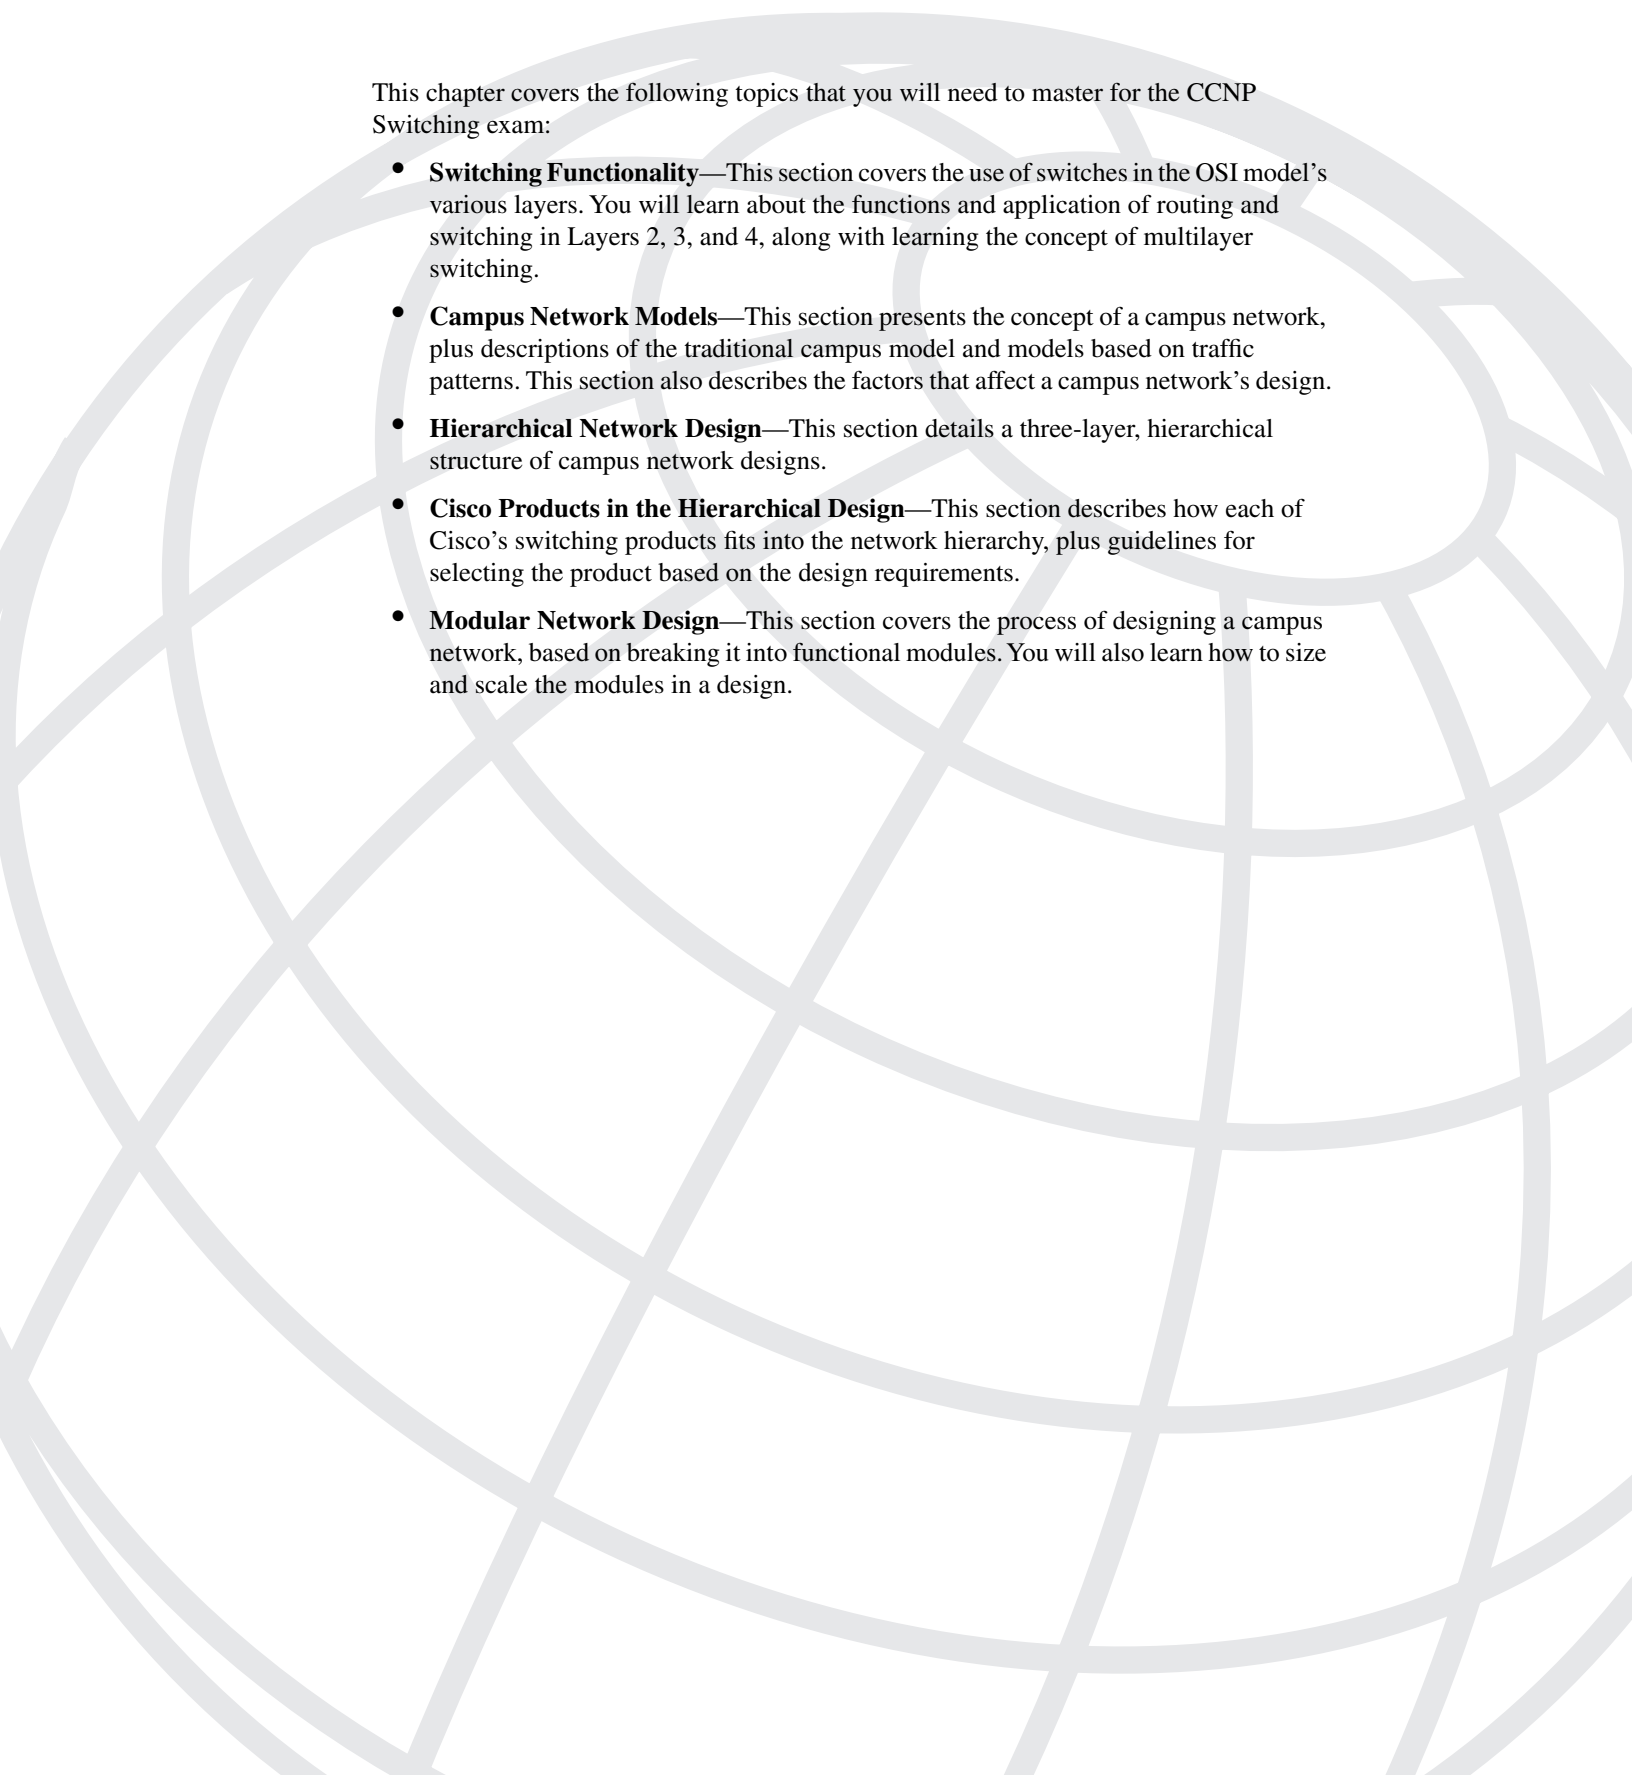

If you feel that you know a fair amount about switching topics already but are worried about the topics you have not worked with, some study strategies to follow are

- Use this book exactly as described in the opening pages of each of Chapters 2 through 13. Each chapter begins with a quiz that helps you assess the basic topics you need to study. The quiz also directs you to the appropriate sections of the chapter to increase your knowledge on specific topics rather than requiring you to read the entire chapter.
- Seriously think about studying and reviewing several chapters that cover “core” switching topics, because these topics are often complex and difficult to remember. Suggested chapters are
 - Chapter 2, “Campus Network Design Models”
 - Chapter 4, “VLANs and Trunking” (VTP)
 - Chapter 5, “Redundant Switch Links” (Spanning-Tree Protocol)
 - Chapter 8, “Multilayer Switching”
 - Chapters 10, “Multicast Networks,” and 11, “Configuring Multicast Networks”

- Use Chapter 14, “Scenarios for Final Preparation,” as a “last pass” strategy. After reviewing all other material, quizzes, and “Q&A” sections at the end of each chapter, set aside time to go through the scenarios. Don’t expect to be able to answer all the scenario questions without having to refer to the chapters though. The scenarios were designed to make you think about a wide variety of topics and to provide some further structure for reviewing the book material.

Conclusion

The *CCNP Switching Exam Certification Guide* is designed to help you attain CCNP certification by successfully preparing you to pass the Switching Exam. This book is the Switching Exam certification book from the only Cisco-authorized publisher. We at Cisco Press believe that this book certainly can help you prepare for and pass the Switching Exam—but the real work is up to you! We trust that your time will be well spent.



This chapter covers the following topics that you will need to master for the CCNP Switching exam:

- **Switching Functionality**—This section covers the use of switches in the OSI model's various layers. You will learn about the functions and application of routing and switching in Layers 2, 3, and 4, along with learning the concept of multilayer switching.
- **Campus Network Models**—This section presents the concept of a campus network, plus descriptions of the traditional campus model and models based on traffic patterns. This section also describes the factors that affect a campus network's design.
- **Hierarchical Network Design**—This section details a three-layer, hierarchical structure of campus network designs.
- **Cisco Products in the Hierarchical Design**—This section describes how each of Cisco's switching products fits into the network hierarchy, plus guidelines for selecting the product based on the design requirements.
- **Modular Network Design**—This section covers the process of designing a campus network, based on breaking it into functional modules. You will also learn how to size and scale the modules in a design.

Campus Network Design Models

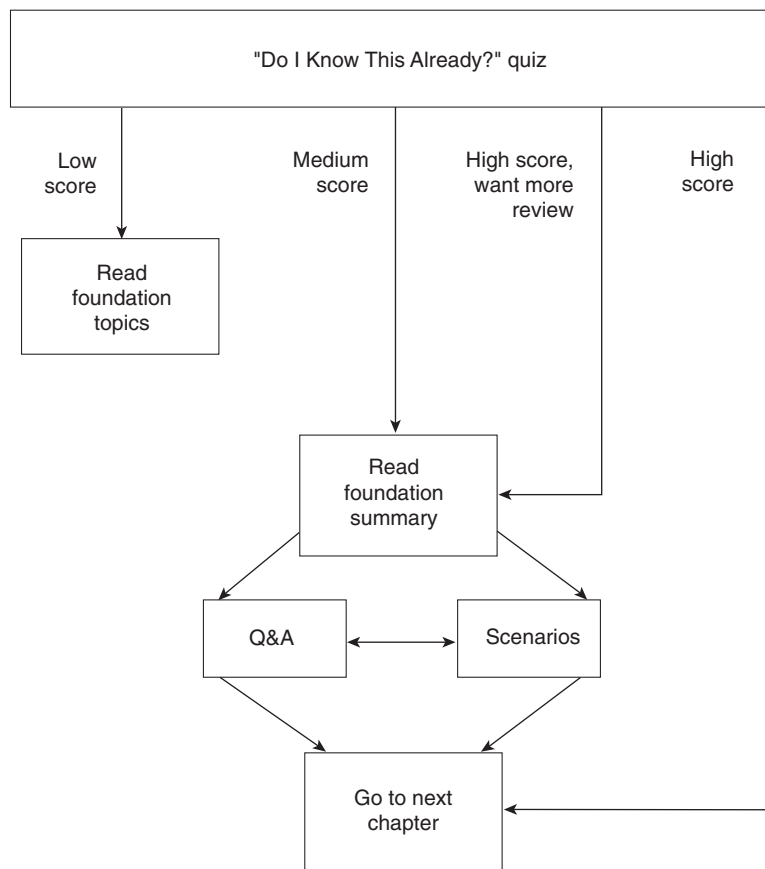
As campus networks have grown and technologies have matured, network engineers and architects have many more options to consider than the bridges, hubs, and routers traditionally put in place. Switches can be used to improve network performance in many ways. It is not enough, however, to simply replace existing shared networks with switched networks. The switching function alone alleviates congestion and increases bandwidth (in addition to more complex capabilities) if properly placed and designed.

This chapter presents a logical design process that can be used to build a new campus network or to modify and improve an existing network. A set of building blocks is introduced and is used to organize and streamline even a large, complex campus network. These building blocks are then properly placed using several campus design models to provide maximum efficiency, functionality, and scalability.

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and answers for all your work with this book in one place for easy reference.
- Take the “Do I Know This Already?” quiz and write down facts and concepts (even if you never look at the information again).
- Use the diagram in Figure 2-1 to guide you to the next step.

Figure 2-1 *How to Use This Chapter*

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide which parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The quiz helps you make good choices of how to spend your limited study time. The quiz is sectioned into five smaller “quizlets,” which correspond to the five major headings in the “Foundation Topics” section of the chapter. Although your answer may differ somewhat from the answers given, finding out if you have the basic understanding of topics presented in this chapter is what’s most important. You will find that the questions are open-ended rather than the multiple choice questions found on the exams. This setup will help you focus more on understanding the subject matter rather than memorizing details.

Use the scoresheet in Table 2-1 to record your score.

Table 2-1 *Scoresheet for Quiz and Quizlets*

Quizlet Number	Foundation Topics Sections Covering These Questions	Questions	Score
1	Switching Functionality	1–2	
2	Campus Networks, Traffic Pattern Models	3–6	
3	The Hierarchical Design Model	7–8	
4	Cisco Products in the Hierarchical Design	9–10	
5	Modular Network Design	11–12	
All questions		1–12	

1 Describe the differences between Layer 2, Layer 3, and Layer 4 switching.

2 What is multilayer switching (MLS)?

3 What is the 20/80 rule of networking?

4 What is a collision domain? Where does it exist in a switched LAN?

5 What is a broadcast domain? Where does it exist in a switched LAN?

6 What is a VLAN, and why is it used?

7 In which OSI layer do devices in the distribution layer typically operate?

8 How many layers are required in the hierarchical campus network design model?

9 Which Cisco switch products should be used in the distribution layer of a campus network?

10 When might a Catalyst 5000 be selected for use in a wiring closet? What attributes make it a good choice?

11 What building blocks are used to build a scalable campus network?

12 What are two types of core or backbone designs?

You can find the answers to the quiz in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections,” on page 477. The suggested choices for your next step are as follows:

- **6 or less overall score**—Read the entire chapter. This includes the “Foundation Topics” and “Foundation Summary” sections, the “Q&A” section, and the scenarios at the end of the chapter.
- **7–9 overall score**—Begin with the “Foundation Summary” section and then follow up with the “Q&A” section and the scenarios at the end of the chapter.
- **10 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” section and then go to the “Q&A” section and the scenarios at the end of the chapter. Otherwise, move on to Chapter 3, “Basic Switch and Port Configuration.”

Foundation Topics

Switching Functionality

To understand how switches and routers should be chosen and placed in a network design, you should first understand how to take advantage of data communication at different layers.

The OSI model separates data communication into seven layers, as shown in Table 2-2. Each layer has a specific function and a specific protocol so that two devices can exchange data on the same layer. A *protocol data unit (PDU)* is the generic name for a block of data that a layer on one device exchanges with the same layer on a peer device. A PDU is encapsulated in a layer's protocol before it is made available to a lower-level layer or unencapsulated before being handed to a higher-level layer.

Table 2-2 *Layers of Data Communications*

OSI Layer	Protocol Data Unit	Mechanism to Process PDU
7 (application)		
6 (presentation)		
5 (session)		
4 (transport)	TCP segment	TCP port
3 (network)	Packet	Router
2 (data link)	Frame	Switch/bridge
1 (physical)		

In Table 2-2, Layers 2, 3, and 4 are represented by the data link, network, and transport layers, respectively, with PDU's *frame*, *packet*, and *TCP segment*. When a TCP segment (Layer 4) needs to be transmitted to another station, the TCP segment is encapsulated as a packet (Layer 3) and further encapsulated as a frame (Layer 2). The receiving station will unencapsulate Layers 2 and 3 before processing the original TCP segment.

The layered protocols also apply to networking devices. For example, a Layer 2 device will transfer data by looking at Layer 2's PDU header information. Any upper-layer protocol will not be looked at or even understood. Layer-specific devices are discussed in detail in the following sections.

Layer 2 Switching

Devices that forward frames at Layer 2 involve the following functions:

- MAC addresses are learned from the source addresses of incoming frames.

- A table of MAC addresses and their associated bridge/switch ports is built and maintained.
- Broadcast and multicast frames are flooded out to all ports.
- Frames destined to unknown locations are flooded out to all ports.
- Bridges and switches communicate with each other using the Spanning-Tree Protocol to eliminate bridging loops.

A Layer 2 switch performs essentially the same function as a transparent bridge. However, a switch can have many ports and can perform *hardware-based bridging*. Frames are forwarded using specialized hardware called *application-specific integrated circuits (ASICs)*. This hardware gives switching great scalability, with wire-speed performance, low latency, low cost, and high port density.

As long as frames are being switched between two Layer 1 interfaces of the same media type, such as two Ethernet connections or an Ethernet connection and a Fast Ethernet connection, the Layer 2 frames do not have to be modified. However, if the two interfaces are different media, such as Ethernet and Token Ring or Ethernet and Fiber Distributed Data Interface (FDDI), the Layer 2 switch must translate the frame contents before sending out the Layer 1 interface.

Layer 2 switching is used primarily for workgroup connectivity and network segmentation. Traffic between users and servers in a workgroup can be contained within the switch. In addition, the number of stations on a network segment can be reduced with a switch, minimizing the collision domain size.

One drawback to Layer 2 switching is that it cannot be scaled very effectively. Switches must forward broadcast frames to all ports, causing large switched networks to become large broadcast domains. In addition, the Spanning-Tree Protocol can have a slow convergence time when the switch topology changes. It also can block certain switch ports, preventing data transfer. (Chapter 5, “Redundant Switch Links,” discusses the Spanning-Tree Protocol in further detail.) Layer 2 switching alone cannot provide an effective, scalable network design.

Layer 3 Routing

Devices involved in Layer 3 routing perform the following functions:

- Packets are forwarded between networks, based on Layer 3 addresses.
- An optimal path is determined for a packet to take through a network to the next router.
- Packet forwarding involves a table lookup of the destination network, next-hop router address, and the router’s own outbound interface.
- An optimal path can be chosen from among many possibilities.
- Routers communicate with each other using *routing protocols*.

By nature, routers do not forward broadcast packets and only forward multicast packets to destinations that are multicast clients. This action provides control over broadcast propagation and offers *segmentation* of the network into areas of common Layer 3 addressing.

Logical addressing is possible on a network with routers because the Layer 3 (network layer) address uniquely identifies a device only at the network layer of the OSI model. Actual frame forwarding occurs using the Layer 2 or data link address of devices. Therefore, some method must exist to associate a device's data link (MAC) address with its network layer (IP) address. A router must also have addresses from both layers assigned to each of its interfaces connected to a network. This assignment gives the router the functionality to support the logical network layer addresses assigned to the physical networks.

In addition, a router must examine the Layer 3 header of each packet before making a routing decision. Layer 3 security and control can be implemented on any router interface by using the source and destination addresses, protocol, or other Layer 3 attribute to make decisions on whether to limit or forward the packets.

Layer 3 routing is generally performed by microprocessor-based engines, which require CPU cycles to examine the network layer header of each packet. The routing table of optimal paths to Layer 3 networks can also be a large table of dynamic values, requiring a finite look-up delay. Although a router can be placed anywhere in a network, the router can become a bottleneck due to a latency of packet examination and processing.

Layer 3 Switching

Devices involved in Layer 3 switching perform the following functions:

- Packets are forwarded at Layer 3 just as a router would do.
- Packets are switched using specialized hardware ASICs for high-speed and low latency.
- Packets can be forwarded with security control and Quality of Service (QoS) using Layer 3 address information.

Layer 3 switches are designed to examine and forward packets in high-speed LAN environments. Whereas, a router might impose a bottleneck to forwarding throughput, a Layer 3 switch can be placed anywhere in the network.

Layer 4 Switching

Devices involved in Layer 4 switching perform the following functions:

- Packets are forwarded using hardware switching, based on both Layer 3 addressing and Layer 4 application information.
- Layer 3 protocol types (UDP or TCP, for example) in packet headers are examined.
- Layer 4 segment headers are examined to determine application port numbers.

Switching at Layer 4 allows finer control over the movement of types of information. For example, traffic can be prioritized according to the source and destination port numbers, and QoS can be defined for end users. Therefore, video or voice data can be switched at a higher level of service with more bandwidth availability than file transfer or HTTP traffic. Layer 4 port numbers for source and destination can also perform traffic accounting.

A Layer 4 switch must also allocate a large amount of memory to its forwarding tables. Layer 2 and Layer 3 devices have forwarding tables based on MAC and network addresses, making those tables only as large as the number of network devices. Layer 4 devices, however, must also keep track of application protocols and conversations occurring in the network. Their forwarding tables become proportional to the number of network devices multiplied by the number of applications.

Multilayer Switching (MLS)

Devices involved in MLS perform the following functions:

- Packets are forwarded in hardware that combines Layer 2, Layer 3, and Layer 4 switching.
- Packets are forwarded at wire speed.
- The traditional Layer 3 routing function is provided as *route one, switch many*. Routing sets up a network conversation, while hardware switches the *traffic flow* at high speeds.

Cisco switches perform multilayer switching at Layer 3 and Layer 4. At Layer 3, the Catalyst family of switches will cache traffic flows based on IP addresses. At Layer 4, the traffic flows are cached based on source and destination addresses, in addition to source and destination ports. All switching is performed in hardware, providing equal performance at both Layer 3 and Layer 4 switching.

Campus Network Models

A *campus network* is an enterprise network consisting of many LANs in one or more buildings, all connected and all usually in the same geographic area. A company typically owns the entire campus network, as well as the physical wiring. Campus networks usually consist of Ethernet, Token Ring, and FDDI LANs and higher-speed Fast Ethernet, Fast EtherChannel, and Gigabit Ethernet LANs.

An understanding of traffic flow is a vital part of the campus network design. While high-speed LAN technologies can be leveraged to improve any traffic movement, the emphasis should be on providing an overall design tuned to known, studied, or predicted traffic flows. The network traffic then can be effectively moved and managed, and the campus network can be scaled to support future needs.

The next sections present various network models that can be used to classify and also to design campus networks. Beginning with traditional shared networks, the models build on each other to leverage traffic movement and provide predictable behavior.

Shared Network Model

Campus networks have traditionally been constructed of a single LAN for all users to connect to and use. All devices on the LAN were forced to share the available bandwidth. LAN media such as Ethernet and Token Ring both have distance limitations, as well as limitations on the number of devices that could be connected to a single LAN.

Network availability and performance both declined as the number of connected devices increased. For example, an Ethernet LAN required all devices to share the available 10-Mbps half-duplex bandwidth. Ethernet also used the carrier sense multiple access collision detect (CSMA/CD) scheme to determine when a device could transmit data on the shared LAN. If two or more devices tried to transmit at the same time, network collisions occurred and all devices had to become silent and wait to retransmit their data. This type of LAN is a *collision domain* because all devices were susceptible to collisions. Token Ring LANs are not susceptible to collisions because they are deterministic and allow stations to transmit only when they receive a “token” that passes around the ring.

One solution used to relieve network congestion was to segment or divide a LAN into discrete collision domains. This solution used transparent bridges, which only forwarded Layer 2 data frames to the network segment where the destination address was located. Bridges enabled the number of devices on a segment to be reduced, lessened the probability of collisions on segments, and increased the physical distance limitations by acting as a repeater.

Bridges normally forward frames to the LAN segment where the destination address is located. However, frames containing the broadcast MAC address (ff:ff:ff:ff:ff:ff) must be flooded out to all connected segments. Broadcast frames are usually associated with requests for information or services, including network service announcements. IP uses broadcasts for Address Resolution Protocol (ARP) requests to ask what MAC address is associated with a particular IP address. Other examples of broadcast frames include IPX Get Nearest Server (GNS) requests, Service Advertising Protocol (SAP) announcements, Routing Information Protocol (RIP—both IP and IPX) advertisements, and NetBIOS name requests. A *broadcast domain* is a group of network segments where a broadcast is flooded.

Multicast traffic is traffic that is destined for a specific set or group of users, regardless of their location on the campus network. Multicast frames must be flooded to all segments because they are a form of broadcast. Although end users must join a multicast group to enable their applications to process and receive the multicast data, a bridge must flood the traffic to all segments because it doesn't know which stations are members of the multicast group. Multicast frames will use shared bandwidth on a segment, but will not force the use of CPU resources on every connected device. Only the CPUs that are registered as multicast group members will actually process those frames. Some multicast traffic is sporadic, as in the case of various

routing protocol advertisements, while other traffic such as Cisco IP/TV multicast video can consume most or all the network resources with a steady stream of real-time data.

Broadcast traffic presents a two-fold performance problem on a bridged LAN because all broadcast frames flood all bridged network segments. First, as a network grows, the broadcast traffic can grow in proportion and monopolize the available bandwidth. Secondly, all end-user stations must listen to, decode, and process every broadcast frame. This function is performed by the CPU, which must look further into the frame to see with which upper layer protocol the broadcast is associated. While today's CPUs are robust and might not show a noticeable degradation from processing broadcasts, forcing unnecessary broadcast loads upon every end user is not wise.


NOTE For a discussion of analysis performed by Cisco on the effects of various protocol broadcasts on CPU performance, refer to <http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd20e.htm>

LAN Segmentation Model

Referred to as *network segmentation*, localizing the traffic and effectively reducing the number of stations on a segment is necessary to prevent collisions and broadcasts from reducing a network segment's performance. By reducing the number of stations, the probability of a collision decreases because fewer stations can be transmitting at a given time. For broadcast containment, the idea is to provide a barrier at the edge of a LAN segment so that broadcasts cannot pass or be forwarded on outward. The network designer can provide segmentation by using either a router or a switch.

Routers can be used to connect the smaller subnetworks and either route Layer 3 packets or bridge Layer 2 packets. The effect of collisions can be improved with fewer stations on each segment. A router cannot propagate a collision condition from one segment to another. As well, broadcasts are not forwarded to other subnets by default, unless bridging (or some other specialized feature) is enabled on the router. Figure 2-2 shows an example of how a campus network can be segmented physically by a router. Although broadcasts are contained, the router becomes a potential bottleneck because it must process and route every packet leaving each subnet.

Another option is to replace shared LAN segments with switches. Switches offer greater performance with dedicated bandwidth on each port. A switch can be thought of as a very fast multiport bridge. Each switch port becomes a separate collision domain, and will not propagate collisions to any other port. However, broadcast and multicast frames are flooded out all switch ports unless more advanced switch features are invoked. Multicast switch features are covered in Chapter 11, "Configuring Multicast Networks."

Figure 2-2 *Network Segmentation with a Router*

To contain broadcasts and segment a broadcast domain, implement virtual LANs (VLANs) within the switched network. A switch can logically divide its ports into isolated segments. VLANs are groups of switch ports (and the end devices they are connected to) that communicate as if attached to a single shared-media LAN segment. By definition, a VLAN becomes a single broadcast domain. VLAN devices don't have to be physically located on the same switch or in the same building, as long as the VLAN itself is somehow connected between switches end-to-end. Figure 2-3 shows how a network can be segmented into three broadcast and collision domains using three VLANs on a switch. Note that stations on a VLAN cannot communicate with stations on another VLAN in the figure—the VLANs are truly isolated.

By default, all ports on a switch are assigned to a single VLAN. With additional configuration, a switch can assign its ports to many specific VLANs. Each VLAN, although present on the same switch, is effectively separated from other VLANs. Frames will not be forwarded from one VLAN to another. To communicate between VLANs, a router (or Layer 3 device) is required as illustrated by Figure 2-4.

Figure 2-3 Segmentation Using VLANs






Figure 2-4 Routing Traffic with VLANs



Ports on the switch have been grouped and assigned to three VLANs. A port from each VLAN also connects to the router. The router then forwards packets between VLANs through these ports. Note that each switch link in the figure supports two VLANs. Because a switch link can be configured only for one VLAN, it has been configured for *trunking*, or carrying multiple VLANs. (Trunking is discussed in Chapter 4, “VLANs and Trunking.”)

To gain the most benefit from routed approaches and VLAN approaches, most campus networks are now built with both LAN switches and routers. Again, the Layer 2 switches are generally placed where the small broadcast domains are located, linked by routers that provide Layer 3 functionality. In this manner, broadcast traffic can be controlled or limited. Users also can be organized and given access to common workgroups, while traffic between workgroups can be interconnected and secured. Figure 2-5 illustrates the structure of a typical routed and switched campus network.

Figure 2-5 Typical Campus Network Structure



Network Traffic Models

To design and build a successful campus network, you must gain a thorough understanding of the traffic generated by applications in use, plus the traffic flow to and from the user communities. All devices on the network will produce data to be transported across the network. Each device could involve many applications that generate data with differing patterns and loads.

Applications such as electronic mail, word processing, printing, file transfer, and most web browsers bring about data traffic patterns that are predictable from source to destination. However, newer applications such as videoconferencing, TV or video broadcasts, and IP telephony have a more dynamic user base, which makes traffic patterns difficult to predict or model.

Traditionally, users with similar applications or needs have been placed in common workgroups, along with the servers they access most often. Whether these workgroups are logical (VLAN) or physical networks, the idea is to keep the majority of traffic between clients and servers limited to the local network segment. In the case of the switched LANs connected by routers mentioned earlier, both clients and servers would be connected to a Layer 2 switch in the proximity of the workgroup. This connection provides good performance while minimizing the traffic load on the routed network backbone.

This concept of network traffic patterns is known as the *80/20 rule*. In a properly designed campus network, 80 percent of the traffic on a given network segment is local (switched). No more than 20 percent of the traffic is expected to move across the network backbone (routed).

If the backbone becomes congested, the network administrator will realize that the 80/20 rule is no longer being met. What recourses are available to improve network performance again? Upgrading the campus backbone is not a desirable option, due to the expense and complexity. The whole idea behind the 80/20 rule is to keep traffic off the backbone in the first place. Instead, the administrator can implement the following solutions:

- Reassign existing resources to bring the users and servers closer together.
- Move applications and files to a different server to stay within a workgroup.
- Move users logically (assigned to new VLANs) or physically to stay near their workgroups.
- Add more servers, which can bring resources closer to the respective workgroups.

Needless to say, conforming modern campus networks to the 80/20 rule has become difficult for the network administrator. Newer applications still use the client/server model, but server portions have been centralized in most enterprises. For example, databases, Internet and intranet technologies, and electronic mail are all available from centralized servers. Not only do these applications involve larger amounts of data, they also require a greater percentage of traffic to cross a network backbone to reach common destinations—quite a departure from the 80/20 rule.

This new model of campus traffic has become known as the *20/80 rule*. Now, only 20 percent of the traffic is local to the workgroup, while at least 80 percent of the traffic is expected to travel off the local network and across the backbone.

This shift in traffic patterns puts a greater burden on the Layer 3 technology of the campus backbone. Now, because traffic from anywhere on the network can be destined for any other part of the network, the Layer 3 performance ideally should match the Layer 2 performance.

Generally, Layer 3 routing involves more processing resources because the data packets must be examined in greater depth. This added computation load can create bottlenecks in the campus network, unless carefully designed.

Likewise, a campus network with many VLANs can become difficult to manage. Before, VLANs were used to logically contain common workgroups and common traffic. With the 20/80 rule, end devices need to communicate with many other VLANs. Measuring traffic patterns and redesigning the campus network becomes too cumbersome just to keep up with the 20/80 rule model.

Predictable Network Model

Ideally, a network needs to be designed with a predictable behavior in mind to offer low maintenance and high availability. For example, a campus network needs to recover from failures and topology changes quickly and in a predetermined manner. The network should be scalable to easily support future expansions and upgrades. With a wide variety of multiprotocol and multicast traffic, the network should be able to support the 20/80 rule from a traffic standpoint. In other words, the network should be designed around traffic flows instead of a particular type of traffic.

Traffic flows in a campus network can be classified as three types, based on where the network service is located in relation to the end user. Table 2-3 lists these types along with the extent of the campus network that is crossed.

Table 2-3 *Types of Network Services*

Service Type	Location of Service	Extent of Traffic Flow
Local	Same segment/VLAN as user	Access layer only
Remote	Different segment/VLAN as user	Access to distribution layers
Enterprise	Central to all campus users	Access to distribution to core layers


The terms access layer, distribution layer, and core layer are each distinct components of the hierarchical network design model. The network is divided into logical levels or layers, according to function. These terms and the hierarchical network design are discussed in the next section.

Hierarchical Network Design

The campus network can be structured so that each of the three types of traffic flows or services outlined in Table 2-3 can be best supported. Cisco has devised a hierarchical approach to network design that enables network designers to logically create a network by defining and using layers of devices. The resulting network is efficient, intelligent, scalable, and easily managed.

The hierarchical model breaks a campus network down into three distinct layers, as illustrated in Figure 2-6.

Figure 2-6 *Hierarchical Network Design*



These layers are the *access layer*, *distribution layer*, and the *core layer*. Each layer has attributes that provide both physical and logical network functions at the appropriate point in the campus network. Understanding each layer and its functions or limitations is important so that the layer can be properly applied in the design process.

Access Layer

The access layer is present where the end users are connected into the network. Devices in this layer should have the following capabilities:

- Low cost
- High port density
- Scalable uplinks to higher layers
- User access functions—VLAN membership and traffic filtering based on MAC addresses
- Resiliency through multiple uplinks

Distribution Layer

The distribution layer provides interconnection between the access and core layers of the campus network. Devices in this layer should have the following capabilities:

- High Layer 3 throughput for packet handling
- *InterVLAN routing* through Layer 3 operations

- Media translation to transport data between dissimilar access layer media types
- Security and *policy-based connectivity* functions through access lists or packet filters

The Core Layer

The core layer of a campus network provides connectivity of all distribution layer devices. The core, sometimes referred to as the *backbone*, must be able to switch traffic as efficiently as possible. Core devices should have the following attributes:

- Very high throughput
- No unnecessary packet manipulations (access lists, packet filtering)
- No Layer 3 processing, unless required and very fast
- Redundancy and resiliency for high availability

Cisco Products in the Hierarchical Design

Before delving into the design practices needed to build a hierarchical campus network, you should have some idea of the actual devices that can be placed at each layer. Cisco has switching products tailored for layer functionality, as well as the size of campus network.

For the purposes of this discussion, a large campus can be considered to span across several or many buildings in a single location. A medium campus might have one or several collocated buildings, while a small campus might have only a single building.

Cisco products should be chosen based on the functionality that is expected at each layer of a small, medium, or large campus. The products available at press time are described in the sections that follow and are summarized in table form for comparison. Try not to get lost in the details of the tables. Rather, try to understand which switch fits into which layer for a given size of network.

Although campus network design is presented as a three-layer approach (access, distribution, and core layers), the hierarchy can be collapsed or simplified in certain cases. For example, small- or medium-sized campus networks may not have the size, interVLAN routing, or volume requirements that would require the functions of all three layers. Here, the distribution and core layers could be combined for simplicity and cost savings. In this case, switch products should be chosen based on the distribution layer features and access layer aggregation port densities needed.

Access Layer Switches

Recall that access layer devices should have these features:

- High port density to connect to end users
- Low cost
- Uplinks to higher layers of the campus network
- Layer 2 services (traffic filtering and VLAN membership)

For small campus networks, the Catalyst 1900 or 2820 series switches can be used as access layer devices. Their smaller port densities can connect 10BaseT users and hubs, while connecting to distribution layer switches with 100BaseX uplinks. The Catalyst 2900XL and 3500XL switches are useful to provide access to groups of less than 50 users and servers. Both switch families offer high performance backplanes for efficient switching and Gigabit Ethernet uplinks to distribution layer switches. The 3500XL family is also stackable, using Gigabit Ethernet links as a shared bus to add port density in an access layer wiring closet. These switch families also offer a greater feature set, including QoS and switch clustering for improved performance and management.

The Catalyst 4000 series switches provide advanced enterprise access layer functions. These switches can be used to connect groups of less than 100 users and servers, or up to 36 Gigabit Ethernet devices. Greater Layer 2 functionality is provided as security, multicast support, and QoS.

For large campuses, the Catalyst 5000 series can provide access layer connectivity. This family of switches is completely modular, with many network media modules to choose from. Mixing up to 11 modules in a single-chassis, media translation can occur as “any-to-any” (Gigabit Ethernet-to-ATM, Token Ring-to-Fast Ethernet, FDDI-to-ATM, and so on) switching. Generally, the 5000 series is used to provide access to more than 100 end users.

Table 2-4 lists each Catalyst switch family suitable for the access layer, along with the maximum port densities and backplane speeds.

Table 2-4 *Catalyst Switches for the Access Layer*

Catalyst Model	Max Port Density	Uplinks	Max Backplane	Other Features
1900	24 10BaseT	2 100BaseX	1 Gbps	Fast EtherChannel
2820	24 10BaseT	100BaseX/FDDI/ATM	1 Gbps	Fast EtherChannel
2900 XL	48 10/100	2 100 or 1000BaseX or ATM	24 Gbps	QoS

continues

Table 2-4 *Catalyst Switches for the Access Layer (Continued)*

Catalyst Model	Max Port Density	Uplinks	Max Backplane	Other Features
3500 XL	48 10/100	2 1000BaseX	10 Gbps	Stackable Gigabit
4000	96 10/100 36 1000BaseX	100 or 1000BaseX	24 Gbps	Security, QoS
5000	396 10/100	Any	3.6 Gbps	Modular, “any-to-any” switching

Distribution Layer Switches

Switches used in the distribution layer should offer these features:

- Aggregation of access layer devices
- High Layer 3 throughput—InterVLAN routing
- Robust Layer 3 functionality
- Security
- Media translation

In the distribution layer, uplinks from all access layer devices are aggregated, or come together. Therefore, the distribution layer switches must be capable of processing the total volume of traffic from all the connected devices. These switches should have a port density of high-speed links to support the collection of access layer switches.

VLANs and broadcast domains converge at the distribution layer, requiring routing, filtering, and security. The switches at this layer must be capable of performing multilayer switching with high throughput. This performance is usually accomplished with a route processor within the switch, but is also possible using an external router as the route processor. Only certain Catalyst switch models can provide multilayer switching; be sure to understand which ones can do this. (Chapter 8, “Multilayer Switching,” covers this topic in greater detail.)

The Catalyst 2926G can serve as a distribution layer switch for up to 24 10/100BaseT access layer uplinks, as might be found in small- to mid-sized networks. MLS is performed by a combination of imbedded *NetFlow* logic and an external Cisco router. A Layer 3 data flow between two endpoints is discovered and processed by the router, as would be expected. However, once the initial frame of the data flow is routed, the NetFlow logic caches the frame-forwarding information. Subsequent frames are switched at wire speed rather than involving the router.

NOTE

Based on port density and certain functionality, many Catalyst switches can be used in more than one layer of a campus network. For example, because the Catalyst 2926G offers only a fixed 24-port 10/100BaseT configuration with two Gigabit Ethernet uplinks, you might want to use it in wiring closets or the access layer to connect workgroups or hubs. The Gigabit Ethernet uplinks would then be used as links to distribution layer switches.

However, what sets these switches apart is the *MLS* or *Layer 3 switching* function. As you've learned, Layer 3 switching is best reserved for the distribution layer.

The Catalyst 2926G will be replaced by the Catalyst 2948G-L3 switch, with 48 10/100BaseT ports and two Gigabit Ethernet uplinks. Again, this switch would fit into the distribution layer of a small- to mid-sized campus. However, the Layer 3 switching function is performed completely within the switch rather than relying on an external router for assistance. The 2948G-L3 switch uses a technique called *Cisco Express Forwarding (CEF)*, which was developed for higher-end switch routers. Here, the switch's RISC processor keeps a topology map of the entire network and distributes that map to each port-based ASIC. Layer 3 switching can be performed at wire speed in hardware on the ASICs for IP, IPX, and IP multicast protocols. Other protocols are Layer 2 switched (bridged) at wire speed on the ASICs.

The Catalyst 4908G-L3 offers the same functionality as the 2948G-L3, but with eight Gigabit Ethernet ports. Layer 3 switching is also performed using CEF in the port-based ASICs. This switch can be used in the distribution layer of mid-sized networks, using the Gigabit Ethernet ports to aggregate access layer devices.

For larger campus networks, the Catalyst 5000/5500 and 6000/6500 families offer high densities of Fast and Gigabit Ethernet for the distribution layer. A fully populated Catalyst 5509, for example, can support up to 38 Gigabit Ethernet ports or 288 10/100 Ethernet ports. The Catalyst 5000 family can perform MLS using either an integrated Route Switch Module (RSM), Route Switch Feature Card (RSFC), or an external router, coupled with the integrated NetFlow Feature Card (NFFC).

The Catalyst 6000 family offers much higher performance and port density that can be used in large distribution layers. For example, the Catalyst 6509 can support up to 130 Gigabit Ethernet ports or 384 10/100 Ethernet ports. MLS can be performed using an integrated Multilayer Switch Feature Card (MSFC), providing a throughput of up to 150 million packets per second.

Table 2-5 in the section "Product Summary" provides information on Cisco distribution layer switch products based on campus size.

Core Layer Switches

Let's recall the features required in core layer switches:

- Fast data transport
- No “expensive” Layer 3 processing
- Very high throughput
- No unnecessary packet manipulations (access lists and packet filtering)
- No Layer 3 processing, unless required and very fast
- Redundancy and resiliency for high availability

Devices in the core layer or backbone of a campus network should be optimized for high-performance Layer 2 or Layer 3 switching. Because the core layer must handle large amounts of campus-wide data (due to the new 20/80 rule of traffic flow), the core layer should be designed with simplicity and efficiency in mind.

For small campus networks, the Catalyst 5000 or 6000 family can be used in the core layer. These switches are modular and provide high port densities of Fast and Gigabit Ethernet to aggregate access layer uplinks. If the distribution and core layers are combined, both of these switch families can support MLS with integrated modules.

In medium and large campus networks, the Catalyst 6000 family can be used. Again, high port densities of Gigabit Ethernet are possible. This family of switches has high-performance, scalable switching from 32 Gbps to 256 Gbps. If MLS is required, integrated modules can be used to provide 6 Mpps to 150 Mpps throughput. Layer 3 security, QoS, and routing protocol support are also available.

The Catalyst 8500 series can also be used in large campus networks that have a mix of high capacity interface types. These switches, termed *switch routers*, should be used in multiservice environments where high performance and QoS are necessary for the delivery of various types of data, such as voice or video. In addition, the 8500 series offer connectivity to a wide range of network media, much like traditional routers do. The media supported include Fast and Gigabit Ethernet, native ATM, and Packet over Sonet (PoS). Layer 3 switching is performed with CEF so that the switching/routing function is distributed across all interfaces. Thus, all routing protocols are made available through a full Cisco IOS implementation. The Catalyst 8500 supports up to 64 T1/E1, 128 OC3, 32 OC12, or 8 OC48 interfaces, and up to 128 Fast Ethernet or 64 Gigabit Ethernet interfaces.

Table 2-5 in the section, “Product Summary,” provides information on Cisco core layer switch products based on campus size.

Product Summary

As a quick review, see Table 2-5 for a summary of the various Catalyst switch families that can be used for various applications. The table is broken down by campus network size and by campus network layer. The application of a particular switch in a network layer is a matter of choice and is not required. For example, if an access layer wiring closet in a small campus network had 200 users attached, choosing a single Catalyst 5000 might make more sense than several Catalyst 1900s. In this case, the size of the access layer workgroup dictates the choice of switch and port density more than the overall size of the campus network.

Table 2-5 *Summary of Catalyst Switch Products and Typical Layer Applications*

Campus Size	Layer	Catalyst Switch	Key Features
Any	Access	1900	< 25 users 10BaseT 100FX uplinks
		2820	< 25 users 10BaseT 100FX/FDDI/ATM uplinks
		2900XL	< 50 users 10/100 100 or GE uplinks
		3500XL	< 50 users 10/100 GE uplinks Stackable
		4000	< 100 users 10/100 or GE GE uplinks
		5000	> 100 users any media Any type uplinks Low price/port
Small Campus	Distribution	2926G	< 25 10/100 access devices GE uplinks MLS
		2948G-L3	< 50 10/100 access devices GE uplinks Layer 3 switching (CEF)
		4908G-L3	8 GE access devices Layer 3 switching (CEF)
		5000	High 100 and GE densities Any media supported MLS

continues

Table 2-5 *Summary of Catalyst Switch Products and Typical Layer Applications (Continued)*

Campus Size	Layer	Catalyst Switch	Key Features
		6000	High 100 and GE densities High performance MLS Scalable for future growth
	Core	Usually combined with distribution	
Medium Campus	Distribution	5000	High 100 and GE densities Any media supported MLS
		6000	High 100 and GE densities High performance MLS Scalable for future growth
	Core	5000	High port densities Any media
		6000	High GE densities High performance Security and QoS Scalability
		8500	High performance 100/GE/ATM/SONET Security and QoS Scalability
Large Campus	Distribution	6000	High 100 and GE densities MLS Security and QoS High performance
	Core	6000	High 100 and GE densities Security and QoS High performance
		8500	100/GE/ATM/SONET Nonblocking Multilayer Switching (CEF) High performance

Modular Network Design

A campus network can be designed in a logical manner, using a modular approach. In this approach, each layer of the hierarchical network model can be broken down into basic functional units. These units, or modules, can then be sized appropriately and connected together, while allowing for future scalability and expansion.

Campus networks can be divided into the following basic elements:

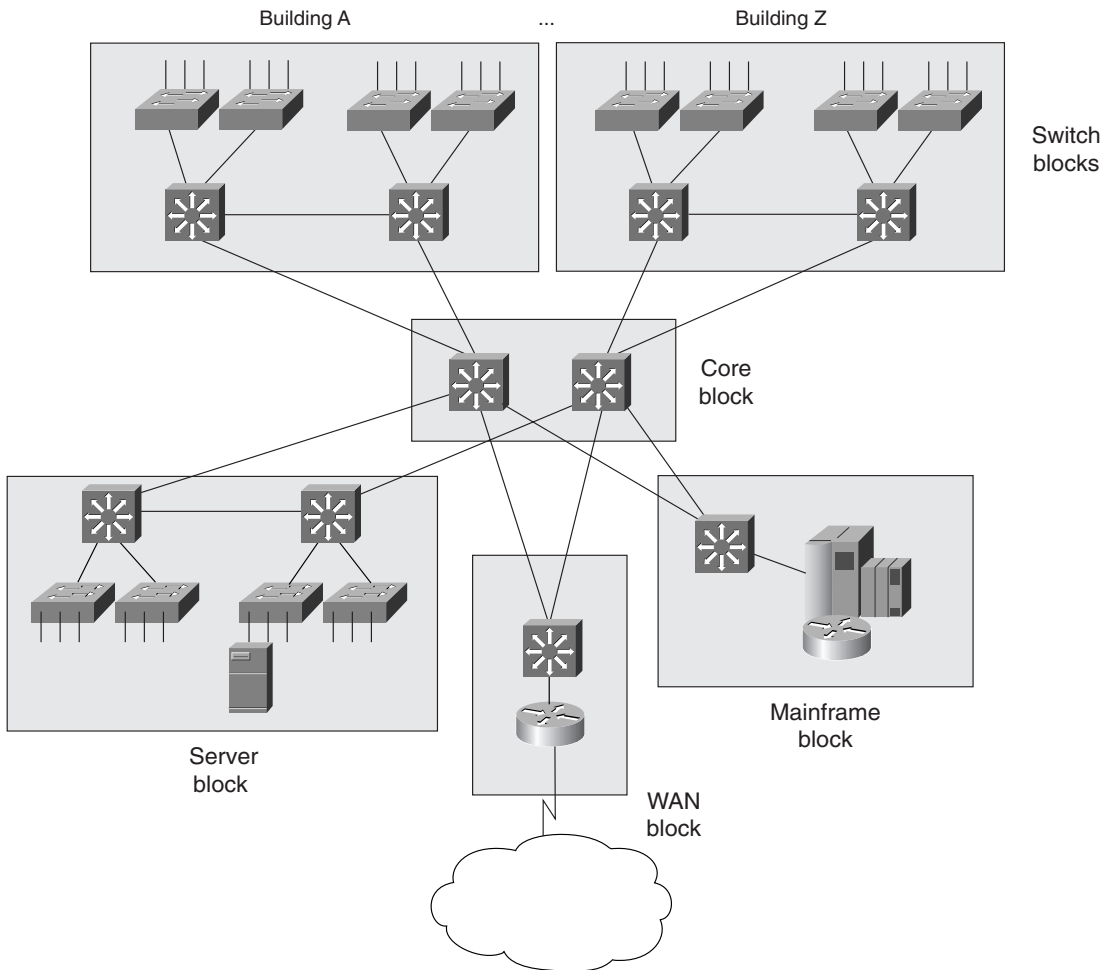
- Switch block
- Core block

Other related elements can exist. Although these elements don't contribute to the overall function of the campus network, they can be designed separately and added to the network design. These elements are

- Server block
- WAN block
- Mainframe block

Figure 2-7 shows the basic structure of a modular campus design. Notice how each of the building block elements can be confined to a certain area or function. Also notice how each is connected into the core block.

Figure 2-7 Modular Approach to Campus Network Design



The Switch Block

Recall how a campus network is divided into access, distribution, and core layers. Basically, the switch block contains switching devices from the access and distribution layers. All switch blocks then connect into the core block, providing end-to-end connectivity across the campus.

Switch blocks contain a balanced mix of Layer 2 and Layer 3 functionality, as might be present in the access and distribution layers. Layer 2 switches located in wiring closets (access layer) connect end users to the campus network. With one end user per switch port, each user receives dedicated bandwidth access.

Upstream, each access layer switch connects to devices in the distribution layer. Here, Layer 2 functionality transports data between all connected access switches at a central connection point. Layer 3 functionality can also be provided in the form of routing and other networking services (security, QoS, and so on). Therefore, a distribution layer device can be one of the following:

- A combination of a switch and an external router
- A multilayer switch

These Layer 3 distribution devices are discussed in more detail in Chapter 7, “InterVLAN Routing.”

The distribution layer also shields the switch block from certain failures or conditions in other parts of the network. For example, broadcasts will not be propagated from the switch block into the core and into other switch blocks. Therefore, the Spanning-Tree Protocol will be confined to each switch block, where a VLAN is bounded, keeping the Spanning Tree domain well defined and controlled.

Access layer switches can support VLANs by assigning individual ports to specific VLAN numbers. In this way, stations connected to the ports configured for the same VLAN will appear on the same subnet. However, remember that a single VLAN can support multiple subnets. Because the switch ports are configured for a VLAN number only (and not a network address), any station connected to a port can present any subnet address range. The VLAN will function as traditional network media and allow any network address to be connected.

In this network design model, VLANs should not be extended beyond distribution switches. The distribution layer should always be the boundary of VLANs, subnets, and broadcasts. Although Layer 2 switches can extend VLANs to other switches and other layers of the hierarchy, this activity is discouraged. VLAN traffic should not traverse the network core. (*Trunking*, or the capability to carry many VLANs over a single connection, is discussed in Chapter 4.)

Sizing a Switch Block

Containing access and distribution layer devices, the switch block is simple in concept. You should consider several factors, however, to determine an appropriate size for the switch block. The range of switch devices available makes the size of the switch block very flexible. At the access layer, switch selection is usually based on port density or the number of connected users.

The distribution layer must be sized according to the number of access layer switches that are collapsed or brought into a distribution device. Factors to consider are

- Various types and patterns of traffic
- Amount of Layer 3 switching capacity at the distribution layer
- Number of users connected to the access layer switches

- Geographical boundaries of subnets or VLANs
- Size of Spanning Tree domains

Designing a switch block based solely on the number of users or stations that are contained within the block is usually inaccurate. As a rule of thumb, no more than 2000 users should be placed within a single switch block. Though useful for an initial estimate of a switch block's size, this idea doesn't take into account the many dynamic processes that occur on a functioning network.

Instead, switch block size should be primarily based on:

- Traffic types and behavior
- Size and number of common workgroups

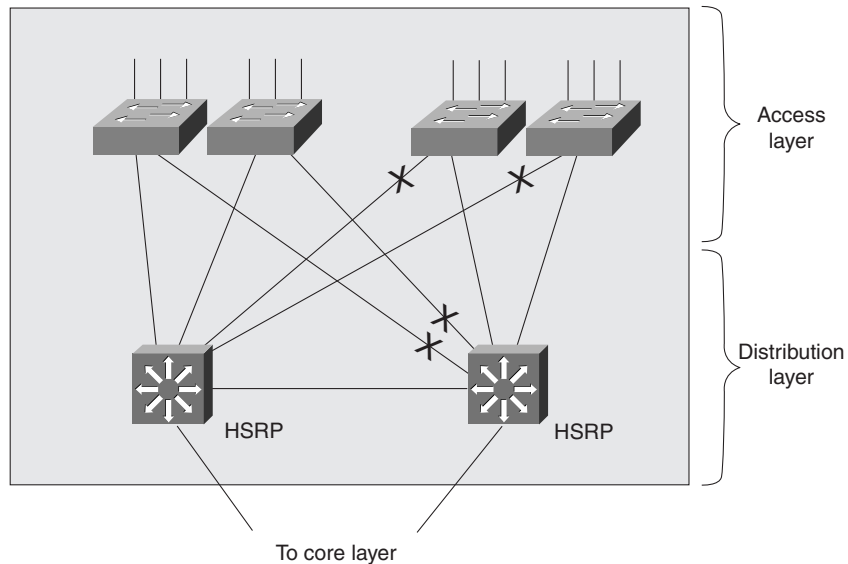
Due to the dynamic nature of networks, sizing a switch block too large to handle the load that is placed upon it is possible. Also, the number of users and applications on a network tend to grow over time. A provision to break up or downsize a switch block is necessary. Again, base these decisions on the actual traffic flows and patterns present in the switch block. These parameters can be estimated, modeled, or measured with network analysis applications and tools.

Generally, a switch block is too large if the following conditions are observed:

- The routers at the distribution layer become traffic bottlenecks. This congestion could be due to the volume of interVLAN traffic, intensive CPU processing, or switching times required by policy or security functions (access lists, queuing, and so on).
- Broadcast or multicast traffic slows down the switches and routers in the switch block. Broadcast and multicast traffic must be replicated and forwarded out many ports. This process requires some overhead in the router or switch and can become too great if significant traffic volumes are present.

Access switches can have one or more redundant links to distribution layer devices. This situation provides a fault tolerant environment, where access layer connectivity is preserved on a secondary link if the primary link fails. Chapter 5 discusses this matter. Generally, two distribution switches should be provided in each switch block for redundancy, with each access layer switch connecting to the two distribution switches with dual links.

Figure 2-8 shows a typical switch block design. Only one of the two links from each access layer switch will be in use at any time. At Layer 2, the Spanning-Tree Algorithm will keep one link in a blocking state and will fail over to the redundant link if the primary link fails. The Spanning-Tree Protocol is discussed in Chapter 5. At Layer 3, the two distribution switches can use Cisco's Hot Standby Router Protocol (HSRP) to provide an active IP gateway and a standby gateway. HSRP is discussed in Chapter 9, "Overview of Hot Standby Routing Protocol."

Figure 2-8 *Design of a Typical Switch Block*

The Core Block

A core block is required to connect two or more switch blocks in a campus network. Because all traffic passing to and from all switch blocks, server blocks, the WAN block, and the Internet must cross the core block, the core must be as efficient and resilient as possible. The core is the basic foundation of the campus network and carries much more traffic than any other block.

A network core can use any technology (frame, cell, or packet) to transport campus data. Some network designs utilize ATM in the core, while others are based on Ethernet. Because most campus networks are now using Gigabit Ethernet as a core technology, Ethernet core blocks will be reviewed at length. (Refer to Chapter 6, "Trunking with ATM LANE," for a discussion of ATM LANE features.)

Recall that the distribution layer provides Layer 3 functionality, and that the core is usually designed with Layer 2 devices. Therefore, individual IP subnets are used to connect all distribution and core switches. At least two subnets should be used to provide resiliency and load balancing into the core, although a single VLAN could be used. As VLANs end at the distribution layer, they are routed into the core.

The core block could consist of a single Layer 2 switch, taking in the two redundant links from the distribution layer switches. Due to the importance of the core block in a campus network, two or more identical switches should be implemented in the core to provide redundancy. As well, either Layer 2 or Layer 3 devices can be used in the core, depending on design requirements. This topic will be discussed further in the following sections.

The links between layers should also be designed to carry at least the amount of traffic load handled by the distribution switches. The links between core switches in the same core subnet should be of sufficient size to carry the aggregate amount of traffic coming into the core switch. Consider the average link utilization, but allow for future growth. An Ethernet core allows for simple and scalable upgrades of magnitude (Ethernet, Fast Ethernet, Fast EtherChannel, Gigabit Ethernet, Gigabit EtherChannel, and so on).

Two basic core block designs are presented in the following sections, each designed for the size of a campus network:

- Collapsed core
- Dual core


Collapsed Core

A collapsed core block is one where the core layer of the hierarchy is collapsed into the distribution layer. Here, both distribution and core functions are provided within the same switch devices. This situation is usually found in smaller campus networks, where a separate core layer (and additional cost or performance) is not warranted.

Figure 2-9 shows the basic collapsed core design. Although the distribution and core layer functions are performed in the same device, keeping these functions distinct and properly designed is important. Note also that the collapsed core is not an independent building block, but is integrated into the distribution layer of the individual standalone switch blocks.

In the collapsed core design, each access layer switch has a redundant link to each distribution/core layer switch. All Layer 3 subnets present in the access layer terminate at the Layer 3 ports of the distribution switches, as in the basic switch block design. The distribution/core switches are connected to each other by one or more links, completing a path to be used during a redundancy failover. Spanning Tree will keep one of the redundant links to the access layer blocked to prevent Layer 2 bridging loops.

Figure 2-9 Collapsed Core Block Design



However, at Layer 3, redundancy is provided through HSRP for IP. The two distribution switches will share a common default gateway address, but only one will be active at any time. In the event of a distribution/core switch failure, connectivity to the core will be maintained as the redundant Layer 3 switch takes over.

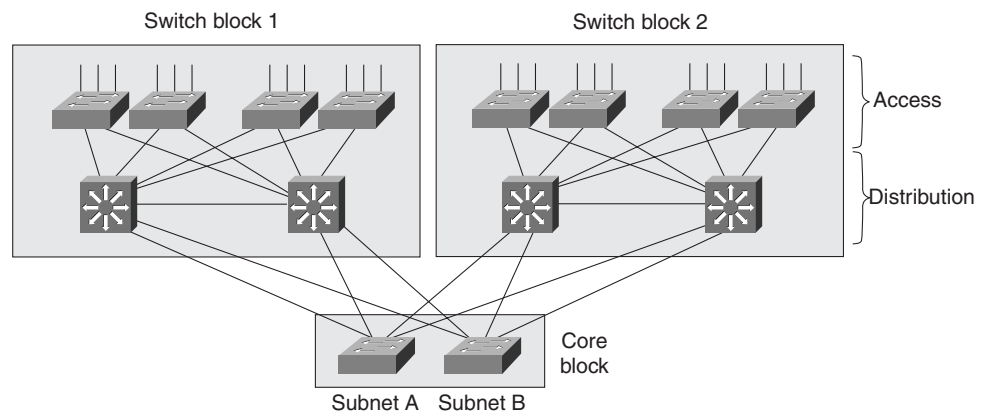
Why be concerned about the differences between Layer 2 and Layer 3 redundancy? Although the distribution/core switches have Layer 3 functionality, as in the case of MLS, understanding how MLS works in switches like the Catalyst 5000 and 6000 families is useful. A Layer 3 routing decision is first made on a traffic flow (either by an integrated or external router). The switches cache this information, and Layer 3 switching begins after a path is determined by the routing decision. Therefore, both Layer 2 and Layer 3 operations are still occurring on the distribution/core switches, each with different redundancy requirements. Chapter 8 covers MLS in more detail.

Dual Core

A dual core connects two or more switch blocks in a redundant fashion. Although the collapsed core can connect two switch blocks with some redundancy, the core is not scalable when more switch blocks are added. Figure 2-10 illustrates the dual core. Notice that this core appears as an independent module and is not merged into any other block or layer.

Normally, the dual core is built with Layer 2 switches to provide the simplest and most efficient throughput. Building a dual core with Layer 3 is possible, as discussed in the section “Layer 3 Core” later in the chapter. The dual core uses two identical switches to provide redundancy. Redundant links connect the distribution layer portion of each switch block to each of the dual core switches. Note the absence of any links between the two core switches. In a Layer 2 core, the switches are not linked to avoid any bridging loops.

Figure 2-10 *Dual Network Core Design*



In the dual core, each distribution switch has two equal-cost paths to the core, providing twice the available bandwidth. Both paths remain active because the distribution layer uses Layer 3 devices that can manage equal-cost paths in routing tables. In fact, the Layer 3 path determination across the core occurs without any reliance on Spanning Tree at all. Designing the core without links between the core switches removes any possibility of loops and eliminates the need for Spanning Tree in the core. The routing protocol in use determines the availability or loss of a neighboring Layer 3 device. Therefore, if one core switch fails, the routing protocol will reroute traffic using an alternate path through the remaining core switch.

Core Size in a Campus Network

The dual core is made up of redundant switches, and is bounded and isolated by Layer 3 devices. Routing protocols determine paths and maintain the operation of the core. As with any network, you must pay some attention to the overall design of the routers and routing protocols in the network. As routing protocols propagate updates throughout the network, network topologies might be undergoing change. The size of the network (the number of routers) then affects routing protocol performance, as updates are exchanged and network convergence takes place.

While the network shown in Figure 2-10 might look small with only two switch blocks of two Layer 3 switches (route processors within the distribution layer switches) each, large campus networks can have many switch blocks connected into the core block. Layer 2 devices are used in the core with usually only a single VLAN or subnet across the core. Therefore, all route processors connect into a single broadcast domain at the core. Each route processor must communicate with and keep information about each of its directly connected peers. Thus, most routing protocols have practical limits on the number of peer routers that can be connected.

Because two equal-cost paths from each distribution switch into the core, each router forms two peer relationships with every other router. Therefore, the actual maximum number of switch blocks that can be supported is half the number of distribution layer routers. For example, if routing protocols such as Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP) could support a maximum of 50 peers, only 25 switch blocks could be connected into the core. Also keep in mind that other types of campus modules (server blocks, WAN blocks, mainframe blocks, and so forth) connect into the core and create additional peer relationships.

One final core size consideration is related to the routing protocols and their support of equal-cost paths. In the case of dual core design, these paths must lead to isolated VLANs or subnets if a routing protocol supports two equal-cost paths. In other words, each path must be connected to a separate physical core switch. Two equal-cost paths are used in a dual core design with two Layer 2 switches. Likewise, a routing protocol that supports six equal-cost paths requires that the six distribution switch links be connected to exactly six Layer 2 devices in the core. Although this setup sounds complicated, it gives six times the redundancy and six times the available bandwidth into the core.

This leads to a final design point for the actual core switch—scale the core switch to match the incoming load. At a minimum, the core switch must be able to handle switching each of its incoming distribution links at 100% capacity.

Core Scalability

As the number of switch blocks increases, the core block must also be capable of scaling without redesign. Traditionally, hierarchical network designs have used Layer 2 switches at the access layer, Layer 3 devices at the distribution layer, and Layer 2 switches at the core. This design has been very cost effective and has provided high-performance connectivity between switch blocks in the campus.

Network growth dictates more switch blocks, which in turn requires more distribution switches with redundant paths into the core. The core must then be scaled to support the redundancy and the additional campus traffic load.

Providing redundant paths from the distribution switches into the core block allows the Layer 3 distribution switches to identify several equal-cost paths across the core. If the number of core switches must be increased for scalability, the number of equal-cost paths can become excessive—more than the routing protocols can handle.

Because the core block is formed with Layer 2 switches, the Spanning-Tree Protocol is used to prevent bridging loops. Two design decisions can be made in the core:

- Interconnect the core switches with redundant links.
- Interconnect the core switches only with distribution devices (not other core switches).


The first design decision creates the possibility of bridging loops by forming a triangle between a distribution switch and two core switches. Therefore, Spanning Tree must be enabled to dismantle the loop. Blocking one of the redundant distribution/core links effectively removes both redundancy and load balancing. Also, the Spanning Tree can be very slow to converge (more than 50 seconds) if the network topology changes. Large portions of the core block can become isolated while the network reconverges.

The second design decision allows Spanning Tree to continue to run within the core, but loops will never form because the core switches are not linked together. The Layer 3 devices in the distribution layer will use the equal-cost paths into the core to the maximum advantage. However, as switch blocks are added, you must take care not to exceed the number of equal-cost paths that can be supported by the routing protocols. As well, the switch block routers will appear to all sit on a single subnet because they are connected to a Layer 2 switch domain in the core. As discussed earlier, the number of routers that can be directly connected to peers is limited. One solution to this problem is to break up the core into multiple VLANs so that router peering is reduced in each VLAN.

Layer 3 Core

Layer 3 switching can also be used in the core to fully scale the core block for large campus networks. This approach also overcomes the problems of slow convergence, load balancing limitations, and router peering limitations. Figure 2-11 shows a network design using a Layer 3 core. Notice that the network structure is identical to one using a Layer 2 dual core. The main difference is that the core devices operate at Layer 3. Also notice that the core switches can have direct links to each other. Because of Layer 3 functionality, the direct links do not impose any bridging loops.

Figure 2-11 *Layer 3 Core Design*



With a Layer 3 core, the path determination intelligence occurs in both the distribution and core layers, allowing the number of core devices to be increased for scalability. Redundant paths also can be used to interconnect the core switches without concern for Layer 2 bridging loops.

Router peering problems are also overcome as the number of routers connected to individual subnets is reduced. Distribution devices are no longer considered peers with all other distribution devices. Instead, a distribution device peers only with a core switch on each link into the core. This advantage becomes especially important in very large campus networks involving more than 100 switch blocks.

The main concerns with implementing a Layer 3 core block are cost and performance. The Layer 3 devices required are more expensive than Layer 2 devices. The Layer 3 devices also need to have switching latencies comparable to their Layer 2 counterparts. Using a Layer 3 core also adds additional routing hops to cross-campus traffic. However, the increased scalability and flexibility of a Layer 3 core far outweigh any slight performance risks. In fact, with MLS and CEF performed in hardware, the additional overhead is negligible.

Foundation Summary

The Foundation Summary is a collection of tables and figures that provides a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final preparation before the exam, the following tables and figures will hopefully be a convenient way to review the day before the exam.

Table 2-6 *Layers of Data Communications*

OSI Layer	Protocol Data Unit	Mechanism to Process PDU
7 (application)		
6 (presentation)		
5 (session)		
4 (transport)	TCP segment	TCP port
3 (network)	Packet	Router
2 (data link)	Frame	Switch/bridge
1 (physical)		

Table 2-7 *Types of Network Services*

Service Type	Location of Service	Extent of Traffic Flow
Local	Same segment/VLAN as user	Access layer only
Remote	Different segment/VLAN as user	Access to distribution layers
Enterprise	Central to all campus users	Access to distribution to core layers

Table 2-8 *Comparison of Hierarchical Layers*

Layer	Attributes
Access	High port density to connect to end users Low cost Uplinks to higher layers of the campus network Layer 2 services (traffic filtering and VLAN membership)
Distribution	Aggregation of access layer devices High Layer 3 throughput—InterVLAN routing Robust Layer 3 functionality Security Media translation
Core	Fast data transport No “expensive” Layer 3 processing

Table 2-9 *Catalyst Switches for the Access Layer*

Catalyst Model	Max Port Density	Uplinks	Max Backplane	Other features
1900	24 10BaseT	2 100BaseX	1 Gbps	Fast EtherChannel
2820	24 10BaseT	100BaseX/FDDI/ATM	1 Gbps	Fast EtherChannel
2900 XL	48 10/100	2 100 or 1000BaseX or ATM	24 Gbps	QoS
3500 XL	48 10/100	2 1000BaseX	10 Gbps	Stackable Gigabit
4000	96 10/100 36 1000BaseX	100 or 1000BaseX	24 Gbps	Security, QoS
5000	396 10/100	Any	3.6 Gbps	Modular, “any-to-any” switching

Table 2-10 *Summary of Catalyst Switch Products and Typical Hierarchical Layer Applications*

Campus Size	Layer	Catalyst Switch	Key Features
Any	Access	1900	< 25 users 10BaseT 100FX uplinks
		2820	< 25 users 10BaseT 100FX/FDDI/ATM uplinks
		2900XL	< 50 users 10/100 100 or GE uplinks
		3500XL	< 50 users 10/100 GE uplinks Stackable
		4000	< 100 users 10/100 or GE GE uplinks
		5000	> 100 users any media Any type uplinks Low price/port
Small Campus	Distribution	2926G	< 25 10/100 access devices GE uplinks MLS
		2948G-L3	< 50 10/100 access devices GE uplinks Layer 3 switching (CEF)
		4908G-L3	8 GE access devices Layer 3 switching (CEF)
		5000	High 100 and GE densities Any media supported MLS
		6000	High 100 and GE densities High performance MLS Scalable for future growth
	Core	Usually combined with distribution	

continues

Table 2-10 *Summary of Catalyst Switch Products and Typical Hierarchical Layer Applications (Continued)*

Campus Size	Layer	Catalyst Switch	Key Features
Medium Campus	Distribution	5000	High 100 and GE densities Any media supported MLS
		6000	High 100 and GE densities High performance MLS Scalable for future growth
	Core	5000	High port densities Any media
		6000	High GE densities High performance Security and QoS Scalability
		8500	High performance 100/GE/ATM/SONET Security and QoS Scalability
	Large Campus	Distribution	6000
Core		6000	High 100 and GE densities Security and QoS High performance
		8500	100/GE/ATM/SONET Non-blocking MLS (CEF) High performance

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however, they are designed to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to these questions can be found in Appendix A, on page 477.

- 1 Where is the most appropriate place to connect a block of enterprise servers? Why?

- 2 Describe the differences between Layer 2, Layer 3, and Layer 4 switching.

- 3 What problems occur as switch blocks are added to a Layer 2 core design?

- 4 What is multilayer switching (MLS)?

- 5 How can redundancy be provided at the switch and core block layers? (Consider physical means, as well as functional methods using protocols, algorithms, and so on.)

6 What is the 20/80 rule of networking?

7 What factors should be considered when sizing a switch block?

8 What is a collision domain? Where does it exist in a switched LAN?

9 What are the signs of an oversized switch block?

10 What is a broadcast domain? Where does it exist in a switched LAN?

11 What are the attributes and issues of having a collapsed core block?

12 What is a VLAN, and why is it used?

13 When would a Layer 3 core block be desirable or necessary?

14 In which OSI layer do devices in the distribution layer usually operate?

15 What is network segmentation? When is it necessary? How is it done in a campus network design?

16 How many layers are required in the hierarchical campus network design model?

17 How many switches are sufficient in a core block design?

18 Which Cisco switch products should be used in the distribution layer of a campus network?

19 List three methods used for Layer 3 switching in Cisco products.

20 When might a Catalyst 5000 be selected for use in a wiring closet? What attributes make it a good choice?

21 Which Cisco Catalyst switches can be used in the access layer? (Consider the most important attributes of access layer switches.)

22 What building blocks are used to build a scalable campus network?

23 Which Cisco switch family has the most scalable performance?

24 What are two types of core or backbone designs?

25 Which Cisco switch family is the most flexible for network media and translation?

Scenarios

Scenario 2-1: Small Campus Network Design

A small company is housed in a single building with two floors. There are about 20 users per floor, all using Ethernet (10/100). All users need access to email and database servers that are located on the first floor. Future growth in the company is possible. Be sure to consider scalable growth in the network design. The tasks for this scenario are as follows:

- 1 Sketch out a network design for this company, based on the hierarchical design model.
- 2 What Cisco switch devices can be used in the access layer?
- 3 What Cisco switch devices can be used in the distribution layer?
- 4 Will a core layer be necessary? If so, what products will be used?

Scenario 2-2: Medium Campus Network Design


A company owns two buildings with three floors each. On each floor, there are 100 users organized by department. All but one floor are Ethernet users, while one floor has Token Ring users that access the corporate mainframe. An enterprise server farm is currently in use, connected by FDDI. The corporation would like to migrate to a fully switched environment for maximum performance. The mainframe users will continue to use Token Ring. The server farm will remain connected by FDDI during network migration, but can then use other Ethernet technologies (Fast Ethernet, Fast EtherChannel, and so on). The tasks for this scenario are as follows:

- 1 Sketch out the components for the new design, based on the hierarchical design model. The client corporation wants to deploy the same product in all locations to maximize their investment protection and ease of management.
- 2 What Cisco switch products can be used at each layer of the design? Because all switch devices will use the same base platform, note which media will be used in each location. In addition, note where Layer 3 or MLS will be used and the Cisco devices needed for that functionality.
- 3 Where should the enterprise resources (mainframe, server farm, and so forth) be placed?

Scenario 2-3: Large Enterprise Campus Network Design

A large corporate campus has an existing routed and switched network environment in place. Although the actual network is much larger, the basic structure is shown in Figure 2-12. Both the server farm and the mainframe are considered enterprise resources. All switches in use are Layer 2-only devices.

Figure 2-12 Diagram for Scenario 2-3



The tasks for this scenario are as follows:

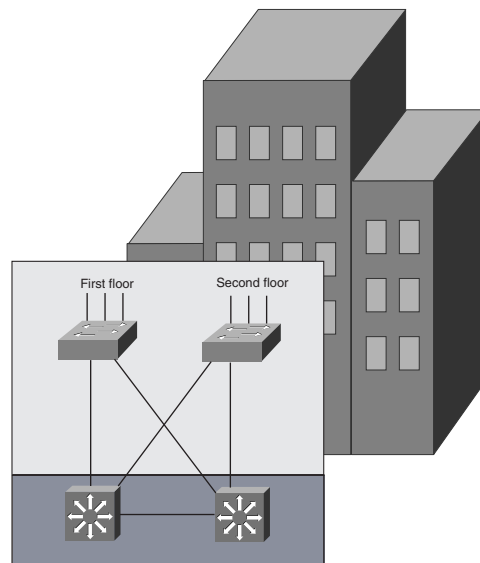
- 1 What improvements can be made to the network design? (Don't worry about IP addressing schemes or problems during migration to a new design.)
- 2 Redraw this network so that it follows the hierarchical campus network design model.
- 3 Where should the mainframe connect? Why?
- 4 Where should the servers connect? Why?
- 5 Where should the WAN connect to the network? Why?
- 6 Are the routers in use best suited where they are? Could they be replace or re-deployed elsewhere? How?
- 7 How has redundancy been addressed? What improvements could be suggested?
- 8 What types of core design might be appropriate for a large campus like this? What considerations should be made for core connectivity?
- 9 How should new resources (a workgroup, a server, a user, and so on) be connected to the network if all access layer ports are occupied? What problems would occur if the new resources were connected into the core or directly to a distribution layer switch?

Scenario Answers

Scenario 2-1 Answers: Small Campus Network Design

- 1 See Figure 2-13 for a network drawing. For a network of this size, a collapsed core can be used, as shown in the diagram.

Figure 2-13 *Network Diagram for Scenario 2-1 Solution*

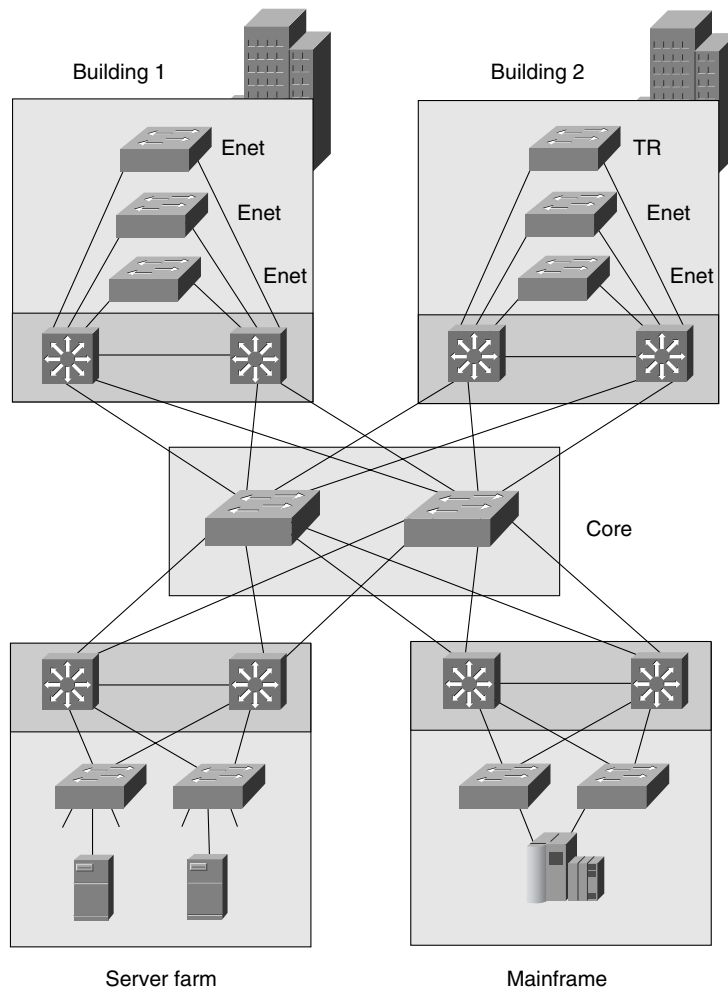


- 2 Access layer switches could be Catalyst 1900, if 10-Mbps Ethernet will be used for a long time. Otherwise, Catalyst 2900XL or 3500XL switches could be used for 10/100 connectivity.
- 3 At the distribution layer, either Catalyst 2926G, 2948G-L3, or Catalyst 5000s with an RSM could be used. (The current size of this network does not warrant Layer 3 functionality; a flat, switched network would work fine until the network grows.)
- 4 A core layer is not necessary. With the collapsed core design, however, growth and scalability are easily managed by adding switch blocks.

Scenario 2-2 Answers: Medium Campus Network Design

- 1 See Figure 2-14 for a network drawing. Each building is considered to be a switch block, with two distribution switches each for redundancy. Two switches are used in the core for a dual core design.

Figure 2-14 Network Diagram for Scenario 2-2 Solution



- 2 Because the customer requires the same switch platform for all locations, the Catalyst 5000 series becomes a natural choice. Fast Ethernet media is used from all access layer switches to the end users. Token Ring is used on one floor of Building 2 to support those


users. FDDI can be used to connect the servers to their distribution switches until a media migration is performed. Gigabit Ethernet is used to connect all distribution switches into the core. For MLS, Catalyst 5000s are used at the distribution layer with integrated RSM.

- 3 The enterprise server farm is incorporated as a server block, connected directly into the two core switches. The mainframe, because it is an enterprise resource, is placed in a mainframe block with redundant links into its distribution switches and into the core.

Scenario 2-3 Answers: Large Enterprise Campus Network Design

- 1 This network could be improved by moving toward a hierarchical campus design. A Layer 2 backbone is being used with two switches, though they are not redundant. The distribution layer uses traditional routers, causing a potential bottleneck for all traffic passing into the core. All other enterprise resources connect into the distribution layer routers, probably because the routers inherently support WAN connections and Layer 3 routing.
- 2 See Figure 2-15 for a new network design drawing.



Figure 2-15 Network Diagram for Scenario 2-3 Solution



- 3 The mainframe, because it is an enterprise resource, should become its own switch block. The mainframe should connect into the core for maximum performance and availability.

- 4 The server farm should be migrated into an independent server block and connected directly to the core. Therefore, all users across the campus can access any server efficiently. All users then become the same “distance” from the server farm, regardless of location.
- 5 The WAN should connect to one or more routers set aside for WAN purposes. This WAN block would then connect directly into the core, as other switch blocks do—maximizing availability of WAN access while offering scalability as the WAN grows.
- 6 This network is not making the best use of the routers where they are. The routers should be relocated into a WAN block, as mentioned previously. Layer 3 switching should be implemented in the distribution layer, in place of the routers.
- 7 Evidently, redundancy has not been addressed in this network design. No redundant links exist between access and distribution devices, between distribution and core devices, and between core devices themselves. In addition, the server farm and mainframe have only a single link each into the network.
- 8 Redundancy could be implemented with redundant links between all switch layers. A redundant link should be configured between distribution switches. Redundant links should be present between the WAN, server, and mainframe blocks into the core.
- 9 For this network, a dual core would probably be sufficient. Core connectivity could be provided using Layer 2 switches and redundant links to each switch block. Otherwise, Layer 3 switches could be used in the core, with redundant links between the core switches.

If the access layer is port-bound, a new access layer switch should be added into a switch block. Connecting users, servers, or other resources directly into the distribution layer or into the core layer will work, but this connection defeats the purpose of the building block design—scalability and organization. As more devices are connected into the upper layers of the network design, any sense of scalability is lost and troubleshooting becomes difficult.



This chapter covers the following topics that you will need to master for the CCNP Switching exam:

- **Desktop Connectivity with Ethernet**—This section covers the Ethernet, Fast Ethernet, and Gigabit Ethernet network media technologies.
- **Desktop Connectivity with Token Ring**—The Token Ring LAN media is discussed in detail, along with its use in switched networks.
- **Connecting Switches**—This section discusses the physical cabling and connectivity used with Catalyst switches, including console, Ethernet, and Token Ring interfaces.
- **Switch Management**—This section presents the basic Catalyst switch configuration and administration commands. In addition, this section also covers techniques for interswitch communication.
- **Switch Port Configuration**—This section covers the switch commands that can be used to configure a LAN port for use.

Basic Switch and Port Configuration

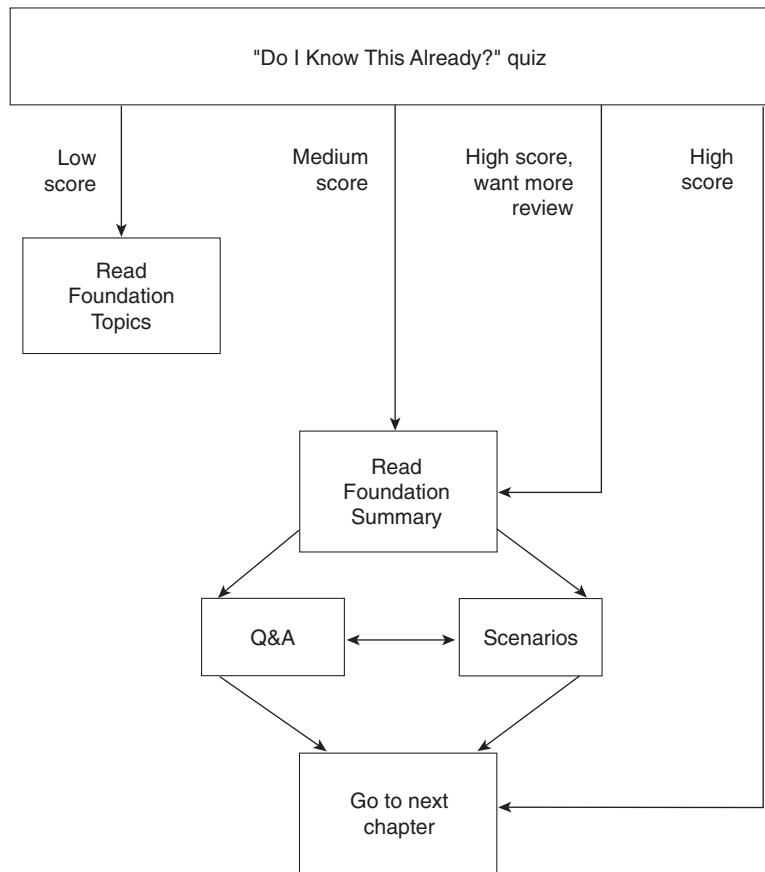
Chapter 2, “Campus Network Design Models,” dealt with the logical processes that can be used to design a campus network. Connections between switch blocks were discussed, such that traffic could be efficiently transported across the campus. Single connections, load balancing, and redundant paths were used to connect switches in modular blocks for complete connectivity. However, these paths were only functional paths—no specifics were presented about how much traffic could be handled, or what physical capabilities were supported. These topics become important when you begin to size traffic loads and actually connect Cisco switch devices.

This chapter presents the various “desktop” network technologies that can be used to establish switched connections within the campus network. As well, you will learn about switch management and the administration commands required to successfully manage switches. Finally, the chapter details the switch commands required for configuring desktop LAN ports.

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and the answers for all your work with this book in one place for easy reference.
- Take the “Do I Know This Already?” quiz, and write down facts and concepts, even if you never look at the information again.
- Use the diagram in Figure 3-1 to guide you to the next step.

Figure 3-1 *How to Use This Chapter*

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The quiz helps you make good choices of how to spend your limited study time. The quiz is sectioned into five smaller “quizlets,” which correspond to the five major headings in the Foundation Topics section of the chapter. Although your answer may differ somewhat from the answers given, finding out if you have the basic understanding of what is presented in this chapter is more important. You will find that these questions are open-ended, rather than multiple choice as found on the exams. Thus, you can focus more on understanding the subject matter than on memorizing details.

Use the scoresheet in Table 3-1 to record your score.

Table 3-1 *Scoresheet for Quiz and Quizlets*

Quizlet Number	Foundation Topics Sections Covering These Questions	Questions	Score
1	Desktop Connectivity with Ethernet	1–6	
2	Desktop Connectivity with Token Ring	7–8	
3	Connecting Switches	9	
4	Switch Management	10–12	
5	Switch Port Configuration	13–14	
All questions		1–14	

1 What are the different Ethernet technologies and their associated IEEE standards?

2 What benefits result with switched Ethernet over shared Ethernet?

3 At what layer are traditional 10 Mbps Ethernet, Fast Ethernet, and Gigabit Ethernet different?

4 Describe Cisco’s EtherChannel technology.

5 In a campus network, where is Fast Ethernet typically used? Where is Gigabit Ethernet typically used?

6 What is the maximum length of a Category 5 100BaseTX cable?

7 Name a type of Token Ring segmentation.

8 What part of a Token Ring frame specifies the exact path the frame should take to reach its destination?

9 What is the purpose of a Gigabit Interface Converter (GBIC)?

10 What must be done to a switch before Telnet access is allowed?

- 11 What type of user interface or command set does the Catalyst 5000 family of switches support? What type is the Catalyst 3500XL?

- 12 What protocol is used by a Catalyst switch to learn about neighboring routers and switches?

- 13 What port speeds can be assigned to a Fast Ethernet switch port?

- 14 What port speeds can be assigned to a Token Ring switch port?

The answers to the quiz are found in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections,” on page 477. The suggested choices for your next step are as follows:

- **7 or less overall score**—Read the entire chapter. The sections include the “Foundation Topics” and “Foundation Summary”, plus the Q&A section at the end of the chapter.
- **8–10 overall score**—Begin with the “Foundation Summary” section, and then go to the Q&A section at the end of the chapter.
- **11 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” section and then go to the Q&A section at the end of the chapter. Otherwise, move to the next chapter.

Foundation Topics

Desktop Connectivity with Ethernet

This section provides a review of the various “flavors” of Ethernet and their application in a campus network. Recall how the bandwidth requirements for each segment of the network are determined by the types of applications in use, the traffic flows within the network, and the size of the user community served. Ethernet scales to support increasing bandwidths, and should be chosen to match the need at each point in the campus network. As network bandwidth requirements grow, the links between access, distribution, and core layers can be scaled to match the load.

Other network media technologies available include Fiber Distribution Data Interface (FDDI), Copper Distribution Data Interface (CDDI), Token Ring, and Asynchronous Transfer Mode (ATM). Although these media are commonly used, Ethernet is emerging as the most popular choice in installed networks. Ethernet is chosen because of its low cost, availability, and scalability to higher bandwidths. Token Ring is discussed later in this chapter, while ATM is covered in Chapter 6, “Trunking with ATM LANE.”

Ethernet

Ethernet is a LAN technology based on the Institute of Electrical and Electronics Engineers (IEEE) 802.3 standard. Ethernet offers a bandwidth of 10 Mbps between end users. In its most basic form, Ethernet is a shared media that becomes both a collision and a broadcast domain. As the number of users on the shared media increases, so does the probability that a user is trying to transmit data at any given time. Ethernet is based on the carrier sense multiple access collision detect (CSMA/CD) technology, which requires that transmitting stations back off for a random period of time when a collision occurs. The more crowded an Ethernet segment becomes, the less efficient it becomes.

Ethernet switching addresses this problem by dynamically allocating a dedicated 10 Mbps bandwidth to each of its ports. The resulting increased network performance occurs by reducing the number of users connected to an Ethernet segment.

Although the principle of switched Ethernet is to offer full dedicated bandwidth to each connected device, assuming that network performance will improve across the board when switching is introduced is a common mistake. For example, consider a workgroup of users connected by a shared media Ethernet hub. These users regularly access an enterprise server located elsewhere in the campus network. To improve performance, the decision is made to replace the hub with an Ethernet switch so that all users get dedicated 10 Mbps connections. Because the switch offers dedicated bandwidth for connections between the end user devices connected to their ports, any user-to-user traffic would probably see improved performance. However, the enterprise server is still located elsewhere in the network, and all of the switched

users must still share available bandwidth across the campus to reach it. As discussed in Chapter 2, rather than throwing raw bandwidth at a problem, a design based on careful observation of traffic patterns and flows will offer a better solution.

Because switched Ethernet can remove the possibility of collisions, stations do not have to listen to each other in order to take a turn transmitting on the wire. Instead, stations can operate in *full-duplex* mode—transmitting and receiving simultaneously. Full-duplex mode further increases network performance, with a net throughput of 10 Mbps in each direction, or 20 Mbps total on each port.

Ethernet cabling involves the use of unshielded twisted-pair (UTP) wiring (10BaseT Ethernet), usually restricted to an end-to-end distance of 100 meters (328 feet) between active devices. Keeping cable lengths as short as possible in the wiring closet will also reduce noise and crosstalk when many cables are bundled together.

In a campus network environment, Ethernet is usually used in the access layer, between end user devices and the access layer switch. Many networks still use Ethernet to connect end users to shared media hubs, which then connect to access layer switches. Ethernet is not typically used at either the distribution or core layer.

NOTE Other cabling technologies are used in Ethernet applications (10Base2, 10Base5, 10BaseF, and so on) though they are not discussed here. For the most part, 10BaseT with UTP wiring is the most commonly used. A useful web site for further reading about Ethernet technology is Charles Spurgeon's Ethernet Web Site, at wwwhost.ots.utexas.edu/ethernet/

Fast Ethernet

Rather than require campuses to invest in a completely new technology to gain increased bandwidth, the networking industry developed a higher-speed Ethernet based on the existing Ethernet standards. Fast Ethernet operates at 100 Mbps and is based on the IEEE 802.3u standard. The Ethernet cabling schemes, CSMA/CD operation, and all upper-layer protocol operations have been maintained with Fast Ethernet. The net result is the same data link Media Access Control (MAC) layer merged with a new physical layer.

In the campus network, Fast Ethernet can be used to link access and distribution layer switches. The larger bandwidth can support the aggregate traffic from multiple Ethernet segments in the access layer. Fast Ethernet can also be used to connect distribution layer switches to the core, with either single or multiple redundant links. It can also be used to connect faster end user workstations to the access layer switch, and to provide improved connectivity to enterprise servers. In other words, Fast Ethernet can be successfully deployed at all layers within a campus network.

Cabling for Fast Ethernet can involve either UTP or fiber. Specifications for Fast Ethernet define the media types and distances as shown in Table 3-2.

Table 3-2 *Cabling Specifications for Fast Ethernet*

Technology	Wiring Type	Pairs	Cable Length
100BaseTX	EIA/TIA Category 5 UTP	2	100 m
100BaseT2	EIA/TIA Category 3,4,5 UTP	2	100 m
100BaseT4	EIA/TIA Category 3,4,5 UTP	4	100 m
100BaseFX	Multimode fiber (MMF) 62.5 micron core, 125 micron outer cladding (62.5/125)	1	400 m half duplex or 2000 m full duplex
	Single-mode fiber (SMF)	1	10 km

Full-Duplex Fast Ethernet

As with traditional Ethernet, the natural progression to improve performance is to use full-duplex operation. Fast Ethernet can provide 100 Mbps in each direction on a switched connection, for 200 Mbps total throughput. This throughput is only possible when a workstation or server is directly connected to a switch port, or when two switches directly connect to each other.

The Fast Ethernet specification also offers backward compatibility to support traditional 10 Mbps Ethernet. To provide this support, two devices at each end of a network connection can automatically negotiate link capabilities so that they both can operate at a maximum common level. This negotiation involves the detection and selection of the highest physical layer technology (available bandwidth) and half-duplex or full-duplex operation. Even if one of the devices uses a fixed configuration, the other device can detect this and match the capabilities.

Autonegotiation uses the priorities shown in Table 3-3 for each mode of Ethernet to determine which technology to agree upon. If both devices can support more than one technology, then the technology with the highest priority will be used. For example, if two devices can support both 10BaseT and 100BaseTX, both devices will use the higher priority 100BaseTX mode.

Table 3-3 *Autonegotiation Selection Priorities*

Priority	Ethernet Mode
7	100BaseT2 (full duplex)
6	100BaseT2 (half duplex)
5	100BaseTX (full duplex)
4	100BaseT4
3	100BaseTX
2	10BaseT (full duplex)
1	10BaseT

NOTE

To assure proper configuration at both ends of a link, Cisco recommends that the appropriate values for transmission speed and duplex mode be manually configured on switch ports.

Cisco provides one additional capability to Fast Ethernet, which allows several Fast Ethernet links to be bundled together for increased throughput. *Fast EtherChannel (FEC)* allows two to eight full-duplex Fast Ethernet links to act as a single physical link, for 400- to 1600-Mbps bandwidth. This technology is described in greater detail in Chapter 5, “Redundant Switch Links.”

For further reading about Fast Ethernet technology, refer to Cisco’s web site: www.cisco.com/warp/public/cc/so/neso/lno/lmnsso/feth_tc.htm

Gigabit Ethernet

Fast Ethernet can be scaled by an additional order of magnitude with the use of Gigabit Ethernet (which supports 1,000 Mbps or 1 Gbps) using the same IEEE 802.3 Ethernet frame format as before. This scalability allows network designers and managers to leverage existing knowledge and technologies to install, migrate, manage, and maintain Gigabit Ethernet networks.

However, the physical layer has been modified to increase data transmission speeds. Two technologies were merged together to gain the benefits of each: The IEEE 802.3 Ethernet standard and the American National Standards Institute (ANSI) X3T11 FibreChannel. IEEE 802.3 provided the foundation of frame format, CSMA/CD, full duplex, and other characteristics of Ethernet. FibreChannel provided a base of high-speed ASICs, optical components, and encoding/decoding and serialization mechanisms. The resulting protocol is termed *IEEE 802.3z Gigabit Ethernet*.

Gigabit Ethernet supports several cabling types, referred to as *1000BaseX*. Table 3-4 lists the cabling specifications for each type.

In a campus network, Gigabit Ethernet can be used in the switch block, the core block, and in the server block. In the switch block, it is used to connect access layer switches to distribution layer switches. In the core, it connects the distribution layer to the core switches, and also interconnects the core devices. For a server block, a Gigabit Ethernet switch in the server block can provide high-speed connections to individual servers.

Table 3-4 *Gigabit Ethernet Cabling and Distance Limitations*

GE Type	Wiring Type	Pairs	Cable Length
1000BaseCX	Shielded Twisted Pair (STP)	1	25 m
1000BaseT	EIA/TIA Category 5 UTP	4	100 m
1000BaseSX	Multimode fiber (MMF) with 62.5 micron core; 850 nm laser	1	275 m
	MMF with 50 micron core; 1300 nm laser	1	550 m
1000BaseLX/LH	MMF with 62.5 micron core; 1300 nm laser	1	550 m
	Single-mode fiber (SMF) with 50 micron core; 1300 nm laser	1	550 m
	SMF with 9 micron core; 1300 nm laser	1	10 km
1000BaseZX	SMF with 9 micron core; 1550 nm laser	1	70 km
	SMF with 8 micron core; 1550 nm laser	1	100 km

Finally, Cisco has extended the concept of Fast EtherChannel to bundle several Gigabit Ethernet links to act as a single physical connection. *Gigabit EtherChannel (GEC)* allows two to eight full-duplex Gigabit Ethernet connections to be aggregated, for up to 16 Gbps throughput. Port aggregation and the EtherChannel technology are described further in Chapter 5.

NOTE The Gigabit Ethernet Alliance offers further reading about Gigabit Ethernet, and its operation, migration, and standards. Refer to the web site: www.gigabit-ethernet.org

Desktop Connectivity with Token Ring

Token Ring is also a LAN technology that provides shared media access to many connected stations. Rather than sharing a common bus or “wire” as Ethernet does, Token Ring stations are arranged in a ring, in a daisy-chain fashion. A token is passed from station to station around the ring, giving the current token holder permission to transmit a frame onto the ring. Once the frame is sent, it is passed around the ring until it is received again by the source. The sending station is responsible for removing the frame from the ring and for introducing a new token to the next neighboring station.

Notice that only one station can transmit at a given time—the one with the token. This restriction prevents a Token Ring network from ever becoming a collision domain. Stations can expect to receive the token at regular intervals as it circulates the ring. This feature makes Token Ring deterministic and useful for delay sensitive protocols. Frames can be sent to a broadcast

MAC address, like Ethernet, causing all stations on the ring to listen. Therefore, a token ring is a broadcast domain.

A Token Ring network offers a bandwidth of 4 Mbps or 16 Mbps. At the higher rate, stations are allowed to introduce a new token as soon as they finish transmitting a frame. This *early token release* increases efficiency by letting more than one station transmit a frame during the original token's round trip. One station is elected to be the *ring monitor*, to provide recovery from runaway frames or tokens. The ring monitor will remove frames that have circled the ring once, if no other station removes them.

Traditional Token Ring networks use *multistation access units (MSAUs)* to provide connectivity between end user stations. MSAUs have several ports that a station can connect to, with either a *B connector* for Type 2 cabling or an *RJ-45 connector* for Category 5 UTP cabling. Internally, the MSAU provides station-to-station connections to form a ring segment. The *Ring-In* and *Ring-Out* connectors of a MSAU can be chained to other MSAUs to form a complete ring topology.

Token Ring Bridging

To form larger networks, Token Rings are interconnected with bridges. Although a *transparent bridge* (or one that forwards frames based solely on MAC addresses) can be used, IBM designed and introduced Token Ring differently. *Source-route bridges* are used to forward frames between rings, based on a predetermined path. The source station includes the exact ring-and-bridge path within the frame so that specific bridges will forward the frame to the appropriate rings. Rings must be uniquely numbered and identified with the campus network, with a number between 1 and 4095. Bridges, however, do not have to be unique across the network, as long as two bridges with the same number do not connect to the same ring. Bridges are numbered 1 through 15.

The steps to determine the path a frame should take are as follows:

- Step 1** The source station first sends a test frame to see if the destination is on the local ring. If the destination responds, the source knows that it is local. If there is no response, the source station will send an *all routes explorer (ARE)* frame, which will cause all bridges to forward the frame to all rings.
- Step 2** Within the frame is a *routing information field (RIF)*. The RIF carries a record of bridges and rings traversed along the way. As the frame is forwarded, each bridge will append its bridge number and the next ring number to the RIF in the frame.
- Step 3** The destination will then reply to each ARE frame it receives, so that the source will receive a confirmation of every possible path to the destination.

For future transmissions, the source can choose the path it thinks is best (quickest response, least number of bridge hops, largest maximum transmission unit (MTU), combination of factors, and so forth). These frames will contain the exact path desired by the source station, in the form of a RIF.

In certain scenarios, hybrid bridging can be provided. Sometimes, both *source-route bridging* (SRB) and transparent bridging must occur. Here, *source-route transparent bridging* (SRT) forwards a frame according to a RIF, if present, or according to MAC address tables if it finds no RIF.

As in Ethernet switching, Token Rings can also be segmented by dividing a ring across several switch ports. While this feature increases the available bandwidth on a ring segment, it requires more in-depth forwarding decisions. Token ring switching, or more properly termed *source-route switching*, forwards frames according to a combination of MAC addresses and RIF contents.

Source-route switching differs from other forms of bridging in that it only looks at the RIF and never updates or adds to the RIF. Instead, the switch learns *route descriptors*, or the ring/bridge combinations that specify the next-hop destinations from incoming frames. The source-route switch then associates the route descriptors and MAC addresses (if needed) with outbound ports closest to the destination. When subsequent frames are received on other ports, the route descriptor is quickly indexed to lookup the outbound port.

In this fashion, source-route switching supports parallel source-route paths to destinations. The number of MAC addresses to be learned is lessened, because route descriptors point to the next-hop ports. The actual operation of source-route switching is much like virtual LANs with Ethernet. For this reason, further discussion of source-route switching is presented in Chapter 4, “VLANs and Trunking.”

Table 3-5 summarizes the attributes of each type of Token Ring connectivity and segmentation method.

Table 3-5 *Token Ring Segmentation Methods*

Method	Forwarding Decision	Frame Modification	Ring Numbering
Transparent bridging	MAC address		N/A
Source-route bridging	RIF	RIF	Ring numbers must be unique among bridge ports.
Source-route transparent bridging	MAC address or RIF	RIF	Ring numbers must be unique among bridge ports.
Source-route switching	Route descriptor		Ring numbers can be same across switch ports (single ring can be segmented on several ports).

Connecting Switches

Switch deployment in a network involves two steps: physical connectivity and switch configuration. This section describes the connections and cabling requirements for devices in a switch block. Cable connections must be made to the console port of a switch in order to make initial configurations. Physical connectivity between switches and end users involves cabling for the various types of LAN ports.

Console Port Cables/Connectors

A terminal emulation program on a PC is usually required to interface with the console port on a switch. Various types of console cables and console connectors are associated with each Cisco switch family.

All Catalyst switch families use an RJ-45-to-RJ-45 *rollover cable* to make the console connection between a PC (or terminal or modem) and the console port. A rollover cable is made so that pin 1 on one RJ-45 connector goes to pin 8 on the other RJ-45 connector, pin 2 goes to pin 7, and so forth. In other words, the cable remains flat while the two RJ-45 connectors point in opposite directions.

To connect the PC end, the rollover cable plugs into an RJ-45 to DB-9 or DB-25 “Terminal” adapter (or a DB-25 “Modem” adapter for a modem connection). At the switch end, the rollover cable plugs directly into the RJ-45 jack of the console port. This situation is true for the Catalyst 1900, 2820, 2900, 3500, 2926G, 2948G, 4912G, 5000 Supervisor IIG/III/IIIG, and the 6000 switches.

On the Catalyst 4003, 5000 Supervisor I/II, and the 8500 switches, the rollover cable must connect to an RJ-45 to DB-25 “Modem” adapter. These switches have a DB-25 console port connector that is a female DCE.

Once the console port is cabled to the PC, terminal, or modem, a terminal emulation program can be started or a user connection can be made. The console ports on all switch families require an asynchronous serial connection at 9600 baud, 8 data bits, no parity, 1 stop bit, and no flow control.

Ethernet Port Cables/Connectors

Catalyst switches support a variety of network connections, including all forms of Ethernet. In addition, Catalyst switches support several types of cabling, including UTP and optical fiber.

On Catalyst 1900 and 2820 series switches, the Ethernet ports are fixed-speed with 12 or 24 10BaseT and one or two 100BaseTX or 100BaseFX ports. The 10BaseT ports can be connected only to other 10BaseT-capable devices (including 10/100 autosensing devices), and the 100BaseX to other 100BaseX-capable devices. The 10BaseT and 100BaseTX ports use Category 5 UTP cabling and RJ-45 connectors.

The 100BaseFX ports use two-strand multimedia fiber (MMF) with SC connectors to provide connectivity. The SC connectors on the fiber cables are square in shape. These connectors snap in and out of the switch port connector as the connector is pushed in or pulled out. One fiber strand is used as a transmit path and the other as a receive path. Therefore, the transmit fiber on one switch device should connect to the receive fiber on the other end.

The remainder of the Catalyst switch families support 10/100 autosensing (using Fast Ethernet autonegotiation) and Gigabit Ethernet. Switched 10/100 ports use RJ-45 connectors on Category 5 UTP cabling to complete the connections. These ports can be connected to other 10BaseT, 100BaseTX, or 10/100 autosensing devices. UTP cabling is arranged so that RJ-45 pins 1,2 and 3,6 form two twisted pairs. These pairs are connected straight through to the far end.

In order to connect two 10/100 switch ports back-to-back, as in an access layer to distribution layer link, a Category 5 UTP *crossover cable* must be used. In this case, RJ-45 pins 1,2 and 3,6 are still twisted pairs, but 1,2 on one end connect to 3,6 on the other end, and 3,6 on one end connect to 1,2 on the other end.

NOTE

Because UTP Ethernet connections use only pairs 1,2 and 3,6, some cable plant installers only connect these pairs and leave the remaining two pair positions empty. While this move provides Ethernet connectivity, it is not good practice for future needs. Instead, all four pairs of the RJ-45 connector should be connected end-to-end. For example, a full four-pair UTP cable plant can be used for either Ethernet or Token Ring connectivity, without rewiring. (Token Ring UTP connections use pairs 3,6 and 4,5.) Also, to be compatible with the new IEEE 802.3ab standard for Gigabit Ethernet over copper, all four pairs must be used end-to-end.

Gigabit Ethernet Port Cables/Connectors

Gigabit Ethernet connections take a different approach by providing modular connectivity options. Catalyst switches with Gigabit Ethernet ports have standardized rectangular openings that accept Gigabit Interface Converters (GBICs). GBIC modules provide the media personality for the port so that various types of cables can be connected. In this way, the switch chassis is completely modular and requires no major change to accept a new media type. Instead, the appropriate GBIC module is hot-swappable and is plugged into the switch to support the new media. GBICs are available for the following Gigabit Ethernet media:


- **1000BaseSX GBIC**—short wavelength connectivity using SC fiber connectors and MMF for distances up to 550 meters (1804 feet).
- **1000BaseLX/LH GBIC**—long wavelength/long haul connectivity using SC fiber connectors and either MMF or single-mode fiber (SMF); MMF can be used for distances up to 550 meters (1804 feet) and SMF can be used for distances up to 10 km (32,810 feet).

- **1000BaseZX GBIC**—extended distance connectivity using SC fiber connectors and SMF; can be used for distances up to 70 km and even to 100 km when used with premium grade SMF.
- **GigaStack GBIC**—uses a proprietary connector with a high-data-rate copper cable with enhanced signal integrity and electromagnetic interference (EMI) performance; provides a GBIC-to-GBIC connection between stacking Catalyst switches or between any two Gigabit switch ports over a short distance.

CAUTION The fiber-based GBICs always have the receive fiber on the left SC connector and the transmit fiber on the right SC connector. These GBICs could produce invisible laser radiation from the transmit SC connector. Therefore, always keep unused SC connectors covered with the rubber plugs and do not look directly into the SC connectors.

Figure 3-2 illustrates a fiber-based GBIC module and how one is installed in a Gigabit Ethernet switch port.

Figure 3-2 *Gigabit Interface Converter*



Token Ring Port Cables/Connectors

Catalyst switches support UTP Token Ring connections. These ports operate at either 4 or 16 Mbps, in several half and full-duplex modes. RJ-45 connectors on Category 5 UTP cabling use twisted pairs 3,6 and 4,5. These pairs are connected straight through to the far end.

Switch Management

Cisco Catalyst switch devices can be configured to support many different requirements and features. When a PC is connected to the serial console port, configuration is generally done with a terminal emulator application on the PC. Further configurations can be performed through a Telnet session across the LAN or through a web-based interface. These topics will be covered in later sections.

Catalyst switches support one of two types of user interface for configuration: Cisco IOS-based commands, and **set**-based, command-line interface (CLI) commands. The IOS-based commands (found in Catalyst 1900/2820, 2900XL, and 3500XL) are similar to many IOS commands used on Cisco routers. However, the CLI commands (found in 2926G, 4000, 5000 and 6000) use **set** and **clear** commands to change configuration parameters. Both types of user interface are discussed in the sections that follow.

Identifying the Switch

All switches come from the factory with a default configuration and a default system name or prompt. This name can be changed so that each switch in a campus network will have a unique identity. This option can be useful when you are using Telnet to move from switch to switch in a network.

Setting the Hostname/System Name on an IOS-Based Switch

To change the host or system name on an IOS-based user interface, enter the following command in configuration mode:

```
Switch(config)# hostname hostname
```

The hostname is a string of 1 to 255 alphanumeric characters. As soon as this command is executed, the system prompt will change to reflect the new hostname.

NOTE Configuration changes made on IOS-based switches apply only to the active *running configuration*, stored in RAM. To make the changes permanent, in effect even after a power cycle, remember to copy the switch configuration into the *startup configuration*, stored in NVRAM. You can do this by using the **copy running-config startup-config** command.

Setting the Hostname/System Name on a CLI-Based Switch

To set the system name on a CLI-based user interface, the system prompt is changed with the following command:

```
Switch(enable) set system name name-string
```

As soon as this command is executed, the system name and the prompt will change to reflect the new value. This prompt is displayed at the beginning of every CLI line.

Passwords and User Access

Normally, a network device should be configured to secure it from unauthorized access. Catalyst switches offer a simple form of security by setting passwords to restrict who can log in to the user interface. Two levels of user access are available: regular login, or *EXEC mode*, and enable login, or *privileged mode*. EXEC mode is the first level of access, which gives access to the basic user interface through any line or the console port. The privileged mode requires a second password and gives access to set or change switch operating parameters or configurations.

Cisco provides various methods for providing device security and user authentication. Many of these methods are more secure and robust than using the login passwords in Chapter 12, “Controlling Access in the Campus Environment,” describes these features in greater detail.

Setting Login Passwords on an IOS-Based Switch

To set the login passwords on a Cisco IOS-based switch interface, enter the following commands in global configuration mode:

```
Switch(config)# enable password level 1 password
Switch(config)# enable password level 15 password
```

Here, the EXEC mode password is set with a privilege level of one (1), while the enable password is set with a privilege level of 15. The password is a string of four to eight alphanumeric characters. Passwords on these switches are not case-sensitive.

To remove a password, use the **no enable password level password** command.

Setting Login Passwords on a CLI-Based Switch

Example 3-1 lists the commands you would enter in enable mode to set the login passwords on a Cisco switch with a CLI-based user interface.

Example 3-1 *Setting the Login Passwords on a Cisco Switch*

```
Switch (enable) set password
Enter old password: oldpassword
Enter new password: newpassword
Retype new password: newpassword
Password changed.
Switch (enable) set enablepass
Enter old password: oldenablepassword
Enter new password: newenablepassword
Retype new password: newenablepassword
Password changed.
Switch (enable)
```

As Example 3-1 demonstrates, “**password**” is the EXEC mode password, and the “**enablepass**” is the privileged mode password. Passwords on these switches *are* case-sensitive.

Remote Access

By default, the switch login passwords allow user access only via the console port. In order to use Telnet to access a switch from within the campus network, to use **ping** to test the reachability of a switch, or to monitor a switch by SNMP, you must perform some configuration for remote access.

Although a switch operates at Layer 2, the switch supervisor processor must maintain an IP stack at Layer 3 for administrative purposes. An IP address and subnet mask can then be assigned to the switch so that remote communications with the switch supervisor are possible.

By default, all ports on a switch are assigned to the same virtual LAN (VLAN) or broadcast domain. The switch supervisor and its IP stack must be assigned to a VLAN before remote Telnet and **ping** sessions will be supported. VLANs are discussed further in Chapter 4.

Enabling Remote Access on an IOS-Based Switch

On a switch with an IOS-based user interface, an IP address can be assigned to the management VLAN (default is VLAN 1) with the following commands in global configuration mode:

```
Switch(config)# interface vlan 1
Switch(config-if)# ip address ip-address netmask
Switch(config-if)# ip default-gateway ip-address
```

As demonstrated by the preceding command syntax, an IP address and subnet mask are assigned to the VLAN1 “interface,” which is really the switch supervisor’s IP stack listening on VLAN1. In order to send packets destined off the local VLAN1 subnet, a default gateway IP address is also assigned.

Again, this default gateway has nothing to do with processing packets that are passed through the switch; rather, the default gateway is only used to forward traffic between a user and the switch supervisor for management purposes.

To view the current switch IP settings, use the **show ip** command.

Enabling Remote Access on a CLI-Based Switch

An IP address can also be configured for in-band management on a switch with a CLI-based user interface by entering the following commands in privileged mode:

```
Switch(enable) set interface sc0 ip-address netmask broadcast-address
Switch(enable) set interface sc0 vlan
Switch(enable) set ip route default gateway
```

The first command line defines the IP address and subnet mask for the switch management interface, **sc0**. The broadcast address must also be given to match the subnet and subnet mask values. In addition, the management interface is assigned to a specific VLAN with the second command line. If this command is not given, the management interface defaults to VLAN1. The third command line assigns a default gateway that will receive any packets destined off the local management interface subnet.

To view the current IP settings, use the **show interface** command.

Communicating Between Switches

Because switch devices are usually interconnected, management is usually simplified if the switches can communicate on some level to become aware of each other. Cisco has implemented protocols on its devices so that neighboring Cisco equipment can be found. As well, some families of switch devices can be clustered and managed as a unit once they discover one another.

Cisco Discovery Protocol

Cisco uses a proprietary protocol on both switches and routers to discover neighboring devices. The Cisco Discovery Protocol (CDP) can be enabled on interfaces to periodically advertise the existence of a device and exchange basic information with directly connected neighbors. The information exchanged in CDP messages includes the device type, links between devices, and the number of ports within each device.

By default, CDP runs on each port of a Cisco switch that is capable of using the SNAP protocol. CDP advertisements occur every 60 seconds by default. CDP communication occurs at the data link layer so that it is independent of any network layer protocol that may be running on a network segment. CDP frames are sent as multicasts, using a destination MAC address of 01:00:0c:cc:cc:cc.

Switches regard the CDP address as a special address designating a multicast frame that should not be forwarded. Instead, CDP multicast frames are redirected to the switch's management port, and are processed by the switch supervisor alone. Therefore, Cisco switches only become aware of other directly connected Cisco devices.

Enabling CDP and Viewing CDP Information on an IOS-Based Switch

CDP is enabled by default on all switch interfaces. To enable CDP, use the following interface configuration command (use the **no** form to disable CDP):

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

To view information learned from CDP advertisements of neighboring Cisco devices, use one of the following commands:

```
Switch# show cdp interface [type module/port]
Switch# show cdp neighbors [type module/port] [detail]
```

The first command displays CDP information pertaining to a specific interface. If the type, module, and port information is omitted, CDP information from all interfaces is listed. The second command displays CDP information about neighboring Cisco devices. If the **detail** keyword is used, all possible CDP information about each neighbor is displayed.

Enabling CDP and Viewing CDP Information on a CLI-Based Switch

CDP is enabled by default. To enable or disable CDP, use the following command:

```
Switch(enable) set cdp {enable | disable} module/port
```

The *module* and *port* parameters are included to enable or disable CDP on individual ports. If these values are excluded, CDP is enabled or disabled on a global basis for all ports on the switch.

To view information learned from CDP advertisements of neighboring Cisco devices, use a form of the following command:

```
Switch(enable) show cdp neighbors [module/port] [vlan | duplex | capabilities | detail]
```

Here, the module and port number can be given to view CDP information on a particular port. The **vlan** keyword displays information about the native VLAN numbers of neighboring devices. The **duplex** keyword displays the duplex type of each neighboring device. Using **capabilities** displays capability codes for the neighboring devices. The **detail** keyword displays all possible CDP information about each neighboring device, including the IP address assigned to the neighboring interface or management interface.

As demonstrated in Example 3-2, the **show cdp neighbors detail** command can be useful when you are connected to a switch and need to know more about what other switches are nearby in a network. Particularly useful are the IP address entries, allowing Telnet access to previously unknown switches.

Example 3-2 Displaying CDP Information for Neighboring Devices

```
Switch(enable) show cdp neighbors 4/4 detail
Port (Our Port):4/4
Device-ID:69046406
Device Addresses:
  IP Address:172.20.25.161
Holdtime:150 sec
Capabilities:TRANSPARENT_BRIDGE SWITCH
Version:
  WS-C5509 Software, Version McpSW: 5.3(0.29)BOU NmpSW: 5.3(0.29)BOU
  Copyright (c) 1995-1999 by Cisco Systems
```


Example 3-2 *Displaying CDP Information for Neighboring Devices (Continued)*

```
Port-ID (Port on Device):4/8
Platform:WS-C6009
VTP Management Domain:unknown
Native VLAN:1
Duplex:half
Switch(enable)
```

For a quick summary of CDP status on all switch ports, use the **show cdp port** command.

Switch Clustering and Stacking

Cisco has also implemented a proprietary method for grouping switches into a management cluster. Up to 16 switch devices can be added into a cluster, regardless of their physical location on the network. In this fashion, an entire cluster of switches can be managed through a single IP address—that of the *command switch*. Cluster management can be performed through HTML, IOS-based, and SNMP-based management interfaces on the command switch.

Cluster discovery takes place once a command switch has been assigned an IP address and configured as a command switch. CDP messages are used to discover neighboring switches that are candidates for cluster membership. Cluster discovery takes place only on switch ports that are assigned and connected to VLAN1. Only the directly connected switch devices will be discovered by the command switch. Other switches daisy-chained behind the directly connected neighbors can be manually added to the cluster.

NOTE

At press time, only the Catalyst 2900 and 3500 switch families (both IOS-based) are capable of cluster operations.

To configure a switch to become the command switch for a cluster, first assign an IP address for the management interface. Then, use the following command:

```
Switch(config)# cluster enable cluster-name
```

Once the command switch has been identified and configured, the cluster discovery can be viewed and managed from a web browser. Refer to the *Cluster Builder* documentation in the Catalyst 2900XL and 3500XL software documentation for further detailed information and examples. (www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29_35xp/scg/kiclust.htm)

Switch Port Configuration

The individual ports on a switch can be configured with various information and settings, as detailed in the following sections.

Identifying Ports

Switch ports can have a text description added to their configuration to help identify them. This description is meant as a comment field only, as a record of port use or other unique information. The port description is shown when the switch configuration is displayed.

Assigning a Port Description on an IOS-Based Switch

To assign a comment or description to an interface on a switch with an IOS-based user interface, enter the following command in interface configuration mode:

```
Switch(config-if)# description description-string
```

If the description string has embedded spaces between words, the entire string must be enclosed between quotation marks. To remove a description, use the **no description** interface configuration command.

Assigning a Port Description on a CLI-Based Switch

For a switch with a CLI-based user interface, assign a port description with:

```
Switch(enable) set port name module/number description-string
```

Here, *module* is the switch module number where the port resides, and *number* is the port number on that module. The description string must be less than 21 characters, and can have embedded spaces with no special treatment. To remove a port description, use the **set port name** *module/number* command, followed by a carriage return (no description string).

Port Speed

Switch ports can be assigned a specific speed through switch configuration commands. Ethernet ports can be set to speeds of *10*, *100*, and *Auto* for autonegotiate mode. Gigabit Ethernet ports are always set to a speed of *1000*. Token Ring ports can be set to speeds of *4*, *16*, and *Auto* for autosensing mode.

NOTE

If a 10/100 Fast Ethernet port is assigned a speed of **auto**, both its speed and duplex mode will be negotiated.

Assigning Port Speed on an IOS-Based Switch

To specify the port speed on a particular Ethernet port, use the following interface configuration command:

```
Switch(config-if)# speed {10 | 100 | auto}
```

Assigning Port Speed on an CLI-Based Switch

On a CLI-based switch, set the port speed with the following command:

```
Switch(enable) set port speed module/number {10 | 100 | auto}  
Switch(enable) set port speed module/number {4 | 16 | auto}
```

The first line applies to Ethernet or Fast Ethernet ports, while the second line applies to Token Ring ports.

Ethernet Port Mode

Ethernet-based switch ports can also be assigned a specific link mode. Therefore, the port operates in half-duplex, full-duplex, or autonegotiated mode. Autonegotiation is only allowed on Fast Ethernet and Gigabit Ethernet ports. In this mode, full-duplex operation will be attempted first, and then half duplex if full duplex was not successful. The autonegotiation process repeats whenever the link status changes.

NOTE A 10-Mbps Ethernet link defaults to half duplex, while a 100-Mbps Fast Ethernet link defaults to full duplex.

Assigning the Ethernet Link Mode on an IOS-Based Switch

To set the link mode on an IOS-based switch port, enter the following command in interface configuration mode:

```
Switch(config-if)# duplex {auto | full | half}
```

If the port is not automatically enabled or activated, use the **no shutdown** interface configuration command. To view the current speed and duplex state of a port, use the **show interface** command.

Assigning the Ethernet Link Mode on a CLI-Based Switch

To set the link mode on a CLI-based switch port, enter the following command:

```
Switch(enable) set port duplex module/number {full | half}
```

If the port is not automatically enabled or activated, use the **set port enable** command. To view the current speed and duplex status of a port, use the **show port** command.

Token Ring Port Mode

Token Ring ports have five modes of operation:

- **Half-duplex concentrator port (hdxcpport)**—The port is connected to a single station in half-duplex mode, similar to a MAU connection.
- **Half-duplex station emulation (hdxstation)**—The port is connected to a media attachment unit (MAU) port, like a regular station.
- **Full-duplex concentrator port (fdxcport)**—The port is connected to a full-duplex station.
- **Full-duplex station emulation (fdxstation)**—The port is connected to another full-duplex Token Ring.
- **Autosensing (auto)**—The port will autosense the operating mode of the connected device or ring.

Assigning the Token Ring Link Mode on a CLI-Based Switch

To set the Token Ring link mode on a CLI-based switch port, enter the following command:

```
Switch(enable) set tokenring portmode module/number {auto | fdxcport | hdxcpport |  
fdxstation | hdxstation}
```

Foundation Summary

The Foundation Summary is a collection of tables and figures that provides a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final prep before the exam, the following tables and figures will hopefully be a convenient way to review the day before the exam.

Table 3-6 *Cabling Specifications for Fast Ethernet*

Technology	Wiring Type	Pairs	Cable Length
100BaseTX	EIA/TIA Category 5 UTP	2	100 m
100BaseT2	EIA/TIA Category 3,4,5 UTP	2	100 m
100BaseT4	EIA/TIA Category 3,4,5 UTP	4	100 m
100BaseFX	Multimode fiber (MMF) 62.5 micron core, 125 micron outer cladding (62.5/125)	1	400 m half duplex or 2000 m full duplex
	Single-mode fiber (SMF)	1	10 km

Table 3-7 *Autonegotiation Selection Priorities*

Priority	Ethernet Mode
7	100BaseT2 (full duplex)
6	100BaseT2 (half duplex)
5	100BaseTX (full duplex)
4	100BaseT4
3	100BaseTX
2	10BaseT (full duplex)
1	10BaseT

Table 3-8 *Gigabit Ethernet Cabling and Distance Limitations*

GE Type	Wiring Type	Pairs	Cable Length
1000BaseCX	Shielded twisted-pair (STP)	1	25 m
1000BaseT	EIA/TIA Category 5 UTP	4	100 m
1000BaseSX	MMF with 62.5 micron core; 850 nm laser	1	275 m
	MMF with 50 micron core; 1300 nm laser	1	550 m
1000BaseLX/LH	MMF with 62.5 micron core; 1300 nm laser	1	550 m
	SMF with 50 micron core; 1300 nm laser	1	550 m
	SMF with 9 micron core; 1300 nm laser	1	10 km
1000BaseZX	SMF with 9 micron core; 1550 nm laser	1	70 km
	SMF with 8 micron core; 1550 nm laser	1	100 km

Table 3-9 *Token Ring Segmentation Methods*

Method	Forwarding Decision	Frame Modification	Ring Numbering
Transparent bridging	MAC address		N/A
Source-route bridging	RIF	RIF	Ring numbers must be unique among bridge ports
Source-route transparent	MAC address or RIF	RIF	Ring numbers must be unique among bridge ports
Source-route switching	Route descriptor		Ring numbers can be same across switch ports (single ring can be segmented on several ports)

Table 3-10 *Switch Management Configuration Commands*

Task	IOS-Based Command	CLI-Based Command
Identify Switch	hostname <i>hostname</i>	set system name <i>name-string</i>
Set EXEC level password	enable password level 1 <i>password</i>	set password
Set privileged level password	enable password level 15 <i>password</i>	set enablepass
Set IP address	interface vlan 1 ip address <i>ip-address netmask</i> ip default-gateway <i>ip-address</i>	set interface sc0 <i>ip-address netmask broadcast-address</i> set interface sc0 <i>vlan</i> set ip route default <i>gateway</i>
CDP	cdp enable show cdp interface [<i>type module/port</i>] show cdp neighbors [<i>type module/port</i>] [detail]	set cdp { enable disable } <i>module/port</i> show cdp neighbors [<i>module/port</i>] [vlan duplex capabilities detail]
Enable cluster	cluster enable <i>cluster-name</i>	N/A

Table 3-11 *Switch Port Configuration Commands*

Task	IOS-Based Command	CLI-Based Command
Identify port	description <i>description-string</i>	set port name <i>module/number description-string</i>
Set port speed (Ethernet)	speed { 10 100 auto }	set port speed <i>module/number</i> { 10 100 auto }
Set port speed (Token Ring)	N/A	set port speed <i>module/number</i> { 4 16 auto }
Set port mode (Ethernet)	duplex { auto full half }	set port duplex <i>module/number</i> { full half }
Set port mode (Token Ring)	N/A	set tokenring portmode <i>module/number</i> { auto fdxport hdxport fdxstation hdxstation }

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however, they are designed to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to these questions can be found in Appendix A, on page 477.

- 1 What are the different Ethernet technologies and their associated IEEE standards?

- 2 What benefits result with switched Ethernet over shared Ethernet?

- 3 When a 10/100 Ethernet link is autonegotiating, which will be chosen if both stations can support the same capabilities—10BaseT full duplex, 100BaseTX half duplex, or 100BaseTX full duplex?

- 4 At what layer are traditional 10-Mbps Ethernet, Fast Ethernet, and Gigabit Ethernet different?

5 Describe Cisco's EtherChannel technology.

6 A switch port is being configured as shown below. What command is needed next to set the port to full-duplex mode?

```
Switch(config)# interface FastEthernet 0/13
Switch(config-if)#
```

7 In a campus network, where is Fast Ethernet typically used? Where is Gigabit Ethernet typically used?

8 What is the maximum length of a Category 5 100BaseTX cable?

9 A CLI-based switch port has been configured for 100 Mbps full-duplex mode, but a link cannot be established. What are some commands that could be used to investigate and correct the problem?

10 Name a type of Token Ring segmentation.

11 What part of a Token Ring frame specifies the exact path the frame should take to reach its destination?

12 What switch command will set the enable-mode password on an IOS-based switch? A CLI-based switch?

13 What is the purpose of a GBIC?

14 What CLI-based commands will allow Telnet and **ping** access to a switch management interface at 192.168.200.10, subnet mask 255.255.255.0, on VLAN 5? Now add a command to allow access between the switch and devices located off the local VLAN 5 subnet, using a router at 192.168.200.1.

15 What must be done to a switch before Telnet access is allowed?

16 What factors determine the choice of a distribution layer switch, its access-to-distribution layer link media, and its Layer 3 processor?

- 17** What type of user interface or command set does the Catalyst 5000 family of switches support? What type is the Catalyst 3500XL?

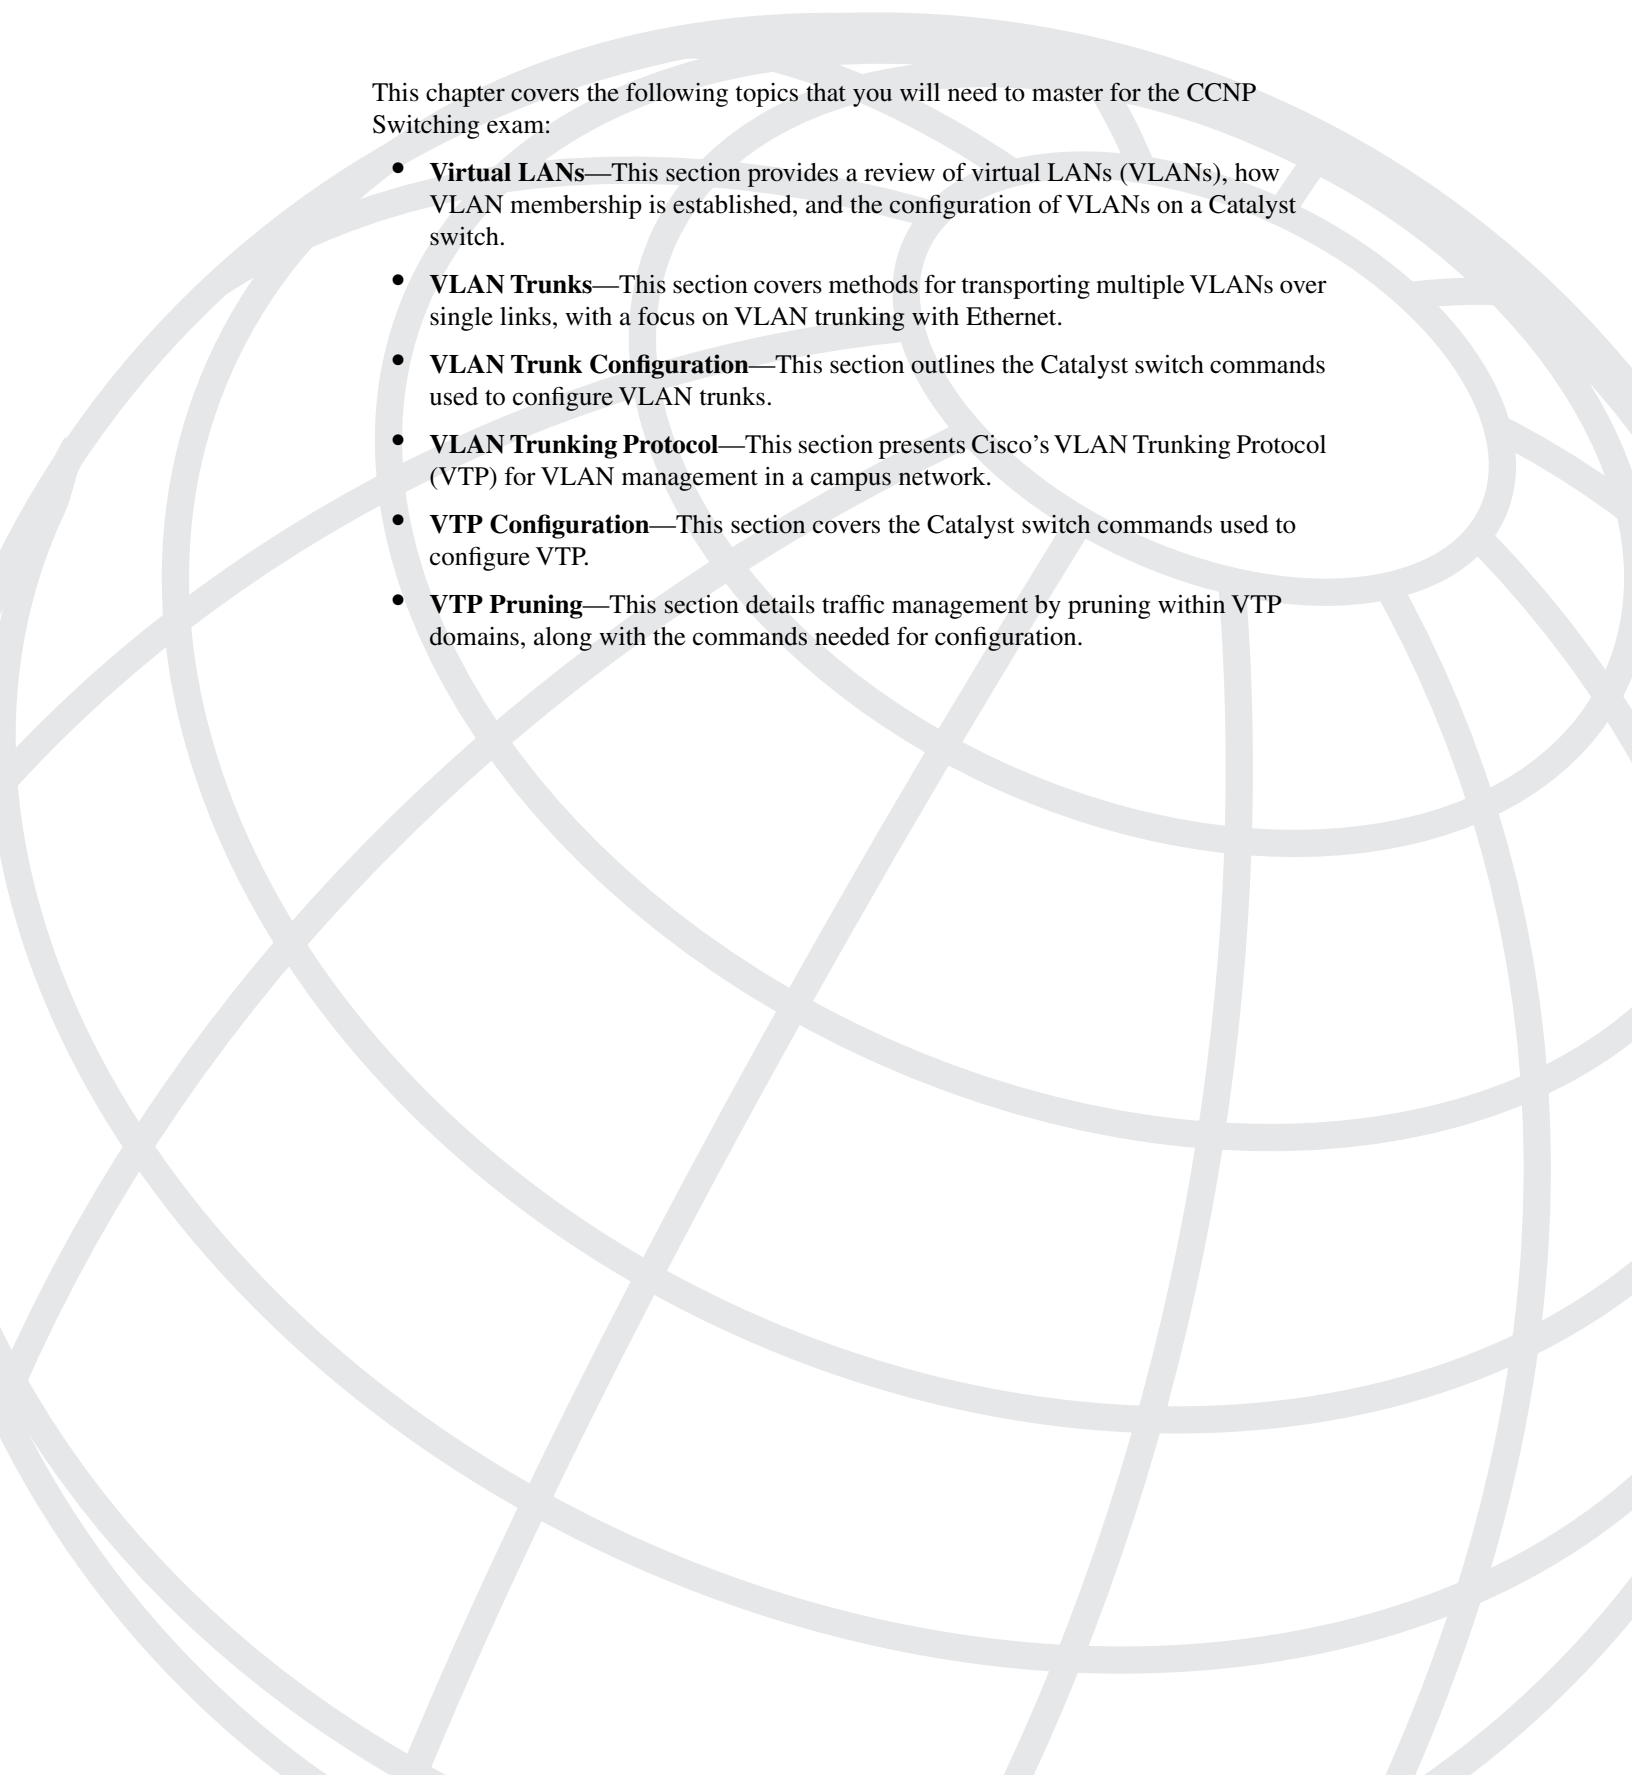

- 18** What protocol is used by a Catalyst switch to learn about neighboring routers and switches?

- 19** What switch command can be used to find the IP addresses of nearby Cisco switches on a network?

- 20** What port speeds can be assigned to a Fast Ethernet switch port?

- 21** What is the purpose of switch clustering? Can clustered switches share switching loads with each other?

- 22** What port speeds can be assigned to a Token Ring switch port?



This chapter covers the following topics that you will need to master for the CCNP Switching exam:

- **Virtual LANs**—This section provides a review of virtual LANs (VLANs), how VLAN membership is established, and the configuration of VLANs on a Catalyst switch.
- **VLAN Trunks**—This section covers methods for transporting multiple VLANs over single links, with a focus on VLAN trunking with Ethernet.
- **VLAN Trunk Configuration**—This section outlines the Catalyst switch commands used to configure VLAN trunks.
- **VLAN Trunking Protocol**—This section presents Cisco's VLAN Trunking Protocol (VTP) for VLAN management in a campus network.
- **VTP Configuration**—This section covers the Catalyst switch commands used to configure VTP.
- **VTP Pruning**—This section details traffic management by pruning within VTP domains, along with the commands needed for configuration.

VLANs and Trunking

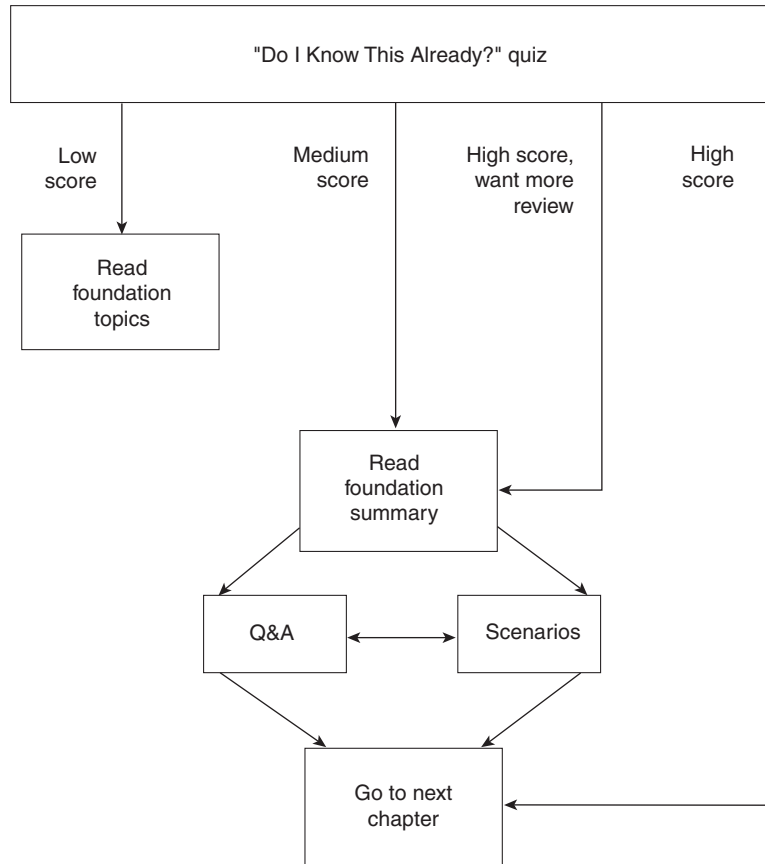
Switched campus networks can be broken up into distinct broadcast domains or virtual LANs (VLANs). A flat network topology, or a network with a single broadcast domain, can be simple to implement and manage. However, flat network topology is not scalable. Instead, the campus can be divided into segments using VLANs, while Layer 3 routing protocols manage interVLAN communication.

This chapter details the process of defining common workgroups within a group of switches. Switch configuration for VLANs is covered, along with the method of identifying and transporting VLANs on various types of links. VLAN administration and management is presented through the configuration of the VLAN Trunking Protocol (VTP).

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and answers for all your work with this book in one place, for easy reference.
- Take the “Do I Know This Already?” quiz and write down facts and concepts (even if you never look at the information again).
- Use the diagram in Figure 4-1 to guide you to the next step.

Figure 4-1 *How to Use This Chapter*

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The quiz helps you make good choices of how to spend your limited study time. The quiz is sectioned into five smaller “quizlets” that correspond to the five major headings in the “Foundation Topics” section of the chapter. Although your answer may differ somewhat from the answers given, finding out if you have the basic understanding that is presented in this chapter is most important. You will find that these questions are open-ended rather than multiple choice as found on the exams. This is done to focus more on understanding the subject matter than on memorizing details.

Use the scoresheet in Table 4-1 to record your score.

Table 4-1 *Scoresheet for Quiz and Quizlets*

Quizlet Number	Foundation Topics Sections Covering These Questions	Questions	Score
1	Virtual LANs	1–4	
2	VLAN Trunks VLAN Trunk Configuration	5–7	
3	VTP VTP Configuration	8–10	
4	VTP Pruning	11–12	
All questions		1–12	

1 What is a VLAN? When is it used?

2 What are two types of VLANs, in terms of spanning areas of the campus network?

3 Generally speaking, what must be configured (both switch and end user device) for a port-based VLAN?

4 What are the components of a Token Ring VLAN?

5 What is a trunk link?

6 What methods of VLAN frame identification can be used on a Catalyst switch?

7 What is the purpose of Dynamic Trunking Protocol (DTP)?

8 What VTP modes can a Catalyst switch be configured for? Can VLANs be created in each of the modes?

9 How many VTP management domains can a Catalyst switch participate in? How many VTP servers can a management domain have?

10 What conditions must exist for two Catalyst switches to be in the same VTP management domains?

11 What is the purpose of VTP pruning?

12 Which VLAN numbers are never eligible for VTP pruning? Why?

The answers to the quiz are found in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections,” on page 477. The suggested choices for your next step are as follows:

- **6 or less overall score**—Read the entire chapter. This reading includes the “Foundation Topics” and “Foundation Summary” sections, the Q&A section, and the scenarios at the end of the chapter.
- **7–9 overall score**—Begin with the “Foundation Summary” section and then follow with the Q&A section and the scenarios at the end of the chapter.
- **10 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” section and then go to the Q&A section and the scenarios at the end of the chapter. Otherwise, move on to the next chapter.

Foundation Topics

Virtual LANs

Consider a network design that consists of Layer 2 devices only. For example, this design could be a single Ethernet segment, an Ethernet switch with many ports, or a network with several interconnected Ethernet switches. A fully Layer 2 switched network is referred to as a *flat network topology*. A flat network is a single broadcast domain, such that every connected device sees every broadcast packet that is transmitted. As the number of stations on the network increases, so does the number of broadcasts.


Due to the Layer 2 foundation, flat networks cannot contain redundant paths for load balancing or fault tolerance. The reason for this is explained in Chapter 5, “Redundant Switch Links.” To gain any advantage from additional paths to a destination, Layer 3 routing functions must be introduced.

A switched environment offers the technology to overcome flat network limitations. Switched networks can be subdivided into virtual LANs (VLANs). By definition, a VLAN is a single broadcast domain. All devices connected to the VLAN receive broadcasts from other VLAN members. However, devices connected to a different VLAN will not receive those same broadcasts.

A VLAN is made up of defined members communicating as a logical network segment. In contrast, a physical segment consists of devices that must be connected to a physical cable segment. A VLAN can have connected members located anywhere in the campus network, as long as VLAN connectivity is provided between all members. Layer 2 switches are configured with a VLAN mapping and provide the logical connectivity between the VLAN members.

Figure 4-2 shows how a VLAN can provide logical connectivity between switch ports.

Two workstations on the left Catalyst switch are assigned to VLAN 1, while a third workstation is assigned to VLAN 100. In this example, there can be no communication between VLAN 1 and VLAN 100. Both ends of the link between the Catalysts are assigned to VLAN 1. One workstation on the right Catalyst is also assigned to VLAN 1. Because there is end-to-end connectivity of VLAN 1, any of the workstations on VLAN 1 can communicate as if they were connected to a physical network segment.

Figure 4-2 VLAN Functionality

VLAN Membership

When a VLAN is provided at an access layer switch, an end user must have some means to gain membership to it. Two membership methods exist on Cisco Catalyst switches: static VLANs and dynamic VLANs.

Static VLANs

Static VLANs offer *port-based* membership, where switch ports are assigned to specific VLANs. End user devices become members in a VLAN based on which physical switch port they are connected to. No handshaking or unique VLAN membership protocol is needed for the end devices; they automatically assume VLAN connectivity when they connect to a port. Normally, the end device is not even aware that the VLAN exists. The switch port and its VLAN are simply viewed and used as any other network segment, with other “locally attached” members on the wire.

Switch ports are assigned to VLANs by the manual intervention of the network administrator, hence the static nature. The ports on a single switch can be assigned and grouped into many VLANs. Even though two devices are connected to the same switch, traffic will not pass between them if they are connected to ports on different VLANs. To perform this function, either a Layer 3 device could be used to route packets or an external Layer 2 device could be used to bridge packets between the two VLANs.

The static port-to-VLAN membership is normally handled in hardware with application-specific integrated circuits (ASICs) in the switch. This membership provides good performance because all port mappings are done at the hardware level with no complex table lookups needed.

Configuring Static VLANs

This section describes the switch commands needed to configure static VLANs. By default, all switch ports are assigned to VLAN 1, are set to be a VLAN type of Ethernet, have a maximum transmission unit (MTU) size of 1500 bytes, and have a Security Association Identifier (SAID) of 100,000 plus the VLAN number.

First, the VLAN must be created on the switch, if it doesn't already exist. Then the VLAN must be assigned to specific switch ports.

NOTE

To create a new VLAN, several prerequisites relating to VTP must be met. The switch must be assigned to a VTP domain and be configured for either *server* or *transparent* VTP mode. VTP is covered in the “VLAN Trunking Protocol” section of this chapter.

To configure static VLANs on an IOS-based switch, you would enter the following commands in enable mode:

```
Switch# vlan database
Switch(vlan)# vlan vlan-num name vlan-name
Switch(vlan)# exit
Switch# configure terminal
Switch(config)# interface interface module/number
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan vlan-num
Switch(config-if)# end
```

The VLAN is created and stored in a database, along with its number and name. To assign a switch port to the VLAN, you would use the **switchport access vlan** interface configuration command. The **switchport mode access** command configures the port for static VLAN membership.

To configure static VLANs on a CLI-based switch, you would enter the following commands in enable mode:

```
Switch(enable) set vlan vlan-num [name name]
Switch(enable) set vlan vlan-num mod-num/port-list
```

The first command creates the VLAN numbered *vlan-num* on the switch and assigns a descriptive name to it. Note that a VLAN and its number are significant only on the local switch, unless some form of VLAN trunking is used to communicate with other switches. If the *name* field is not specified, the switch will create a name based on the VLAN number, in the form of

VLAN0002 for VLAN 2 for example. The second command assigns VLAN *vlan-num* to one or more switch ports, identified with the switch module number and the list of port numbers. For example, the command **set vlan 101 3/1,3-7** would assign ports 3/1, 3/3, 3/4, 3/5, 3/6, and 3/7 to VLAN 101.

To verify VLAN configuration, using the **show vlan** command will output a list of all VLANs defined in the switch, in addition to the ports assigned to each VLAN.

Dynamic VLANs

Dynamic VLANs are used to provide membership based on the MAC address of an end user device. When a device is connected to a switch port, the switch must query a database to establish VLAN membership. A network administrator must assign the user's MAC address to a VLAN in the database of a VLAN Membership Policy Server (VMPS).

With Cisco switches, dynamic VLANs are created and managed through the use of network management tools like CiscoWorks 2000 or CiscoWorks for Switched Internetworks (CWSI). Dynamic VLANs allow a great deal of flexibility and mobility for end users, but require more administrative overhead.

NOTE

Dynamic VLANs are not covered in this text. For more information, refer to the following Cisco resources:

- CLI-based switches: "Configuring Dynamic Port VLAN Membership with VMPS" at www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_5_5/sw_cfg/vmps.htm
 - IOS-based switches: "How VMPS Works" at www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29_35xu/scg/kivlan.htm#xtocid2442355
-

Extent of VLANs

To implement VLANs, you must give some consideration to the number of VLANs you'll need and how best to place them. As usual, the number of VLANs will be dependent on traffic patterns, application types, segmenting common workgroups, and network management requirements.

However, an important factor to consider is the relationship between VLANs and the IP addressing schemes used. Cisco recommends a one-to-one correspondence between VLANs and IP subnets. This recommendation means that if a Class C network address is used for a VLAN, then no more than 254 devices should be in the VLAN. As well, VLANs not extending beyond the Layer 2 domain of the distribution switch is recommended. In other words, the VLAN should not reach across the core of a network and into another switch block. The idea again is to keep broadcasts and unnecessary movement of traffic out of the core block.

VLANs can be scaled in the switch block by using two basic methods: end-to-end VLANs and local VLANs.

End-to-End VLANs

End-to-end VLANs, also called campus-wide VLANs, span the entire switch fabric of a network. They are positioned to support maximum flexibility and mobility of end devices. Users are assigned to VLANs regardless of physical location. As a user moves around the campus, that user's VLAN membership stays the same. This means that each VLAN must be made available at the access layer in every switch block.

End-to-end VLANs should group users according to common requirements. All users in a VLAN should have roughly the same traffic flow patterns, following the 80/20 rule. Recall that this rule estimates that 80 percent of user traffic stays within the local workgroup, while 20 percent is destined for a remote resource in the campus network. Although only 20 percent of the traffic in a VLAN is expected to cross the network core, end-to-end VLANs make it possible for *all* traffic within a single VLAN to cross the core.

Because all VLANs must be available at each access layer switch, VLAN trunking must be used to carry all VLANs between the access and distribution layer switches. (Trunking is discussed in later sections of this chapter.)

Local VLANs

Because most enterprise networks have moved toward the 20/80 rule (where server and intranet/Internet resources are centralized), end-to-end VLANs have become cumbersome and difficult to maintain. The 20/80 rule is reversed—only 20 percent of traffic is local, while 80 percent is destined to a remote resource across the core layer. End users require access to central resources outside their VLAN. Users must cross into the network core more frequently. In this type of network, VLANs are designed to contain user communities based on geographic boundaries, with little regard to the amount of traffic leaving the VLAN.

Local or geographic VLANs range in size from a single switch in a wiring closet to an entire building. Arranging VLANs in this fashion enables the Layer 3 function in the campus network to intelligently handle the inter-VLAN traffic loads. This scenario provides maximum availability by using multiple paths to destinations, maximum scalability by keeping the VLAN within a switch block, and maximum manageability.

VLAN Trunks

At the access layer, end user devices connect to switch ports that provide simple connectivity to a single VLAN each. The attached devices are unaware of any VLAN structure and simply attach to what appears to be a normal physical network segment. Remember, sending information from an access link on one VLAN to another VLAN is not possible without the intervention of an additional device—either a Layer 3 router or an external Layer 2 bridge.

NOTE Note that a switch port can support more than one IP subnet for the devices attached to it. For example, consider a shared Ethernet hub that is connected to a single Ethernet switch port. One user device on the hub may be configured for 192.168.1.1 255.255.255.0, while another is assigned 192.168.17.1 255.255.255.0. Although these subnets are unique communicating on one switch port, they cannot be considered separate VLANs. The switch port supports one VLAN, but multiple subnets can exist on that single VLAN.


A *trunk link*, however, can transport more than one VLAN through a single switch port. Trunk links are most beneficial when switches are connected to other switches or switches are connected to routers.

A trunk link is not assigned to a specific VLAN. Instead, one, many, or all active VLANs can be transported between switches using a single physical trunk link. Connecting two switches with separate physical links for each VLAN is possible. Figure 4-3 shows how two switches might be connected in this fashion.

As VLANs are added to a network, the number of links can quickly grow. A more efficient use of physical interfaces and cabling involves the use of trunking. The right half of the figure shows how one trunk link can replace many individual VLAN links. A trunk link can be associated with a native VLAN, which is used if the trunk link fails for some reason.

Cisco supports trunking on both Fast Ethernet and Gigabit Ethernet switch links, as well as aggregated Fast and Gigabit EtherChannel links. To distinguish between traffic belonging to different VLANs on a trunk link, the switch must have a method of identifying each frame with the appropriate VLAN. Several identification methods are available and are discussed in the next section.

Figure 4-3 *Passing VLAN Traffic Using Single Links Versus Trunk Links*



VLAN Frame Identification

Because a trunk link can be used to transport many VLANs, a switch must identify frames with their VLANs as they are sent and received over a trunk link. *Frame identification*, or *tagging*, assigns a unique user-defined ID to each frame transported on a trunk link. This ID can be thought of as the VLAN number or VLAN “color,” as if each VLAN was drawn on a network diagram in a unique color.

VLAN frame identification was developed for switched networks. As each frame is transmitted over a trunk link, a unique identifier is placed in the frame header. As each switch along the way receives these frames, the identifier is examined to determine to which VLAN the frames belong.

If frames must be transported out another trunk link, the VLAN identifier is retained in the frame header. Otherwise if frames are destined out an access link, the switch removes the VLAN identifier before transmitting the frames to the end station. Therefore, all traces of VLAN association are hidden from the end station.


VLAN identification can be performed using several methods. Each uses a different frame identifier mechanism, and some are suited for specific network media. These methods are described in the sections that follow.

Inter-Switch Link Protocol

The Inter-Switch Link (ISL) protocol is a Cisco proprietary method for preserving the source VLAN identification of frames passing over a trunk link. ISL performs frame identification in Layer 2 by encapsulating each frame between a header and trailer. Any Cisco switch or router device configured for ISL can process and understand the ISL VLAN information. ISL is primarily used for Ethernet media, although Cisco has included provisions to carry Token Ring, FDDI, and ATM frames over Ethernet ISL. (A frame-type field in the ISL header indicates the source frame type.)

When a frame is destined out a trunk link to another switch or router, ISL adds a 26-byte header and a 4-byte trailer to the frame. The source VLAN is identified with a 10-bit VLAN ID in the header. The trailer contains a cyclic redundancy check (CRC) to assure the data integrity of the new encapsulated frame. Figure 4-4 shows how Ethernet frames are encapsulated and forwarded out a trunk link. Because tagging information is added at the beginning and end of each frame, ISL is sometimes referred to as *double tagging*.

Figure 4-4 ISL Frame Identification



If a frame is destined for an access link, the ISL encapsulation (both header and trailer) is removed before transmission. This removal preserves ISL information only for trunk links and devices that can understand the protocol.

IEEE 802.1Q Protocol

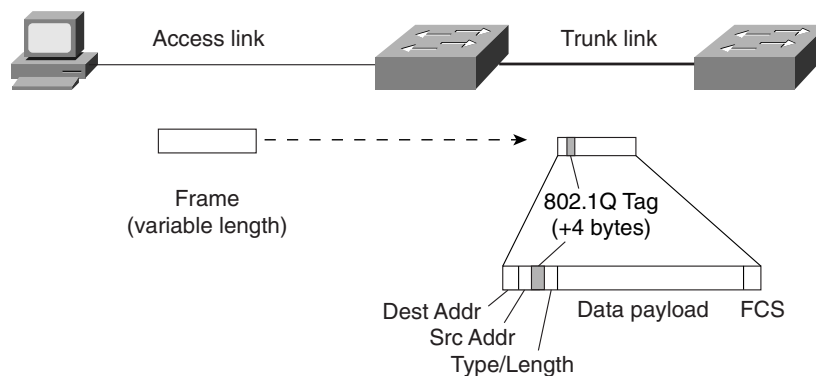
The IEEE 802.1Q protocol can also be used to preserve VLAN associations over trunk links. However, this frame identification method is standardized, allowing VLAN trunks to exist and operate between equipment from multiple vendors.

In particular, the IEEE 802.1Q standard defines an architecture for VLAN use, services provided with VLANs, and protocols and algorithms used to provide VLAN services. Further information about the 802.1Q standard can be found at grouper.ieee.org/groups/802/1/pages/802.1Q.html

Like Cisco ISL, IEEE 802.1Q can be used for VLAN identification with Ethernet trunks. Instead of encapsulating each frame with a VLAN ID header and trailer, 802.1Q embeds its tagging information within the Layer 2 frame. This method is referred to as *single-tagging* or *internal tagging*. 802.1Q also introduces the concept of a *native VLAN* on a trunk. Frames belonging to this VLAN are not encapsulated with tagging information. In the event that an end station is connected to an 802.1Q trunk link, the end station will be able to receive and understand only the native VLAN frames.

In an Ethernet frame, 802.1Q adds a four-byte tag just after the source address field, as shown in Figure 4-5.

Figure 4-5 IEEE 802.1Q Frame Tagging Standard



The first two bytes are used as a Tag Protocol Identifier (TPID). The first two bytes also always have a value of 0x8100 to signify an 802.1Q tag. The remaining two bytes are used as a Tag Control Information (TCI) field. The TCI information contains a 3-bit priority field, which is used to implement class of service functions in the accompanying 802.1Q/802.1p prioritization standard. One bit of the TCI is a Canonical Format Indicator (CFI), flagging whether the MAC addresses are in canonical format. The last 12 bits are used as a VLAN Identifier (VID) to indicate the source VLAN for the frame. The VID can have values from 0 to 4095, but VLAN 0, 1, and 4095 are reserved.

NOTE Note that both ISL and 802.1Q tagging methods have one implication: they add to the length of an Ethernet frame. ISL adds a total of 30 bytes to each frame, while 802.1Q adds 4 bytes. Because Ethernet frames cannot exceed 1518 bytes, the additional VLAN tagging information can cause the frame to be too large. Frames that barely exceed the MTU size are called *baby giant frames*. Switches will usually report these frames as Ethernet errors or oversize frames.

LAN Emulation (LANE)

Trunking VLANs between switches over an Asynchronous Transfer Mode (ATM) link is possible. Here, VLANs are transported using the IEEE LAN Emulation (LANE) standard. LANE is discussed in greater detail in Chapter 6, “Trunking with ATM LANE.”

IEEE 802.10

Cisco offers a proprietary method for transporting VLAN information inside the standard IEEE 802.10 FDDI frame. The VLAN information is carried in the Security Association Identifier (SAID) field of the 802.10 frame.

Dynamic Trunking Protocol

Trunk links on Catalyst switches can be manually configured for either ISL or 802.1Q mode. However, Cisco has implemented a proprietary point-to-point protocol called Dynamic Trunking Protocol (DTP) that will negotiate a common trunking mode between two switches. DTP is available in Catalyst supervisor engine software Release 4.2 and later. DTP negotiation should be disabled if a switch has a trunk link connected to a router because the router cannot participate in the DTP negotiation protocol.

NOTE A trunk link can be negotiated between two switches only if both switches belong to the same VLAN Trunking Protocol (VTP) management domain. VTP is discussed in the “VTP Configuration” section of this chapter. If the two switches are in different VTP domains and trunking is desired between them, the trunk links must be set to on or nonegotiate mode. This setting will force the trunk to be established. These options are explained in the next section.

VLAN Trunk Configuration

By default, all switch ports are non-trunking and operate as access links until some intervention changes the mode. The sections that follow demonstrate the commands necessary to configure VLAN trunks on both an IOS-based and CLI-based switch.

VLAN Trunk Configuration on an IOS-Based Switch

Use the following commands to create a VLAN trunk link on an IOS-based switch:

```
Switch(config)# interface interface mod/port
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport trunk encapsulation {isl | dot1q}
Switch(config-if)# switchport trunk allowed vlan remove vlan-list
Switch(config-if)# switchport trunk allowed vlan add vlan-list
```

Individually, these commands place the switch port into trunking mode, using the encapsulation specified as either **isl** or **dot1q**. The last two commands define which VLANs can be trunked over the link. A list of VLANs is first removed from the trunk because all VLANs (1–1005) are trunked by default. Then, a list of VLANs can be added back into the trunk.

To view the trunking status on a switch port, use the **show interface int mod/port switchport** command.

VLAN Trunk Configuration on a CLI-Based Switch

To create a VLAN trunk link, use the **set trunk** CLI-based command. This command sets the trunking mode and any mode negotiation. The **set trunk** command also identifies the VLANs that will be transported over the trunk link. Trunk configuration uses the following command syntax:

```
Switch(enable) set trunk module/port [on | off | desirable | auto | nonegotiate]
vlan-range [isl | dot1q | dot10 | lane | negotiate]
```

Here, the trunk link is identified by its physical location as the switch module number and port number. The trunking mode can be set to any of the following:

- **on**—This setting places the port in permanent trunking mode. The corresponding switch port at the other end of the trunk should be similarly configured because negotiation is not allowed. The encapsulation or identification mode should also be manually configured.
- **off**—This setting places the port in permanent non-trunking mode. The port will attempt to convert the link to non-trunking mode.
- **desirable**—Selecting this port will actively attempt to convert the link into trunking mode. If the far end switch port is configured to **on**, **desirable**, or **auto** mode, trunking will be successfully negotiated.
- **auto**—The port will be willing to convert the link into trunking mode. If the far end switch port is configured to **on** or **desirable**, trunking will be negotiated. By default, all Fast Ethernet and Gigabit Ethernet links that are capable of negotiating using DTP are configured to this mode. Because of the passive negotiation behavior, the link will never become a trunk, if both ends of the link are left to the **auto** default.
- **nonegotiate**—The port is placed in permanent trunking mode, but no DTP frames are generated for negotiation. The far end switch port must be manually configured for trunking mode.

NOTE Note that in all modes except **nonegotiate**, DTP frames are sent out every 30 seconds to keep neighboring switch ports informed of the link’s mode. On critical trunk links in a network, manually configuring the trunking mode on both ends is best so that the link can never be negotiated to any other state.

By default, a switch will transport all VLANs (1–1000) over a trunk link, even if a VLAN range is specified in the **set trunk** command. There might be times when the trunk link should not carry all VLANs. For example, broadcasts are forwarded to every switch port on a VLAN—including the trunk link because it, too, is a member of the VLAN. If the VLAN doesn’t extend past the far end of the trunk link, propagating broadcasts across the trunk makes no sense.

Therefore, to remove VLANs from a trunk link, use the following command:

```
Switch(enable) clear trunk module/port vlan-range
```

Then, if VLANs need to be added back to the trunk, they can be specified as the *vlan-range* in the **set trunk** command.

Lastly, the trunk encapsulation or identification mode is specified at the end of the **set trunk** command. These values are

- **isl**—VLANs are tagged by encapsulating each frame using the Cisco ISL protocol. This protocol is the default, if no value is specified.
- **dot1q**—VLANs are tagged in each frame using the IEEE 802.1Q standard protocol.
- **dot10**—VLANs are tagged on an FDDI switch port using the IEEE 802.10 protocol.
- **lane**—VLANs are identified on an ATM link using LAN Emulation.
- **negotiate**—On Fast and Gigabit Ethernet ports, the mode will be negotiated to select either ISL or IEEE 802.1Q. ISL is preferred, unless one end of the link is configured for **dot1q**.

To view and verify the trunk configuration on a switch, use the **show trunk** [*module/port*] command. Example 4-1 shows a sample output of trunk information.

Example 4-1 **show trunk** Verifies Trunk Configuration on a Switch

```
Switch> (enable) show trunk
```

Port	Mode	Encapsulation	Status	Native vlan
2/1	auto	dot1q	trunking	1
3/1	auto	isl	trunking	1
3/2	desirable	isl	trunking	1

continues

Example 4-1 show trunk Verifies Trunk Configuration on a Switch (Continued)

```

Port      Vlans allowed on trunk
-----
2/1      1-1000
3/1      1-1000
3/2      1-1000
Port      Vlans allowed and active in management domain
-----
2/1      1-10,20-35,100,201
3/1      1,11-19,100,201
3/2      1,11,15,100,201
Port      Vlans in spanning tree forwarding state and not pruned
-----
2/1      1-10,20-35,100,201
3/1      1000
3/2      1000
Switch> (enable)

```

VLAN Trunking Protocol

As the previous sections have shown, VLAN configuration and trunking on a switch or a small group of switches is fairly easy and straightforward. Campus network environments, however, are usually made up of many interconnected switches. Configuring and managing a large number of switches, VLANs, and VLAN trunks can quickly get out of hand.

Cisco has developed a method to manage VLANs across the campus network. The VLAN Trunking Protocol (VTP) uses Layer 2 trunk frames to communicate VLAN information among a group of switches. VTP manages the addition, deletion, and renaming of VLANs across the network from a central point of control.

VTP Domains

VTP is organized into *management domains* or areas with common VLAN requirements. A switch can belong to only one VTP domain, in addition to sharing VLAN information with other switches in the domain. Similar to VLANs, switches in different VTP domains do not share VTP information.

Switches in a VTP domain advertise several attributes to their domain neighbors. Each advertisement contains information about the VTP management domain, VTP revision number, known VLANs, and specific VLAN parameters. When a VLAN is added to a switch in a management domain, other switches are notified of the new VLAN through VTP advertisements. In this way, all switches in a domain can prepare to receive traffic on their trunk ports using the new VLAN.

VTP Modes

To participate in a VTP management domain, each switch must be configured to operate in one of several modes. The VTP mode will determine how the switch processes and advertises VTP information. The following modes can be used:

- **Server mode**—VTP servers have full control over VLAN creation and modification for their domains. All VTP information is advertised to other switches in the domain, while all received VTP information is synchronized with the other switches. By default, a switch is in VTP server mode. Note that each VTP domain must have at least one server so that VLANs can be created, modified, or deleted, and so that VLAN information can be propagated.
- **Client mode**—VTP clients do not allow the administrator to create, change, or delete any VLANs. Instead, they listen to VTP advertisements from other switches and modify their VLAN configurations accordingly. In effect, this is a passive listening mode. Received VTP information is forwarded out trunk links to neighboring switches in the domain.
- **Transparent mode**—VTP transparent switches do not participate in VTP. While in transparent mode, a switch does not advertise its own VLAN configuration, and a switch does not synchronize its VLAN database with received advertisements. As well, in VTP version 1, a transparent mode switch does not even relay VTP information it receives to other switches. In VTP version 2, transparent switches do forward received VTP advertisements out of their trunk ports, acting as VTP relays.

NOTE

While a switch is in VTP transparent mode, a switch can create and delete VLANs that are local to itself. These VLAN changes, however, will not be propagated to any other switch.

VTP Advertisements

Each switch participating in VTP advertises VLANs, revision numbers, and VLAN parameters on its trunk ports to notify other switches in the management domain. VTP advertisements are sent as multicast frames. The switch intercepts frames sent to the VTP multicast address and processes them with its supervisory processor. VTP frames are forwarded out trunk links as a special case.

Because all switches in a management domain learn of new VLAN configuration changes, a VLAN need only be created and configured on just one VTP server switch in the domain.

By default, management domains are set to use non-secure advertisements without a password. A password can be added to set the domain to secure mode. The same password has to be configured on every switch in the domain so that all switches exchanging VTP information will use identical encryption methods.

The VTP advertisement process starts with configuration revision number 0 (*zero*). When subsequent changes are made, the revision number is incremented before advertisements are sent out. When listening switches receive an advertisement with a greater revision number than is locally stored, the advertisement will overwrite any stored VLAN information. Because of this, forcing any newly added network switches to have revision number zero is important. The VTP revision number is stored in NVRAM and is not altered by a power cycle of the switch. Therefore, the revision number can only be initialized to zero using one of the following methods:

- Change the VTP mode of the switch to *transparent* and then change the mode back to *server*.
- Change the VTP domain of the switch to a bogus name (a non-existent VTP domain) and then change the VTP domain back to the original name.
- Issue a **clear config all** command, which will clear the switch configuration *and* the VTP information stored in NVRAM. Power cycle the switch so that it boots up with a non-existent VTP domain name and a VTP revision number of zero. (*Use caution. This is the most drastic method because it will erase all configuration data.*)

If the VTP revision number is not reset to zero, a new server switch might advertise VLANs as non-existent or deleted. If the advertised revision number happens to be greater than previous legitimate advertisements, listening switches would overwrite good VLAN database entries with null or deleted VLAN status information. This is referred to as a *VTP synchronization problem*.

Advertisements can originate as requests from client-mode switches that want to learn about the VTP database at boot-up time. As well, advertisements can originate from server-mode switches as VLAN configuration changes occur.

VTP advertisements can occur in three forms:

- **Summary advertisements**—VTP domain servers will send summary advertisements every 300 seconds and every time a VLAN topology change occurs. The summary advertisement lists information about the management domain, including VTP version, domain name, configuration revision number, timestamp, MD5 encryption hash code, and the number of subset advertisements to follow. For VLAN configuration changes, summary advertisements are followed by one or more subset advertisements, with more specific VLAN configuration data. Figure 4-6 shows the summary advertisement format.

Figure 4-6 VTP Summary Advertisement Format

Version (1 byte)	Type (Summary Adv) (1 byte)	Number of subset advertisements to follow (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision Number (4 bytes)			
Updater Identity (originating IP address: 4 bytes)			
Update Timestamp (12 bytes)			
MD5 Digest hash code (16 bytes)			

- **Subset advertisements**—VTP domain servers will send subset advertisements after a VLAN configuration change occurs. These advertisements list the specific changes that have been performed, such as creation or deletion of a VLAN, suspending or activating a VLAN, changing the name of a VLAN, and changing the MTU of a VLAN. Subset advertisements can list the following VLAN parameters: status of the VLAN, VLAN type (like Ethernet or Token Ring), MTU, length of the VLAN name, VLAN number, SAID value, and the VLAN name. VLANs are listed individually in sequential subset advertisements. Figure 4-7 shows the VTP subset advertisement format.
- **Advertisement requests from clients**—A VTP client can request any lacking VLAN information. For example, a client switch might be reset and have its VLAN database cleared, its VTP domain membership might be changed, or it might hear a VTP summary advertisement with a higher revision number than it currently has. After a client advertisement request, the VTP domain servers respond with summary and subset advertisements. Figure 4-8 shows the advertisement request format.

Figure 4-7 VTP Subset Advertisement and VLAN Info Field Formats

VTP Subset Advertisement

0	1	2	3
Version (1 byte)	Type (Subset Adv) (1 byte)	Subset sequence number (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision Number (4 bytes)			
VLAN Info Field 1 (see below)			
VLAN Info Field ...			
VLAN Info Field N			

VTP VLAN Info Field

0	1	2	3
Info Length	VLAN Status	VLAN Type	VLAN Name Length
ISL VLAN ID		MTU Size	
802.10 SAID			
VLAN Name (padded with zeros to multiple of 4 bytes)			

Figure 4-8 VTP Advertisement Request Format

0	1	2	3
Version (1 byte)	Type (Adv request) (1 byte)	Reserved (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Starting advertisement to request			

Catalyst switches in server mode use a separate nonvolatile random-access memory (NVRAM) for VTP, different from the configuration NVRAM. All VTP information, including the VTP configuration revision number, is retained even when the switch power is off. In this manner, a switch is able to recover the last known VLAN configuration from its VTP database once it reboots.

VTP Configuration

Before VLANs can be configured, VTP must be configured. By default, every switch will operate in VTP server mode for the management domain *NULL*, with no password or secure mode. The following sections discuss the commands and considerations that should be used to configure a switch for VTP operation.

Configuring a VTP Management Domain

Before a switch is added into a network, the VTP management domain should be identified. If this switch is the first one on the network, the management domain will need to be created. Otherwise, the switch may have to join an existing management domain with other existing switches.

Configuring a VTP Management Domain on an IOS-Based Switch

The following command can be used to assign a switch to a management domain, where the *domain-name* is a text string up to 32 characters long.

```
Switch# vlan database
Switch(vlan)# vtp domain domain-name
```

Configuring a VTP Management Domain on a CLI-Based Switch

Similar to the command to assign a switch to a management domain on an IOS-based switch, the following command does the same for a CLI-based switch:

```
Switch(enable) set vtp [domain domain-name]
```

Configuring the VTP Mode

Next, the VTP mode needs to be chosen for the new switch. The three VTP modes of operation and their guidelines for use are as follows:

- **Server mode**—Server mode can be used on any switch in a management domain, even if other server and client switches are in use. This mode provides some redundancy in the event of a server failure in the domain. However, each VTP management domain must have at least one server. The first server defined in a network also defines the management domain that will be used by future VTP servers and clients. Server mode is the default VTP mode.
- **Client mode**—If other switches are in the management domain, a new switch should be configured for client mode operation. In this way, the switch will learn any existing VTP information from a server.

If this switch will be used as a redundant server, it should start out in client mode to learn all VTP information from reliable sources. If the switch was initially configured for server mode instead, it might propagate incorrect information to the other domain switches. Once the switch has learned the current VTP information, it can be reconfigured for server mode.

- **Transparent mode**—This mode is used if a switch is not going to share VLAN information with any other switch in the network. VLANs can still be created, deleted, and renamed on the transparent switch. However, they will not be advertised to other neighboring switches. VTP advertisements received by a transparent switch will be forwarded on to other switches on trunk links.
- Keeping switches in transparent mode can eliminate the chance for duplicate, overlapping VLANs in a large network with many network administrators. For example, two administrators might configure VLANs on switches in their respective areas, but use the same VLAN identification or VLAN number. Even though the two VLANs have different meanings and purposes, they could overlap if both administrators advertised them using VTP.

Configuring the VTP Mode on an IOS-Based Switch

On an IOS-based switch, the VTP mode can be configured with the following sequence of commands:

```
Switch# vlan database
Switch(vlan)# vtp domain domain-name
Switch(vlan)# vtp {server | client | transparent}
Switch(vlan)# vtp password password
```

Configuring the VTP Mode on a CLI-Based Switch

On a CLI-based switch, the VTP mode can be configured with the following command:

```
Switch(enable) set vtp [domain domain-name] [mode {server | client | transparent}]
[passwd password]
```

If the domain is operating in secure mode, a password can be included in the command line. The password must be a string of 8 to 64 characters.

Configuring the VTP Version

Two versions of VTP are available for use in a management domain. Catalyst switches are capable of running either VTP version 1 or VTP version 2. Within a management domain, the two versions are not interoperable. Therefore, the same VTP version must be configured on each switch in a domain. VTP version 1 is the default protocol on a switch.

If a switch is capable of running VTP version 2, however, a switch may coexist with other version 1 switches, as long as its VTP version 2 is not enabled. This situation becomes important if you want to use version 2 in a domain. Then, only one server mode switch needs to have VTP version 2 enabled. The new version number is propagated to all other version 2-capable switches in the domain, causing them all to automatically enable version 2 for use.

By default, VTP version 1 is enabled. Version 2 can be enabled or disabled using the **v2** option.

The two versions of VTP differ in the features they support. VTP version 2 offers the following additional features over version 1:

- **Version-dependent transparent mode**—VTP version 1 in transparent mode matches the VTP version and domain name before forwarding the information to other switches using VTP. VTP version 2 in transparent mode differs by forwarding the VTP messages without checking the version number. Because only one domain is supported in a switch, the domain name doesn't have to be checked.
- **Consistency checks**—VTP version 2 performs consistency checks on the VTP and VLAN parameters entered from the CLI or by Simple Network Management Protocol (SNMP). This checking helps prevent errors in such things as VLAN names and numbers from being propagated to other switches in the domain. However, no consistency checks are performed on VTP messages that are received on trunk links or on configuration and database data that is read from NVRAM.
- **Token Ring support**—VTP version 2 supports the use of Token Ring switching and Token Ring VLANs. (If Token Ring switching is being used, VTP version 2 must be enabled.) The section “Token Ring VLANs” later in this chapter discusses these topics in detail.
- **Unrecognized Type-Length-Value (TLV) support**—VTP version 2 switches will propagate received configuration change messages out other trunk links, even if the switch supervisor is not able to parse or understand the message. For example, a VTP advertisement contains a *Type* field to denote what type of VTP message is being sent. VTP message type 1 is a summary advertisement, and message type 2 is a subset advertisement. An extension to VTP could be in use that utilizes other message types and other message length values. Instead of dropping the unrecognized VTP message, version 2 will still propagate the information and keep a copy in NVRAM.

Configuring the VTP Version on an IOS-Based Switch

On an IOS-based switch, the VTP version number is configured using the following commands:

```
Switch# vlan database  
Switch(vlan)# vtp v2-mode
```

Configuring the VTP Version on a CLI-Based Switch

On a CLI-based switch, the VTP version number is configured using the following command:

```
Switch(enable) set vtp v2 enable
```

NOTE

Many separate VTP options can be given in a single **set vtp** command, if desired. Consider the following example:

```
Switch(enable) set vtp domain mayberry mode server password opie9 v2 enable
```

Here, one command has set the VTP management domain name to **mayberry**, enabled VTP server mode, set the VTP password to **opie9**, and enabled VTP version 2.

VTP Status

The current VTP parameters for a management domain can be displayed. On an IOS-based switch, use the **show vtp status** command. On a CLI-based switch, use the **show vtp domain** command. Example 4-2 demonstrates some sample output of this command on a CLI-based switch.

Example 4-2 **show vtp domain** Reveals VTP Parameters for a Management Domain

```
Switch> show vtp domain
Domain Name                Domain Index VTP Version Local Mode Password
-----
mydomain                    1           2           server      -

Vlan-count Max-vlan-storage Config Revision Notifications
-----
15          1023              7           disabled

Last Updater   V2 Mode Pruning PruneEligible on Vlans
-----
192.168.1.4    enabled disabled 2-1000
Switch>
```

VTP message and error counters can also be displayed with the **show vtp counters** IOS-based command and the **show vtp statistics** CLI-based command. This command can be used for basic VTP troubleshooting to see if the switch is interacting with other VTP nodes in the domain. Example 4-3 demonstrates some sample output from the **show vtp statistics** command.

Example 4-3 show vtp statistics Reveals VTP Message and Error Counters

```

Switch> show vtp statistics
VTP statistics:
summary advts received      8
subset advts received      11
request advts received     0
summary advts transmitted  1
subset advts transmitted   1
request advts transmitted  0
No of config revision errors 0
No of config digest errors  0

VTP pruning statistics:

Trunk      Join Transmitted  Join Received  Summary advts received from
-----  -----  -----  non-pruning-capable device
3/1
Switch>

```


VTP Pruning

Recall that by definition, a switch must forward broadcast frames out all available ports in the broadcast domain because broadcasts are destined everywhere there is a listener. Multicast frames, unless forwarded by more intelligent means, follow the same pattern. (Multicast switching is covered in detail in Chapter 11, “Configuring Multicast Networks.”)

In addition, frames destined for an address that the switch has not yet learned or has forgotten (the MAC address has aged out of the address table) must be forwarded out all ports in an attempt to find the destination. These frames are referred to as *unknown unicast*.

When forwarding frames out all ports in a broadcast domain or VLAN, trunk ports are included. By default, a trunk link transports traffic from all VLANs, unless specific VLANs are removed from the trunk with the **clear trunk** command. Generally, in a network with several switches, trunk links are enabled between switches and VTP is used to manage the propagation of VLAN information. This scenario causes the trunk links between switches to carry traffic from *all* VLANs—not just from the specific VLANs created.


Consider the network shown in Figure 4-9. When end user HostPC in VLAN 3 sends a broadcast, Catalyst switch C forwards the frame out all VLAN 3 ports, including the trunk link to Catalyst A. Catalyst A, in turn, forwards the broadcast on to Catalysts B and D over those trunk links. Catalysts B and D forward the broadcast out only their access links that have been configured for VLAN 3. If Catalysts B and D don’t have any users in VLAN 3, forwarding that broadcast frame to them would consume bandwidth on the trunk links and processor resources in both switches, only to have switches B and D discard the frames.

Figure 4-9 *Flooding in a Catalyst Switch Network*

VTP pruning makes more efficient use of trunk bandwidth by reducing unnecessary flooded traffic. Broadcast and unknown unicast frames on a VLAN are forwarded over a trunk link only if the switch on the receiving end of the trunk has ports in that VLAN. VTP pruning occurs as an extension to VTP version 1, using an additional VTP message type. When a Catalyst switch has a port associated with a VLAN, the switch sends an advertisement to its neighbor switches that it has active ports on that VLAN. The neighbors keep this information, enabling them to decide if flooded traffic from a VLAN should use a trunk port or not.

Figure 4-10 shows the network from Figure 4-9 with VTP pruning enabled. Because Catalyst B has not advertised its use of VLAN 3, Catalyst A will choose not to flood VLAN 3 traffic to it over the trunk link. Catalyst D has advertised the need for VLAN 3, so traffic will be flooded to it.

Figure 4-10 Flooding in a Catalyst Switch Network Using VTP Pruning



Enabling VTP Pruning on an IOS-Based Switch

By default, VTP pruning is disabled on IOS-based switches. In the VLAN database configuration mode, the **vtp pruning** command can be used to enable pruning.

Enabling VTP Pruning on a CLI-Based Switch

VTP pruning is enabled using the **set vtp pruning enable** command. If this command is used on a VTP server, pruning is enabled for the entire management domain. By default, VTP pruning is disabled. When pruning is enabled with this command, all VLANs become eligible for pruning on all trunk links, if needed. The default list of pruning eligibility can be modified. Like VLAN trunking, you can first clear VLANs from the eligibility list using the **clear vtp pruneeligible** *vlan-range* command. Then, specify the VLANs that can be pruned with the **set vtp pruneeligible** *vlan-range* command.

NOTE

By default, VLANs 2–1000 are eligible for pruning. VLAN 1 has a special meaning because it is normally used as a management VLAN and is never eligible for pruning. In addition, VLANs 1001–1005 are never eligible for pruning.

Example 4-4 shows the use of these commands. VTP pruning is enabled, and VLAN 6 should be eligible for pruning, while VLANs 5, 7, 8, 9, and 10 are not eligible for pruning. Notice how VLANs 1 and 1001–1005 are never eligible for pruning.

Example 4-4 *Enabling VTP Pruning*

```
Switch(enable) set vtp pruning enable
Switch(enable) clear vtp pruneeligible 5-10
Switch(enable) set vtp pruneeligible 6
```

The pruning status on the switch and its VLANs can be displayed with the **show vtp domain** command. Example 4-5 shows some sample output from this command:

Example 4-5 *show vtp domain Command Output Displays VTP Pruning Status on a Switch and Its VLANs*

```
Switch> show vtp domain
Domain Name                Domain Index VTP Version Local Mode Password
-----
accounting                  1            2            server      -

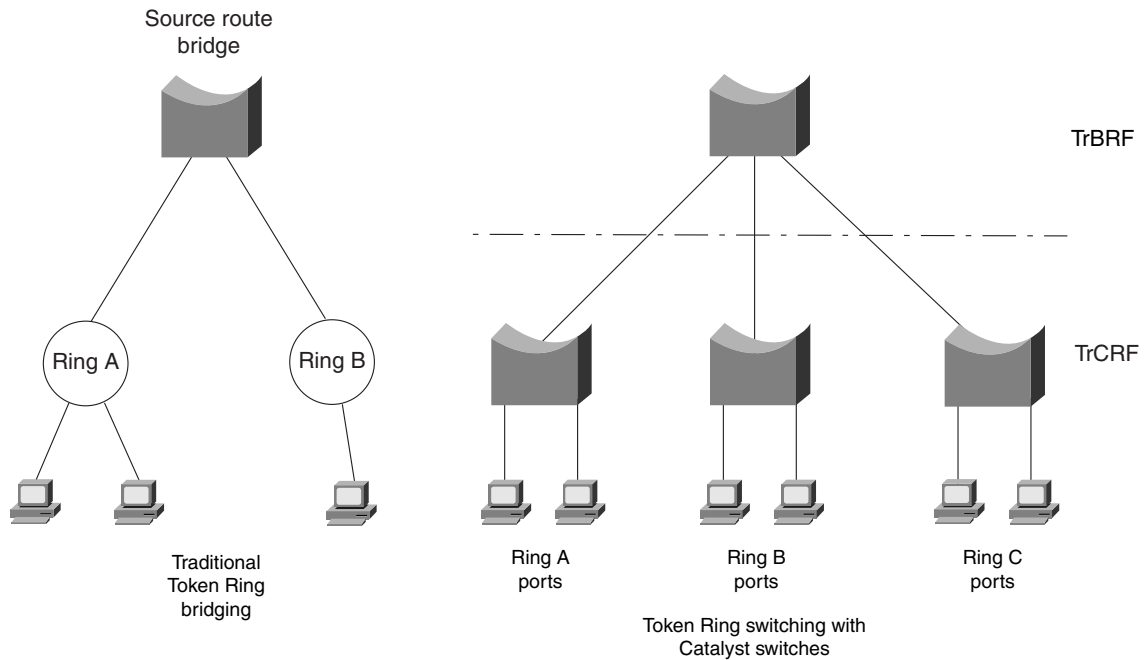
Vlan-count Max-vlan-storage Config Revision Notifications
-----
3           1023           2            disabled

Last Updater   V2 Mode   Pruning   PruneEligible on Vlans
-----
172.16.4.1     disabled enabled   6,11-2000
Switch>
```

Token Ring VLANs

This section discusses VLANs as they are applied to Token Ring networks. Only the Catalyst 5000 and the Catalyst 3900 switches support Token Ring—both using CLI-based commands.

Recall from the discussion in Chapter 3, “Basic Switch and Port Configuration,” the basic topology of Token Ring networks. End stations are connected to multistation access units (MSAUs), which interconnect with other MSAUs to form a ring. Multiple rings can be interconnected by bridges for segmentation and frame forwarding using source-route bridging and the RIF information. Figure 4-11 shows a typical Token Ring network with two rings and a source-route bridge in the left half. The right half of the figure shows a similar network topology with three rings and a source-route bridge, using the Token Ring switching features of Catalyst switches. The functionality of rings and bridges are performed within the switches, using Token Ring switching functions.

Figure 4-11 *Token Ring Networks: Traditional and Switched*

Token Ring switching follows the same topology, but performs the various functions within the switch. Where groups of end stations are connected by MAUs in a ring, the IEEE has defined the Concentrator Relay Function (CRF). The function of a multiport bridge to connect individual rings is defined as the Bridge Relay Function (BRF). These functions as performed by Catalyst switches are further described in the sections that follow.

TrBRF

A Catalyst switch connects *logical Token Ring Concentrator Relay Functions (TrCRFs)* with a logical multiport bridge, or *Token Ring Bridge Relay Function (TrBRF)*. In the hierarchy of bridged Token Rings, each TrCRF must be connected to a *parent TrBRF*. The hierarchical structure of Token Ring VLANs is shown in the right half of Figure 4-11.

By default, the TrBRF interconnects only TrCRFs located on the local switch. However, if trunking is used with ISL encapsulation, the TrBRF can extend to TrCRFs located on other Catalyst switches.

Each TrBRF exists as a special VLAN within a Catalyst switch. A switch can support many TrBRFs, but only one VLAN can be assigned to each TrBRF. By default, one TrBRF is defined as “trbrf-default” on VLAN 1005. Each TrBRF can operate as either a source-route bridge

(SRB), a source-route transparent (SRT) bridge, or both as a mixed mode. Additionally, each TrBRF runs a separate instance of either the IBM or IEEE Spanning-Tree Protocol to prevent bridging loops. The Spanning-Tree Protocol is covered in Chapter 5.

NOTE

To create and use Token Ring VLANs, VTP version 2 must be enabled on all Catalyst switches in the Token Ring domain. Enabling VTP version 2 was discussed in the previous “VTP Configuration” section.

To define a TrBRF on a Catalyst switch, use the following command:

```
Switch(enable) set vlan vlan-num [name name] type trbrf bridge bridge-num [stp  
{ieee | ibm}]
```

The only two required fields for a TrBRF are the VLAN number and the bridge number. Bridge numbers are defined by a single byte value, as a hexadecimal number from 0x1 to 0xf. The default Spanning-Tree Protocol is **ibm**. Notice that the type of bridging is not defined with the TrBRF, although the TrBRF performs the actual bridging function. Instead, the type of bridging is defined at the TrCRF. This way, multiple TrCRFs can be connected by a single parent TrBRF, each bridged with the desired method.

TrCRF

In a Catalyst switch, individual Token Ring ports can be connected to a logical ring, or Token Ring VLAN, by assigning them with identical ring numbers. Internally, the Catalyst performs the TrCRF to maintain the ring connectivity. Frame forwarding between ports on a common ring is performed with source-route switching, using either MAC addresses or route descriptors.

The TrCRF can be confined within a single switch or can be spread across multiple switches, depending on the topology and switch configuration. When a TrCRF is contained completely within a switch, it is referred to as an *undistributed TrCRF*. However, a TrCRF can be distributed across multiple switches if ISL trunking is enabled between switches and TrCRFs with identical VLAN numbers are defined.

By default, one TrCRF is defined on every Catalyst switch as “trcrf-default” on VLAN 1003. The default TrCRF is also assigned to the default TrBRF on VLAN 1005. If ISL trunking is in use, every Token Ring port on every switch will be defined to the same distributed TrCRF. Because only one TrBRF is defined by default, no bridging will occur. Instead, source-route switching will be performed to forward frames between switch ports within the TrCRF.

To define a TrCRF on a Catalyst switch, use the following command:

```
Switch(enable) set vlan vlan-num [name name] type trcrf  
{ring hex-ring-num | decring decimal-ring-num} parent vlan-num
```

Both the ring number and the parent VLAN number must be specified to define a TrCRF. The ring number can be defined as a hexadecimal value of 0x1 to 0xffff with the **ring** option, or as a decimal value of 1 to 4095 with the **decring** option. The parent VLAN number must match the VLAN number assigned to the parent TrBRF.

On the Catalyst 5000, a single TrCRF can be distributed across multiple switches. To enable this feature, use the **set tokenring distrib-crf enable** command.

After a TrCRF VLAN has been created, switch ports can be assigned to it. As with Ethernet switching, use the following command to assign ports to a VLAN:

```
Switch(enable) set vlan vlan-num mod-num/port-num
```

To view the current Token Ring VLAN configuration, use the **show vlan** command. The output of which is demonstrated in Example 4-6.

Example 4-6 show vlan Command Output Displays the Current Token Ring VLAN Configuration

```
Switch(enable) show vlan
```

VLAN Name	Status	Mod/Ports, Vlans
1 default	active	1/1-2
800 brf800	active	801,802
801Floor_1	active	
802Floor_2	active	
1002 fddi-default	active	
1003 trcrf-default	active	2/1-16
1004 fddinet-default	active	
1005 trbrf-default	active	1003

VLAN	Type	SAID	MTU	Parent	RingNo	BrdgNo	Stp	BrdgMode	Trans1	Trans2
1	enet	100001	1500	-	-	-	-	-	0	0
800	trbrf	100800	4472	-	-	0x2	ibm	-	0	0
801	trcrf	100201	4472	800	0x01	-	srb	-	0	0
802	trcrf	100202	4472	800	0x02	-	srb	-	0	0
1002	fddi	101002	1500	-	0x0	-	-	-	0	0
1003	trcrf	101003	4472	1005	0xcc	-	srb	-	0	0
1004	fdnet	101004	1500	-	-	0x0	ieee	-	0	0
1005	trbrf	101005	4472	-	-	0xf	ibm	-	0	0

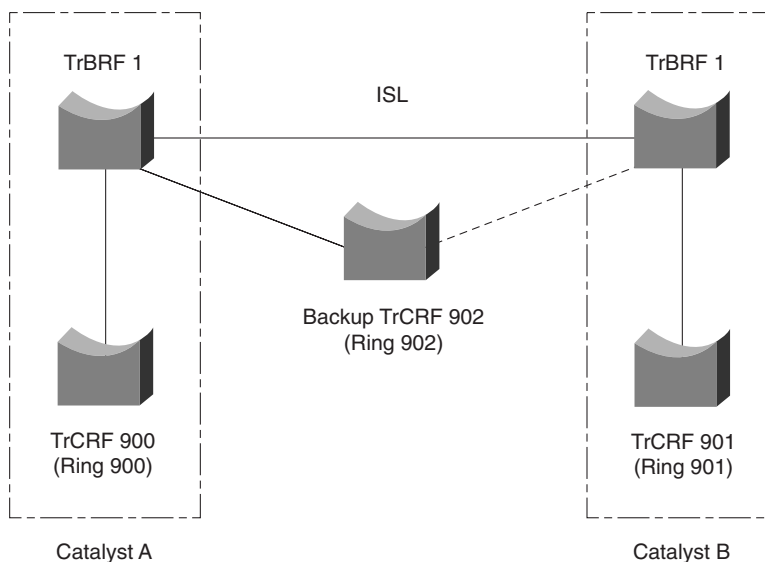
Notice in this example that one TrBRF has been assigned to VLAN 800, and two TrCRFs to VLANs 801 and 802. TrBRF on VLAN 800 shows the two active TrCRFs. In the lower group of output lines, the two TrCRFs on VLANs 801 and 802 show the parent TrBRF as 800 with source-route bridging enabled. Even though the defaults are not being used in this example, the default TrBRF is listed as VLAN 1005, and the default TrCRF is VLAN 1003, assigned to the default TrBRF.

TrCRF Redundancy

Catalyst switches also offer a form of redundancy for Token Ring switching. When two switches are connected by a common TrBRF *and* ISL trunking is enabled, connectivity between the TrCRFs in the switches could be disrupted if the ISL trunk link fails. A *backup TrCRF* can be used to provide a backup path in this case.

For each TrBRF, a single backup TrCRF can be defined with a single port from each connected switch. Only one of the TrCRF ports will be active at all times, while the other ports will be disabled. If the ISL trunk link goes down (along with the common TrBRF), the backup TrCRF links will come up and pass traffic between switches. A backup TrCRF is shown in Figure 4-12. The backup TrCRF is first defined on all switches, assigned to the TrBRF connecting the switches, marked as a backup TrCRF, and then assigned to one port on each switch.

Figure 4-12 A Backup TrCRF



To enable a backup TrCRF, first define a TrCRF that spans between switches. Then assign one port from each switch to the backup TrCRF. Finally, use the **set vlan *vlan-num* backupcrf on** command to enable the backup TrCRF function.

VTP and Token Ring VLANs

Using VTP in a Token Ring network domain will simplify VLAN administration, just as it does for Ethernet. TrCRF information will be propagated to all switches in a management domain.

As well, VTP pruning can also be performed on Token Ring VLANs. Both the default TrBRF (VLAN 1005) and the default TrCRF (VLAN 1003) are always pruning ineligible. VTP pruning is configured on a per-TrBRF basis. When a TrBRF is made pruning-eligible, all TrCRFs connected to it are also made pruning-eligible.

Duplicate Ring Protocol (DRiP)

Catalyst switches also have a mechanism to monitor the use of TrCRFs or ring numbers within a domain of switches. The *Duplicate Ring Protocol (DRiP)* collects and maintains the status of TrCRFs that are interconnected by TrBRFs. This information is used for the following purposes:

- Preventing duplicate ring numbers from being assigned to TrCRFs.
- Filtering All-Routes Explorer (ARE) frames from reentering TrCRFs that they have already visited.
- Operating the backup TrCRF function when an ISL trunk link goes down.

Every switch participating in Token Ring switching sends a DRiP advertisement out all ISL trunk ports every 30 seconds. Advertisements are sent to multicast address 01:00:0C:CC:CC:CC and are sent only on the default VLAN 1. When a switch receives the multicast advertisements, the switch does not forward the advertisements on to other switches over ISL links unless the advertisements contain new information. As well, a switch compares advertisements to the information in its own configuration. If it detects that a TrCRF has already been configured elsewhere, the local TrCRF configuration will be denied.

Foundation Summary

The Foundation Summary is a collection of tables and figures that provides a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final preparation before the exam, these tables and figures will hopefully be a convenient way to review the day before the exam.

Figure 4-13 *VTP Summary Advertisement Format*

Version (1 byte)	Type (Summary Adv) (1 byte)	Number of subset advertisements to follow (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision Number (4 bytes)			
Updater Identity (originating IP address; 4 bytes)			
Update Timestamp (12 bytes)			
MD5 Digest hash code (16 bytes)			

Figure 4-14 VTP Subset Advertisement and VLAN Info Field Formats**VTP Subset Advertisement**

0	1	2	3
Version (1 byte)	Type (Subset Adv) (1 byte)	Subset sequence number (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision Number (4 bytes)			
VLAN Info Field 1 (see below)			
VLAN Info Field ...			
VLAN Info Field N			

VTP VLAN Info Field

0	1	2	3
Info Length	VLAN Status	VLAN Type	VLAN Name Length
ISL VLAN ID		MTU Size	
802.10 SAID			
VLAN Name (padded with zeros to multiple of 4 bytes)			

Figure 4-15 VTP Subset Request Format

0	1	2	3
Version (1 byte)	Type (Adv request) (1 byte)	Reserved (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Starting advertisement to request			

Table 4-2 *VLAN Configuration Commands*

Task	IOS-Based Command	CLI-Based Command
Create VLAN	vlan database vlan <i>vlan-num</i> name <i>vlan-name</i>	set vlan <i>vlan-num</i> [name <i>name</i>]
Assign port to VLAN	interface <i>interface module/number</i> switchport mode access switchport access vlan <i>vlan-num</i>	set vlan <i>vlan-num mod-num/port-list</i>
Display VLANs	show vlan	show vlan
Configure trunk	interface <i>interface mod/port</i> switchport mode trunk switchport trunk encapsulation { isl dot1q } switchport trunk allowed vlan remove <i>vlan-list</i> switchport trunk allowed vlan add <i>vlan-list</i>	set trunk <i>module/port</i> [on off desirable auto nonegotiate] <i>vlan-range</i> [isl dot1q dot10 lane negotiate] clear trunk <i>module/port vlan-range</i>
Display trunks	show interface <i>mod/num</i> switchport	show trunk

Table 4-3 *VTP Configuration Commands*

Task	IOS-Based Command	CLI-Based Command
Configure VTP domain	vlan database vtp domain <i>domain-name</i>	set vtp [domain <i>domain-name</i>]
Configure VTP mode	vlan database vtp domain <i>domain-name</i> vtp { server client transparent } vtp password <i>password</i>	set vtp [domain <i>domain-name</i>] [mode { server client transparent }] [passwd <i>password</i>]
Configure VTP version	vlan database vtp v2-mode	set vtp v2 enable
Display VTP status	show vtp status show vtp counters	show vtp domain show vtp statistics
VTP pruning	vtp pruning	set vtp pruning enable set vtp pruneeligible <i>vlan-range</i> clear vtp pruneeligible <i>vlan-range</i>

Table 4-4 *Token Ring VLAN Configuration Commands*

Task	IOS-Based Command	CLI-Based Command
Define TrBRF	N/A	set vlan <i>vlan-num</i> [name <i>name</i>] type trbrf bridge <i>bridge-num</i> [stp { <i>ieee</i> <i>ibm</i> }]
Define TrCRF	N/A	set vlan <i>vlan-num</i> [name <i>name</i>] type trcrf { ring <i>hex-ring-num</i> decring <i>decimal-ring-num</i> } parent <i>vlan-num</i>
Enable distributed TrCRF	N/A	set tokenring distrib-crf enable
Assign Token Ring ports to TrCRF	N/A	set vlan <i>vlan-num</i> <i>mod-num/port-num</i>
Enable backup TrCRF	N/A	set vlan <i>vlan-num</i> backupcrf on

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however, they are designed to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to these questions can be found in Appendix A, on page 477.

- 1 What is a VLAN? When is it used?

- 2 When a VLAN is configured on a Catalyst switch port, in how much of the campus network will the VLAN number be unique and significant?

- 3 What are two types of VLANs, in terms of spanning areas of the campus network?

- 4 What is the Catalyst CLI-based switch command to configure ports 4/11 and 5/1 through 5/24 for VLAN 2?

5 Generally speaking, what must be configured (both switch and end user device) for a port-based VLAN?

6 What is the default VLAN on all ports of a Catalyst switch?

7 What are the components of a Token Ring VLAN?

8 What is a trunk link?

9 What methods of Ethernet VLAN frame identification can be used on a Catalyst switch?

10 What is the difference between these two trunking methods? How many bytes are added to trunked frames for VLAN identification in each method?

11 What is the purpose of Dynamic Trunking Protocol (DTP)?

12 What CLI-based commands are needed to configure a Catalyst switch trunk port 1/1 to transport only VLANs 100, 200–205, and 300 using IEEE 802.1Q? (Assume that trunking is enabled and active on the port already.)

13 What VTP modes can a Catalyst switch be configured for? Can VLANs be created in each of the modes?

14 Two neighboring switch trunk ports are set to *auto* mode with *ISL* trunking mode. What will the resulting trunk mode become?

15 How many VTP management domains can a Catalyst switch participate in? How many VTP servers can a management domain have?

16 What CLI-based command can be used on a Catalyst switch to verify exactly what VLANs will be transported over a trunk link?

17 What conditions must exist for two Catalyst switches to be in the same VTP management domains?

-
- 18** What are the types of VTP messages or advertisements used by Catalyst switches? What field in these messages determines if a switch should use and record VLAN data in the messages?

- 19** What CLI-based command can be used to configure a Catalyst switch to become a VTP server for the domain “engineering”? The domain should be secured with the password “secret123.”

- 20** What is the purpose of VTP pruning?

- 21** Which VLAN numbers are never eligible for VTP pruning? Why?

- 22** What commands can be used to make only VLANs 300 and 400 eligible for VTP pruning?

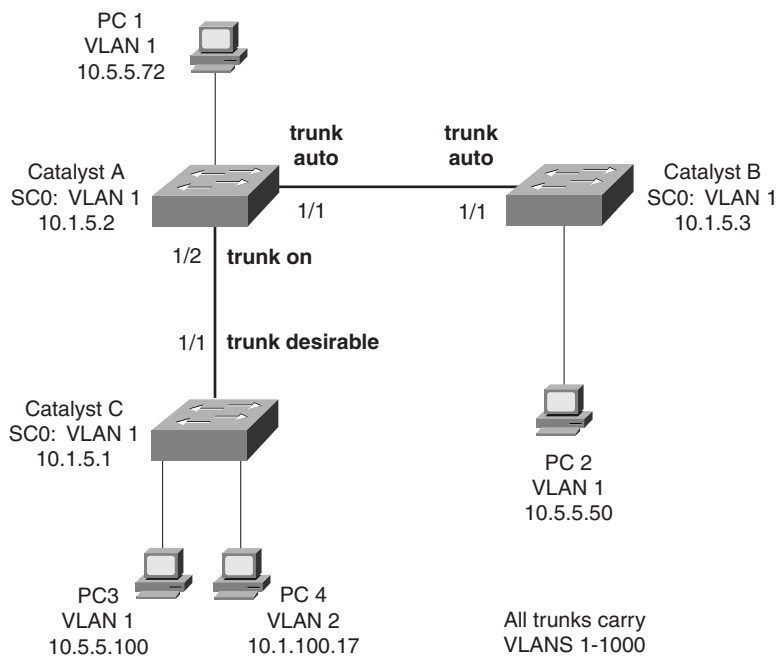
- 23** What are the steps needed to establish Token Ring switching with VLANs?

Scenarios

Scenario 4-1

Consider the network shown in Figure 4-16 and answer the questions that follow.

Figure 4-16 Diagram for Scenario 4-1

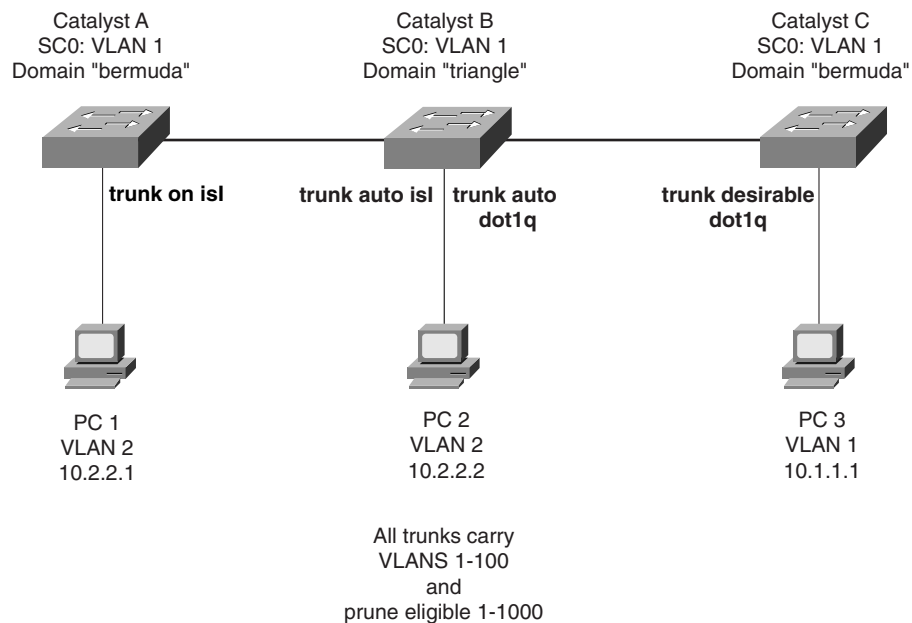


- 1 What is the mode of the link between Catalyst A and Catalyst B?
- 2 Now suppose the network administrator types the command **set trunk 1/1 nonegotiate** on Catalyst B. What will the link mode be now?
- 3 Catalyst B has been given the command **set trunk 1/1 on**. What is the link mode now?
- 4 What is the mode of the link between Catalyst A and Catalyst C?
- 5 Assume that all links between Catalyst switches are in trunking mode, transporting VLANs 1-1000. Can PC-2 ping PC-4?
- 6 Suppose PC-1 begins to generate a broadcast storm. Where would the effects of this storm be experienced in this network? Consider both devices and links. Will PC-4 receive the broadcasts?

Scenario 4-2

See the diagram shown in Figure 4-17 and answer the questions that follow.

Figure 4-17 Diagram for Scenario 4-2



- 1 What is the mode of the link between Catalyst A and Catalyst B?
- 2 Can Catalyst A **ping** Catalyst C? Can PC-1 **ping** PC-2? Why or why not?
- 3 Suppose Catalyst B's VTP domain is now changed to "bermuda." Can Catalyst A **ping** Catalyst C?
- 4 Which Catalyst switches will receive a broadcast from PC-1?
- 5 Where will VLAN1 be pruned? Why?
- 6 Now suppose Catalyst A is a VTP server, Catalyst C is a VTP client, and Catalyst B is configured for VTP transparent mode. All switches are in the "bermuda" management domain. If VLAN14 is created on Catalyst A, which switches will also get VLAN14 created via VTP?
- 7 If VLAN15 is created on Catalyst B, what other switches will also create VLAN15 via VTP?
- 8 If VLAN16 is created on Catalyst C, what will happen?

Scenarios Answers

Scenario Answers 4-1

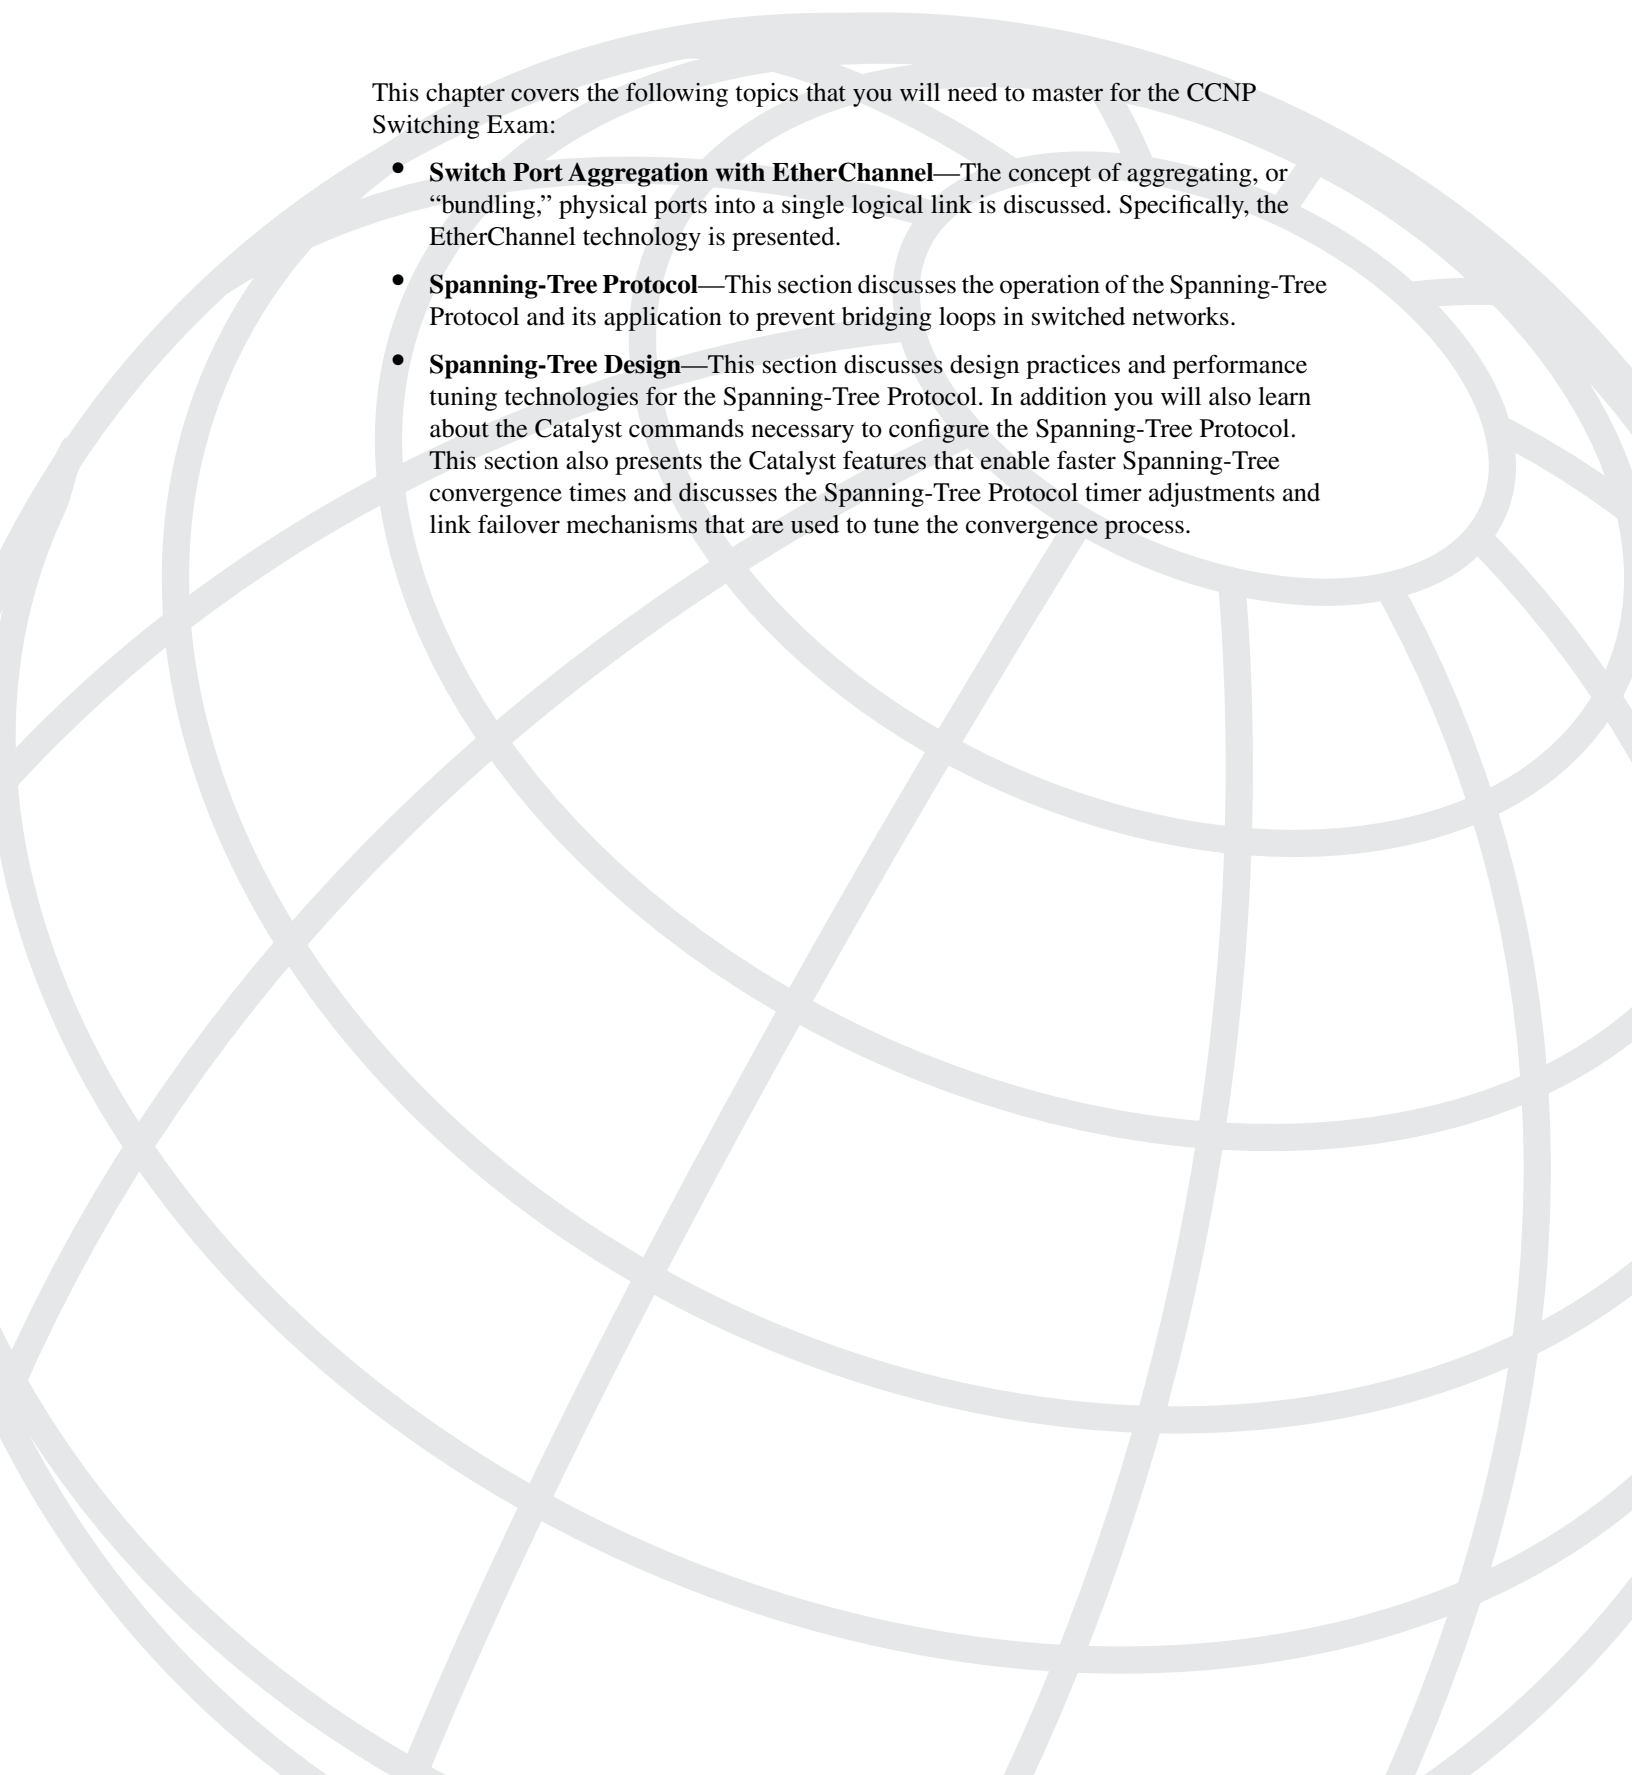

- 1 The link is still an access link, with no trunking established because both switches are set to *auto* mode. The switches are each passively waiting for the other to initiate trunking.
- 2 Trunking is still not established. Catalyst A is waiting to be asked to trunk, and Catalyst B is set to *nonegotiate*. Catalyst B will never try to negotiate trunking.
- 3 Trunking has finally been established.
- 4 Trunking. Catalyst A expects trunking on the link, while Catalyst C actively tries to negotiate trunking.
- 5 No. The two PC devices are connected to different VLANs. Without a router or Layer 3 device connecting the VLANs, no traffic will cross between them.
- 6 All hosts on VLAN1 (PC-1, PC-2, and PC-3) will experience the broadcast storm. All trunk links between switches will transport the broadcast frames. In addition, all switch supervisor CPUs will receive and process the broadcasts, because each switch has its SC0 port assigned to VLAN1. (For this reason, it is recommended to reserve VLAN1 for management traffic only. User-generated broadcasts can overload the switch supervisor to the extent that it can no longer keep track of its management protocols like VTP, CDP, and so forth. Instead, all user traffic should be kept off VLAN1.)

Scenario Answers 4-2

- 1 The link is still an access link, with no trunking established. The two switches would have negotiated trunking, but the switches are configured for different VTP management domains. Neighboring switches must be in the same domain for trunking to be negotiated.
- 2 Catalyst A can **ping** Catalyst B. The SC0 ports on both switches are configured for the same VLAN. Because trunking has not been established between Catalyst A and Catalyst B (due to domain name conflicts), the link is still an access link. Fortunately, the access link has defaulted to VLAN1 so that the two SC0 ports on VLAN1 can communicate.

PC-1 cannot **ping** PC-2, however. Both PCs are in VLAN2, but VLAN2 is not being transported between switches because the trunk link has not been established.
- 3 Yes. Trunk links are now negotiated or established between all switches.
- 4 Catalyst A and Catalyst B. Because Catalyst C has no ports in VLAN2 (where PC-1 resides), VLAN2 will be pruned by Catalyst B and will not cross the trunk link to Catalyst C.

- 5 VLAN1 will not be pruned at all. Although VLAN1 is present on all switches, it is not pruned because VLAN1 is ineligible for pruning by definition. Remember that VLAN1 is usually used for management traffic and should be kept intact so that no switches become isolated.
- 6 Only Catalyst C will create VLAN14 in response to VTP advertisements. Catalyst B in transparent mode will only relay the VTP information without interpreting the information.
- 7 Only Catalyst B will create VLAN15. Because it is in transparent mode, no VLAN activity will be advertised to other neighboring switches. However, Catalyst B is allowed to create, delete, and rename VLANs freely. These VLANs are significant only to the local switch.
- 8 Catalyst C will not allow any VLANs to be created, unless they are learned from a VTP server in the "bermuda" domain. Because it is in VTP client mode, no VLAN changes can be performed from the console.



This chapter covers the following topics that you will need to master for the CCNP Switching Exam:

- **Switch Port Aggregation with EtherChannel**—The concept of aggregating, or “bundling,” physical ports into a single logical link is discussed. Specifically, the EtherChannel technology is presented.
- **Spanning-Tree Protocol**—This section discusses the operation of the Spanning-Tree Protocol and its application to prevent bridging loops in switched networks.
- **Spanning-Tree Design**—This section discusses design practices and performance tuning technologies for the Spanning-Tree Protocol. In addition you will also learn about the Catalyst commands necessary to configure the Spanning-Tree Protocol. This section also presents the Catalyst features that enable faster Spanning-Tree convergence times and discusses the Spanning-Tree Protocol timer adjustments and link failover mechanisms that are used to tune the convergence process.

Redundant Switch Links

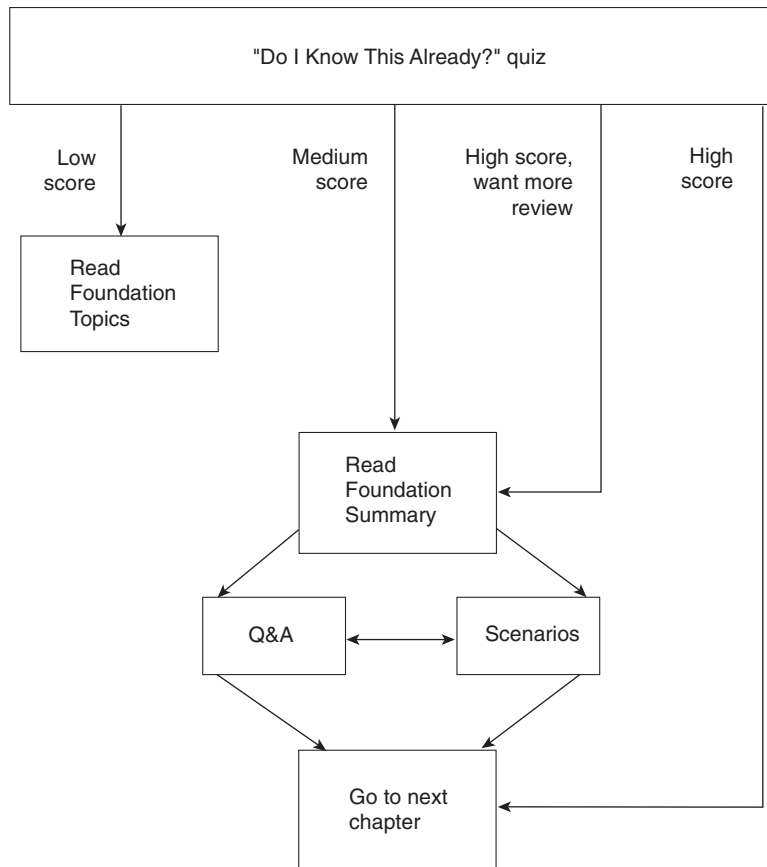
In the previous chapters, you learned about campus network design and connecting and organizing switches into blocks and common workgroups. Using these principles, end users can be given effective access to resources both on and off of the campus network. However, today's mission critical applications and services demand networks that provide high availability and reliability.

This chapter presents technologies that can be used in a campus network to provide higher reliability. Redundancy between switches, fault tolerance and recovery, and timely access are all techniques that are discussed. Each of these makes use of redundant links between switches and switch blocks.

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and the answers for all your work with this book in one place for easy reference.
- Take the “Do I Know This Already?” quiz and write down facts and concepts, even if you never look at the information again.
- Use the diagram in Figure 5-1 to guide you to the next step.

Figure 5-1 *How to Use This Chapter*

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The quiz helps you make good choices of how to spend your limited study time. The quiz is sectioned into three smaller “quizlets,” which correspond to the five major headings in the “Foundation Topics” section of the chapter. Although your answer may differ somewhat from the answers given, it is more important that you find out if you have the basic understanding that is presented in this chapter. You will find that these questions are open-ended rather than multiple choice as found on the exams. This way you will focus more on understanding the subject matter than on memorizing details.

Use the scoresheet in Table 5-1 to record your score.

Table 5-1 *Scoresheet for Quiz and Quizlets*

Quizlet Number	Foundation Topics Sections Covering These Questions	Questions	Score
1	Switch Port Aggregation with EtherChannel	1–3	
2	Spanning-Tree Protocol	4–7	
3	Spanning-Tree Design	8–16	
All questions		1–16	

1 What is EtherChannel? What types of switch links can it be used with?

2 How is traffic distributed over an EtherChannel?

3 What is PAgP used for?

4 What is a bridging loop? Why is it bad?

5 Name two types of Spanning-Tree Protocol messages used to communicate between bridges.

6 What criteria are used to select the following:

- a. Root Bridge
- b. Root Port
- c. Designated Port
- d. Redundant (or Secondary) Root Bridges

7 What conditions cause an STP topology change? What effect does this have on STP and the network?

8 What is the single most important design decision to be made in a network running STP?

9 Where should the Root Bridge be located in a switched network?

10 What happens to a port that is neither a Root Port nor a Designated Port?

11 How is the Root Path Cost calculated for a switch port?

12 What is the maximum number of Root Ports that a Catalyst switch can have?

13 What mechanism is used to set STP timer values for all switches in a network?

14 What parameters can be tuned to influence the selection of a port as a Root or Designated Port?

15 What technology can be useful to decrease the amount of time STP keeps an end user’s workstation in the Blocking state when it powers up?

16 Where should the UplinkFast feature be used in a switched network?

The answers to the quiz are found in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections,” on page 477. The suggested choices for your next step are as follows:

- **8 or less overall score**—Read the entire chapter. This includes the “Foundation Topics” and “Foundation Summary” sections, the Q&A section, and the scenarios at the end of the chapter.
- **9–11 overall score**—Begin with the “Foundation Summary” section and then follow with the Q&A section and the scenarios at the end of the chapter.
- **12 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” section and then go to the Q&A section and the scenarios at the end of the chapter. Otherwise, move to the next chapter.

Foundation Topics

Switch Port Aggregation with EtherChannel

As discussed in Chapter 3, “Basic Switch and Port Configuration,” switches can use Ethernet, Fast Ethernet, or Gigabit Ethernet ports to scale link speeds by a factor of ten. Cisco offers another method of scaling link bandwidth by aggregating or bundling parallel links, termed the *EtherChannel* technology. Two to eight links of either Fast Ethernet (FE) or Gigabit Ethernet (GE) are bundled as one logical link of *Fast EtherChannel (FEC)* or *Gigabit EtherChannel (GEC)*, respectively. This bundle provides a full-duplex bandwidth of up to 1600 Mbps (8 links of Fast Ethernet) or 16 Gbps (8 links of Gigabit Ethernet).

As the Spanning-Tree Protocol (STP) portion of this chapter discusses, switches should never be configured to have multiple links connected to form a loop. EtherChannel avoids this situation by bundling parallel links into a single logical link, which can act as either an access or a trunk link. Switches or devices on each end of the EtherChannel link must understand and use the EtherChannel technology for proper operation.

Although an EtherChannel link is seen as a single logical link, the link does not have an inherent total bandwidth equal to the sum of its component physical links. For example, suppose a FEC link is made up of four full-duplex 100-Mbps Fast Ethernet links. Although it is possible for the FEC link to carry a throughput of 800 Mbps, the single resulting FEC link does not operate at this speed. Instead, traffic is balanced across the individual links within the EtherChannel. Each of these links operates at its inherent speed (200 Mbps full-duplex for FE) but carries only the frames placed on it by the EtherChannel hardware. The load-balancing process is explained further in the next section.

EtherChannel also provides redundancy through the use of the several bundled physical links. If one of the links in the bundle fails, traffic sent through that link will move to an adjacent link. Failover occurs in less than a few milliseconds and is transparent to the end user. As more links fail, more traffic will be moved to further adjacent links. Likewise, as links are restored, the load will be redistributed among the links

Bundling Ports with EtherChannel

Fast EtherChannel is available on the Catalyst 1900, 2820, 2900, 2900XL, 3500XL, 4000, 5000, and 6000 families. Gigabit EtherChannel is supported only on the Catalyst 2900, 2900XL, 4000, 5000, and 6000 families. Most of the switch families support a maximum of four FE or GE links bundled in a single EtherChannel link. However, the Catalyst 6000 family supports up to eight bundled links for a total throughput of 1600 Mbps (FEC) or 16 Gbps (GEC). The Catalyst 6000 also supports up to 128 individual EtherChannel links.

On the majority of Catalyst switch modules, EtherChannel bundles have several configuration restrictions. For example, either two or four ports must be bundled and these ports must be

contiguous on the switch module. Newer switch modules allow the ports to be selected from anywhere on the module or even across modules. Generally, all bundled ports must first belong to the same VLAN. If used as a trunk, bundled ports must all be in trunking mode and pass the same VLANs. As well, each of the ports should have the same speed and duplex settings before they are bundled.

Distributing Traffic in EtherChannel

Traffic in an EtherChannel is statistically load-balanced across the individual links bundled together. However, the load is not necessarily balanced equally across all of the links. Instead, frames are forwarded on a specific link as a function of the addresses present in the frame. Some combination of source and destination addresses (either MAC or IP addresses) is used to form a binary pattern used to select a link number in the bundle.

Switches perform an exclusive-OR (XOR) operation on one or more low-order bits of the addresses to determine what link to use. For example, an EtherChannel consisting of two links bundled together requires the XOR of the last bit of the addresses in the frame. A four-link bundle uses the XOR of the last two bits. Likewise, an eight-link bundle uses the XOR of the last three bits. The outcome of the XOR operation selects the outbound link of the EtherChannel. Table 5-2 shows the results of an XOR on a two-link bundle.

Table 5-2 *Frame Distribution on a Two-Link EtherChannel*

Binary Addresses	Two-Link EtherChannel XOR and Link Number
Addr1: ... xxxxxx0	
Addr2: ... xxxxxx0	... xxxxxx0: Link 0
Addr1: ... xxxxxx0	
Addr2: ... xxxxxx1	... xxxxxx1: Link 1
Addr1: ... xxxxxx1	
Addr2: ... xxxxxx0	... xxxxxx1: Link 1
Addr1: ... xxxxxx1	
Addr2: ... xxxxxx1	... xxxxxx0: Link 0

The XOR operation is performed independently on each bit position in the address value. If the two address values have the same bit value, the XOR result is 0. If the two address bits differ, the XOR result is 1. In this way, frames can be statistically distributed among the links with the assumption that MAC or IP addresses are statistically distributed throughout the network. In a four-link EtherChannel, the XOR is performed on the lower two bits of the address values resulting in a two-bit XOR value (each bit is computed separately) or a link number from 0 to 3.

A conversation between two devices will always be sent through the same EtherChannel link because the two endpoint addresses stay the same. However, when a device talks to several

other devices, chances are that the destination addresses are equally distributed with zeros and ones in the last bit (even and odd address values). This causes the frames to be distributed across the EtherChannel links. Note that a conversation between two end devices to create a load imbalance is possible using one of the links in a bundle because all traffic between a pair of stations will use the same link.

Switches with an Ethernet Bundling Controller (EBC) are limited to distributing frames based on source and destination MAC addresses only. For each frame, the source MAC address is XOR'd with the destination MAC address. Because this is the only choice, no switch configuration is necessary.

Switches such as the IOS-based Catalyst 2900 and 3500XL distribute frames according to a different criteria. By default, EtherChannel frames are distributed by the low-order bits of their source MAC addresses. The administrator can select either source or destination addresses as the distribution method by using the following command (the port group is defined in the next section):

```
Switch (config-if)# port group group-number [distribution {source | destination}]
```

Other switches, such as the Catalyst 6000, offer more flexibility in computing frame distribution. The XOR operation can be performed on either MAC or IP addresses and can be based solely on source or destination addresses or both. Use the following command to configure frame distribution for all EtherChannel switch links:

```
Switch> (enable) set port channel all distribution {ip | mac} [source | destination | both]
```

The default configuration is to use IP addresses, both source and destination. Normally, this action should result in a statistical distribution of frames. However, you should determine if the EtherChannel is imbalanced according to the traffic patterns present. For example, if a single server is receiving most of the traffic on an EtherChannel, the source IP addresses of the stations talking to the server can cause one link to be overused. In the case of a four-link EtherChannel, perhaps two of the four links are overused. Configuring the use of MAC addresses or only the source IP addresses might cause the distribution to be more balanced across all the bundled links.

In applications involving switches like the Catalyst 6000, some EtherChannel traffic may consist of protocols other than IP. For example, IPX or SNA frames may be switched along with IP. Non-IP protocols would need to be distributed according to MAC addresses because IP addresses are not applicable. Here, the switch should be configured to use MAC addresses instead of the IP default.

NOTE A special case results when a router is connected to an EtherChannel because the router will use its own MAC address in all frames that it forwards to many end stations. For the EBC-based switch, this means that the destination MAC address is always the same for frames destined through the router. Usually this won't present a problem because the source MAC addresses are all different. When two routers are forwarding frames to each other, however, both source and destination MAC addresses will remain constant and only one link of the EtherChannel will be used. The flexibility in the Catalyst 6000 switch allows the administrator to select exactly which criteria frames will be distributed. If the MAC addresses are remaining constant, you should choose IP addresses instead.

Port Aggregation Protocol (PAgP)

To provide automatic EtherChannel configuration and negotiation between switches, Cisco developed the *Port Aggregation Protocol (PAgP)*. PAgP packets are exchanged between switches over EtherChannel-capable ports. The identification of neighbors and port group capabilities are learned and are compared with local switch capabilities. Ports that have the same neighbor device ID and port group capability will be bundled together as a bidirectional, point-to-point EtherChannel link.

PAgP will form an EtherChannel only on ports that are configured for either identical static VLANs or trunking. PAgP also dynamically modifies parameters of the EtherChannel if one of the bundled ports is modified. For example, if the VLAN, speed, or duplex mode of a port in an established bundle is changed, PAgP will change that parameter for all ports in the bundle.

When ports are bundled into an EtherChannel, all broadcasts and multicasts are sent over one port in the bundle only. Broadcasts will not be sent over the remaining ports and will not be allowed to return over any other port in the bundle.

Switch ports can be configured for the following modes of PAgP:

- **On**—The ports will always be bundled as an EtherChannel. No negotiation takes place because PAgP packets are not sent or processed.
- **Off**—The ports will never be bundled as an EtherChannel. They will remain as individual access or trunk links. No PAgP packets are sent.
- **Auto**—(*Default*) PAgP packets are sent to negotiate an EtherChannel only if the far end initiates EtherChannel negotiations. Therefore, *auto* mode is a passive mode that requires a neighbor in *desirable* mode. (Two switches in *auto* mode will never negotiate an EtherChannel because each is passively waiting for the other to request an EtherChannel.)
- **Desirable**—PAgP packets are sent to actively negotiate an EtherChannel. This mode starts the negotiation process, and will bring up a successful EtherChannel with another switch in either *desirable* or *auto* mode.

EtherChannel Configuration

Before configuring switch ports into an EtherChannel bundle, you should make sure the switch module supports it. Use the **show port capabilities** [*module/port*] command to do this. (This command is available on Catalyst software versions 4.x and later.) Example 5-1 demonstrates using the **show port capabilities** command to ensure the switch module supports EtherChannel bundling.

Example 5-1 show port capabilities Command Output

Switch (enable) show port capabilities 2	
Model	WS-X5234
Port	2/1
Type	10/100BaseTX
Speed	auto,10,100
Duplex	half,full
Trunk encap type	ISL,802.1Q
Trunk mode	on,off,desirable,auto,nonegotiate
Channel	2/1-2,2/1-4
Broadcast suppression	percentage(0-100)
Flow control	receive-(off,on),send-(off,on)
Security	yes
Membership	static,dynamic
Fast start	yes
Rewrite	yes

On this and other early Ethernet modules, only certain ports can be bundled. Notice that Example 5-1 shows that only ports 2/1 and 2/2 or 2/1 through 2/4 can be bundled. These modules use a hardware chip called the *Ethernet Bundling Controller (EBC)* to manage the EtherChannel ports. Ports to be bundled must belong to the same EBC, according to the specific arrangement of ports on the module. For example, a 24-port module offers three groups of eight ports and a 12-port module offers three groups of four ports. Generally, the EBC requires an EtherChannel to start with the first port of a group. The output of the **show port capabilities** command will show the acceptable port groupings, if they are available.

Newer modules, such as the Catalyst 6000, offer more flexibility with EtherChannel configuration. Ports located anywhere on an EtherChannel-capable module can be bundled along with ports from other modules.

NOTE

Remember the following guidelines that apply to the switch ports that will be grouped into an EtherChannel:

- All ports should be assigned to the same VLAN or configured for trunking (an EtherChannel can be used as a trunk link).
- If the EtherChannel will be a trunk link, all ports should have the same trunk mode and should carry the same VLANs over the trunk.

- All ports should be configured for the same speed and duplex mode.
- Do not configure the ports as dynamic VLAN ports.
- All ports should be enabled; a disabled port will be seen as a failed link, forcing its traffic to be moved to the next available link in the bundle.

EtherChannel Configuration on a CLI-Based Switch

To configure an EtherChannel on a CLI-based switch, use the following command:

```
Switch (enable) set port channel module/port-range mode {on | off | desirable | auto}
```

Ports are grouped into an EtherChannel by specifying them as a range, as in **set port channel 2/1-4 mode on**.

EtherChannel Configuration on an IOS-Based Switch

To configure an EtherChannel on an IOS-based switch, use the following command:

```
Switch (config-if)# port group group-number [distribution {source | destination}]
```

The port must be assigned to a group number, which represents the EtherChannel as a number from 1 to 12.

Displaying EtherChannel Configuration

Information about the current EtherChannel configuration can be displayed using the **show port channel** [*mod/port*] [**info** | **statistics**] command on a CLI-based switch and the **show port group** [*group-number*] command on an IOS-based switch. Example 5-2 demonstrates how the **show port channel info** command can be used to view the current status of EtherChannel links on a CLI-based switch.

Example 5-2 show port channel info Command Output

```
Switch> (enable) show port channel info
Switch Frame Distribution Method: mac both
```

Port	Status	Channel mode	Admin group	Channel id	Speed	Duplex	Vlan
3/29	connected	desirable silent	158	847	a-100	a-full	53
3/30	connected	desirable silent	158	847	a-100	a-full	53
3/31	connected	auto silent	159	848	a-100	a-full	101
3/32	connected	auto silent	159	848	a-100	a-full	101

The first shaded line shows that the switch is using both source and destination MAC addresses to distribute frames across the bundled links. The next set of shaded lines show that switch ports 3/29 and 3/30 are bundled as EtherChannel ID number 847, are operating at 100 Mbps full-duplex (autonegotiated), and are assigned to VLAN 53 only. The second EtherChannel is made up of switch ports 3/31 and 3/32 bundled as EtherChannel ID 848, passing VLAN 101.

Spanning-Tree Protocol

A robust network design not only includes efficient transfer of packets or frames but also considers how to recover quickly from faults in the network. In a Layer 3 environment, the routing protocol(s) in use keeps track of redundant paths to a destination network so that a secondary path can be quickly utilized if the primary path fails. Layer 3 routing allows many paths to a destination to remain up and active and allows load sharing across multiple paths.

In a Layer 2 environment (switching or bridging), however, no routing protocols are used and redundant paths are not allowed. Instead, some form of bridging provides data transport between networks or switch ports. The Spanning-Tree Protocol (STP) is used to provide network link redundancy and load balancing so that a Layer 2 switched network can recover from failures without intervention in a timely manner.

STP is discussed in relation to the problems it solves in the following sections.

Bridging Loops

Recall that a Layer 2 switch mimics the function of a transparent bridge. A transparent bridge must offer segmentation between two networks, while remaining transparent to all the end devices connected to it. For the purpose of this discussion, consider a two-port Ethernet switch and its similarities to a two-port transparent bridge.

A transparent bridge (and the Ethernet switch) must operate as follows:


- The bridge has no initial knowledge of the location of any end device; therefore, the bridge must “listen” to frames coming into each of its ports to figure out on which network a device resides. The source address in an incoming frame is the clue to a device’s whereabouts—the bridge assumes the source device is located behind the port that the frame arrived on. As the listening process continues, the bridge builds a table containing source MAC addresses and the bridge port numbers associated with them.

The bridge has the capability to constantly update its bridging table upon detecting the presence of a new MAC address or upon detecting a MAC address that has changed location from one bridge port to another. The bridge is then able to forward frames by looking at the destination address, looking up the address in the bridge table, and sending the frame out the port where the destination device is located.

- If a frame arrives with the broadcast address as the destination address, the bridge must forward or flood the frame out all available ports. However, the frame is not forwarded out the port that initially received the frame. In this way, broadcasts are able to reach all available networks. A bridge only segments collision domains but does not segment broadcast domains.
- If a frame arrives with a destination address that is not found in the bridge table, the bridge is unable to determine which port to forward the frame to for transmission. This type of frame is known as an *unknown unicast*. In this case, the bridge treats the frame as if it were a broadcast and forwards it out all remaining ports. After a reply to that frame is overheard, the bridge will learn the location of the unknown station and add it to the bridge table for future use.
- Frames that are forwarded across the bridge cannot be modified.

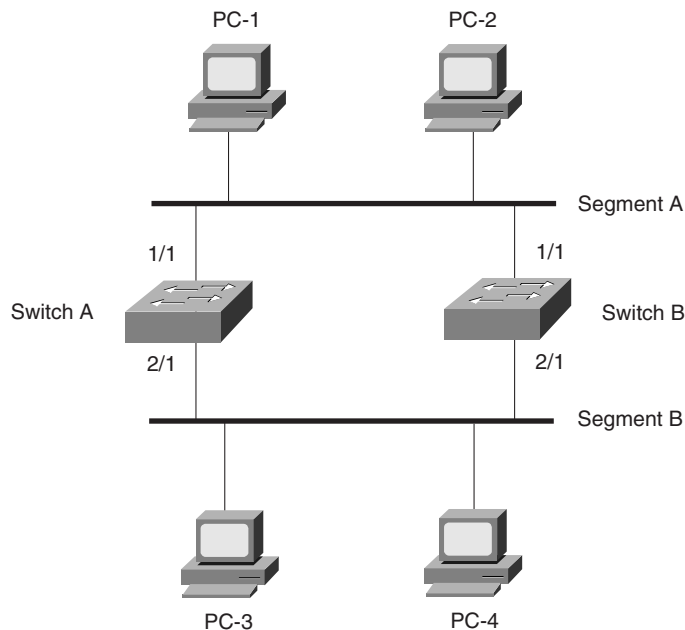
Bridging or switching in this fashion works well. Any frame received, whether to a known or unknown destination, will be forwarded out the appropriate port or ports so that it is very likely to be received successfully at the end device. Figure 5-2 shows a simple two-port switch functioning as a bridge, forwarding frames between two end devices. However, this network design offers no additional links or paths for redundancy, should the switch or one of its links fail.

Figure 5-2 *Transparent Bridging with a Switch*



To add some redundancy, a second switch can be added between the two original network segments, as shown in Figure 5-3. Now two switches offer the transparent bridging function in parallel. Consider what will happen when PC-1 sends a frame to PC-4. For now, assume that both PC-1 and PC-4 are known to the switches and are in their address tables. PC-1 sends the frame out onto network Segment A. Switch A and Switch B both receive the frame on their 1/1 ports. Because PC-4 is already known to the switches, the frame is forwarded out ports 2/1 on each switch onto Segment B. The end result is that PC-4 will receive two copies of the frame from PC-1. This is not ideal, but is not disastrous either.

Figure 5-3 *Redundant Bridging with Two Switches*



Now consider the same process of sending a frame from PC-1 to PC-4. This time, however, neither switch knows anything about PC-1 or PC-4. PC-1 sends the frame to PC-4 by placing it on Segment A. The sequence of events is as follows:

- 1 Both Switch A and Switch B receive the frame on their 1/1 ports. Because PC-1's MAC address has not yet been seen or recorded, each switch records PC-1's MAC address in its address table along with the receiving port number, 1/1. From this information, both switches infer that PC-1 must reside on Segment A.
- 2 Because PC-4's location is unknown, both switches forward the frame out all available ports, or their 2/1 ports, and onto Segment B.

- 3 Each switch places a new frame on its 2/1 port on Segment B. PC-4, located on Segment B, receives the two frames destined for it. However, Switch A hears the new frame forwarded by Switch B, and Switch B hears the new frame forwarded by Switch A.
- 4 Switch A sees that the “new” frame is from PC-1 to PC-4. From the address table, the switch had learned that PC-1 was on port 1/1 or Segment A. However, the source address of PC-1 has just been heard on port 2/1 on Segment B. By definition, the switch must relearn PC-1’s location, which is now incorrectly assumed to be Segment B. (Switch B follows the same procedure, based on the “new” frame from Switch A.)
- 5 At this point, neither Switch A nor Switch B has learned the location of PC-4 because no frames have been received with PC-4 as the source address. Therefore, the frame must be forwarded out all available ports in an attempt to find PC-4. This frame is then sent out Switch A’s 1/1 port and onto Segment A.
- 6 Now both switches relearn PC-1’s location as Segment A, forward the “new” frames back onto Segment B, and the whole process repeats.

This process of forwarding a single frame around and around between two switches is known as a *bridging loop*. Neither switch is aware of the other, so each just happily forwards the same frame back and forth between its segments. Also note that because two switches are involved in the loop, the original frame has been duplicated and now gets sent around in two counter-rotating loops. What stops the frame from being forwarded in this fashion forever? Nothing. PC-4 will begin receiving frames addressed to it as fast as the switches can forward them.

Notice how the learned location of the PCs keeps changing as frames get looped. Even a unicast frame has caused a bridging loop to form, and each switch’s bridge table is repeatedly corrupted with incorrect data.

What would happen if PC-1 had sent a broadcast frame instead? The bridging loops (remember that there are two of them produced by the two parallel switches) will form exactly as before. The broadcast frames will continue to circulate forever. Now, however, every end-user device located on both Segments A and B will receive and process each and every broadcast frame. This type of broadcast storm can easily saturate the network segments and bring every host on the segments to a halt.

The only way to end the bridging loop is to physically break the loop by disconnecting switch ports or by shutting a switch down. Rather than break devastating bridging loops, they should be prevented instead.

Preventing Loops with Spanning-Tree Protocol

Bridging loops form basically because parallel switches (or bridges) are unaware of each other. STP was developed to overcome the possibility of bridging loops so that redundant switches and switch paths could be used for their benefits. In a nutshell, the protocol enables switches to become aware of each other so that they can negotiate a loop-free path through the network.

Loops are discovered before they are opened for use, and redundant links are shut down to prevent the loops from forming. In the case of redundant links, switches can be made aware that a link shut down for loop prevention should be quickly brought up in case of a link failure. This is discussed in later sections of this chapter.

STP is communicated between all connected switches on a network. Each switch executes the Spanning-Tree Algorithm (STA) based on information received from other neighboring switches. The algorithm chooses a reference point in the network and calculates all the redundant paths to that reference point. When redundant paths are found, STA picks one path to forward frames with and disables or blocks forwarding on the other redundant paths.

As its name implies, STP computes a tree structure that spans all switches in a subnet or network. Redundant paths are placed in a *blocking* or standby state to prevent frame forwarding. The switched network is then in a loop-free condition. However, if a *forwarding* port fails or becomes disconnected, the STA will run again to recompute the Spanning-Tree topology so that blocked links can be reactivated.

Spanning-Tree Communication: Bridge Protocol Data Units

STP operates as switches communicate with one another. Data messages are exchanged in the form of *Bridge Protocol Data Units (BPDUs)*. A switch sends a BPDU frame out a port, using the unique MAC address of the port itself as a source address. The switch is unaware of the other switches around it. Therefore, the BPDU frame has a destination address of the well-known STP multicast address 01-80-c2-00-00-00 to reach all listening switches.

There are two types of BPDU: the *Configuration BPDU*, used for Spanning Tree computation; and the *Topology Change Notification (TCN) BPDU*, used to announce changes in the network topology. The Configuration BPDU message contains the fields shown in Table 5-3. The TCN BPDU is discussed in the “Topology Changes” section later in this chapter.

The exchange of BPDU messages works toward the goal of electing reference points as a foundation for a stable Spanning-Tree topology. As well, loops will be identified and removed by placing specific redundant ports in a blocking or standby state. Notice that several key fields in the BPDU are related to bridge (or switch) identification, path costs, and timer values. These all work together so that the network of switches will converge upon a common Spanning-Tree topology and will select the same reference points within the network. These reference points are defined in the following sections.

BPDUs are sent out all switch ports every two seconds so that current topology information is exchanged and loops are identified quickly.

Table 5-3 Configuration BPDU Message Content

Field Description	Number of Bytes
Protocol ID (always 0)	2
Version (always 0)	1
Message Type (Configuration or Topology Change Notification BPDU)	1
Flags	1
Root Bridge ID	8
Root Path Cost	4
Sender Bridge ID	8
Port ID	2
Message Age (in 256ths of a second)	2
Maximum Age (in 256ths of a second)	2
Hello Time (in 256ths of a second)	2
Forward Delay (in 256ths of a second)	2

Electing a Root Bridge

For all switches in a network to agree on a loop-free topology, a common frame of reference must exist to use as a guide. This reference point is called the *Root Bridge*. (The term “bridge” continues to be used even in a switched environment because STP was developed for use in bridges. Therefore, when you see “bridge,” think “switch.”)

The Root Bridge is chosen by an election process among all connected switches. Each switch has a unique *Bridge ID* that it uses to identify itself to other switches. The Bridge ID is an 8-byte value that is made up of the following fields:

- **Bridge Priority (2 bytes)**—The priority or weight of a switch in relation to all other switches. The priority field can have a value of 0 to 65,535 and defaults to 32,768 (or 0x8000) on every Catalyst switch.
- **MAC Address (6 bytes)**—The MAC address used by a switch can come from the Supervisor module, the backplane, or a pool of 1024 addresses that are assigned to every Supervisor or backplane depending on the switch model. In any event, this address is hardcoded, unique, and cannot be changed by the user.

When a switch first powers up, it has a narrow view of its surroundings and assumes that it is the root bridge itself. Obviously, this notion will probably change as other switches check in and enter the election process. The election process then proceeds as follows: Every switch begins by sending out BPDUs with a Root Bridge ID equal to its own Bridge ID and a Sender

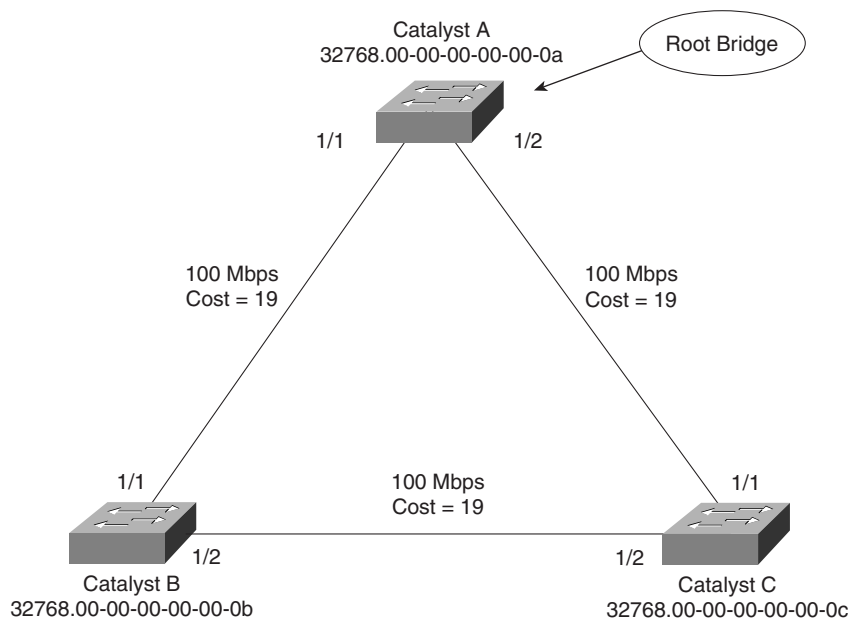
Bridge ID of its own Bridge ID. The Sender Bridge ID simply tells other switches who is the actual sender of the BPDU message.

Received BPDU messages are analyzed to see if a “better” root bridge is being announced. A root bridge is considered better if the Root Bridge ID value is *lower* than another. Again, think of the Root Bridge ID as being broken up into Bridge Priority and MAC address fields. If two Bridge Priority values are equal, then the lower MAC address makes the Bridge ID better. When a switch hears of a better Root Bridge, it replaces its own Root Bridge ID with the Root Bridge ID announced in the BPDU. The switch is then required to nominate the new Root Bridge ID in its own BPDU messages although it will still identify itself as the Sender Bridge ID.

Sooner or later, the election will converge and all switches will agree on the notion that one of them is the Root Bridge. As might be expected, if a new switch with a lower MAC address powers up, it will begin advertising itself as the Root Bridge. Because the new switch does indeed have a lower Bridge ID, all the switches will soon reconsider and record it as the new Root Bridge. Root Bridge election is then an ongoing process, triggered by Root Bridge ID changes in the BPDUs every two seconds.

As an example, consider the small network shown in Figure 5-4. For simplicity, assume that each Catalyst switch has a MAC address of all zeros with the last hex digit equal to the switch label.

Figure 5-4 Example of Root Bridge Election



In this network, each switch has the default Bridge Priority of 32768. The switches are interconnected with Fast Ethernet links, having a default path cost of 19. All three switches try to elect themselves as the root but all of them have equal Bridge Priority values. Therefore, the election is determined by the lowest MAC address—that of Catalyst A.

Electing Root Ports

Now that a reference point has been nominated and elected for the entire switched network, each non-root switch must figure out where it is in relation to the Root Bridge. This action can be performed by selecting only one *Root Port* on each non-root switch.

STP uses the concept of cost to determine many things. Selecting a Root Port involves evaluating the *Root Path Cost*. This value is the cumulative cost of all the links leading to the Root Bridge. A particular switch link has a cost associated with it, too, called the *Path Cost*. To understand the difference between these values, remember that only the Root Path Cost is carried along inside the BPDU. As the path cost travels along, other switches can modify its value to make it cumulative. The Path Cost, however, is not contained in the BPDU. It is known only to the local switch where the port (or “path” to a neighboring switch) resides.

Path Costs are defined as a one-byte value, with the default values shown in Table 5-4. Generally, the higher the bandwidth of a link, the lower the cost of transporting data across it. The original IEEE 802.1D standard defined path cost as 1000 Mbps divided by the link bandwidth in Mbps. These values are shown in the center column of the table. Modern networks commonly use Gigabit Ethernet and OC-48 ATM, which are both either too close to or greater than the maximum scale of 1000 Mbps. The IEEE now uses a non-linear scale for path cost, as shown in the right column of the table.

Table 5-4 STP Path Cost

Link Bandwidth	Old STP Cost	New STP Cost
4 Mbps	250	250
10 Mbps	100	100
16 Mbps	63	62
45 Mbps	22	39
100 Mbps	10	19
155 Mbps	6	14
622 Mbps	2	6
1 Gbps	1	4
10 Gbps	0	2

NOTE Be aware that not all versions of the Catalyst Supervisor code use the newer non-linear scale by default. For example, Catalyst 5000 versions 2.4 and lower use the older linear scale. Catalyst 5000 versions 3.1 and higher, Catalyst 4000 (all versions), and Catalyst 6000 (all versions) use the non-linear scale by default.

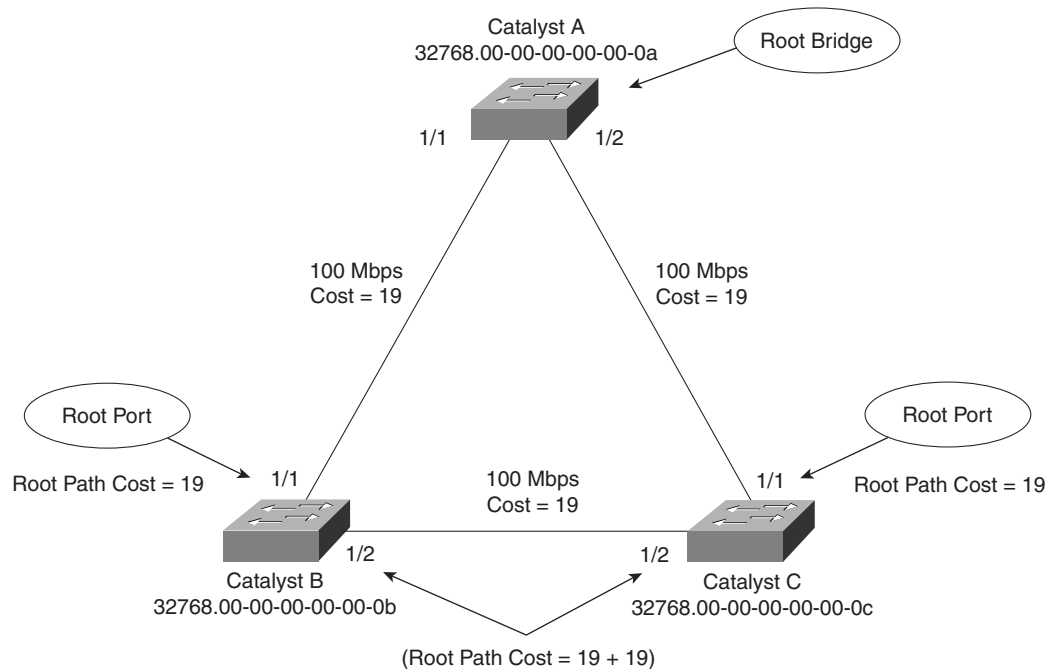
The Root Path Cost value is determined in the following manner:

- 1 The Root Bridge sends out a BPDU with a Root Path Cost value of zero because its ports sit directly on the Root Bridge.
- 2 When the next closest neighbor *receives* the BPDU, it adds the Path Cost of its own port where the BPDU arrived.
- 3 Then the neighbor sends out BPDUs with this new cumulative value as the Root Path Cost.
- 4 This value is incremented by subsequent switch port Path Costs as the BPDU is *received* by each switch on down the line.

NOTE Note the emphasis on incrementing the Root Path Cost as BPDUs are received. When computing the STA manually, remember to compute a new Root Path Cost as BPDUs come in to a switch port, not as they go out.

After incrementing the Root Path Cost, a switch also records the value in its memory. When a BPDU is received on another port and the new Root Path Cost is lower than the previously recorded value, this lower value becomes the new Root Path Cost. In addition, the lower cost tells the switch that the Root Bridge must be closer to this port than it was on other ports. The switch has now determined which of its ports is the closest to the root—the *Root Port*.

Figure 5-5 shows the same network from Figure 5-4 in the process of Root Port selection.

Figure 5-5 Example of Root Port Selection

The Root Bridge, Catalyst A, has already been elected. Therefore, every other switch in the network must choose one port that is closest to the Root Bridge. Catalyst B selects its port 1/1, with a Root Path Cost of 0+19. Port 1/2 is not chosen because its Root Path Cost is 0 (BPDU from Catalyst A) plus 19 (Path Cost of A-C link) plus 19 (Path Cost of C-B link), or a total of 38. Catalyst C makes a similar choice of port 1/1.

Electing Designated Ports

By now, you should begin to see the process unfolding: a starting or reference point has been identified, and each switch “connects” itself toward the reference point with the closest single link. A tree structure is beginning to emerge, but links have only been identified at this point. All links are still connected and could be active, leaving bridging loops.

To remove the possibility of bridging loops, STP makes a final computation to identify one Designated Port on each network segment. Suppose that two or more switches have ports connected to a single common network segment. If a frame appears on that segment, all the bridges will attempt to forward it to its destination. Recall that this behavior was the basis of a bridging loop and should be avoided.

Instead, only one of the links on a segment should forward traffic to and from that segment. This location is the Designated Port. Switches choose a Designated Port based on the lowest cumulative Root Path Cost to the Root Bridge. For instance, a switch always has an idea of its own Root Path Cost, which it announces in its own BPDUs. If a neighboring switch on a shared LAN segment sends a BPDU announcing a lower Root Path Cost, the neighbor must have the Designated Port. If a switch only learns of higher Root Path Costs from other BPDUs received on a port, however, it then correctly assumes that its receiving port is the Designated Port for the segment.

Notice that the whole STP determination process has only served to identify bridges and ports. All ports are still active and bridging loops might still lurk in the network. STP has a set of progressive states that each port must go through, regardless of the type or identification. These states will actively prevent loops from forming and are described in the next section.

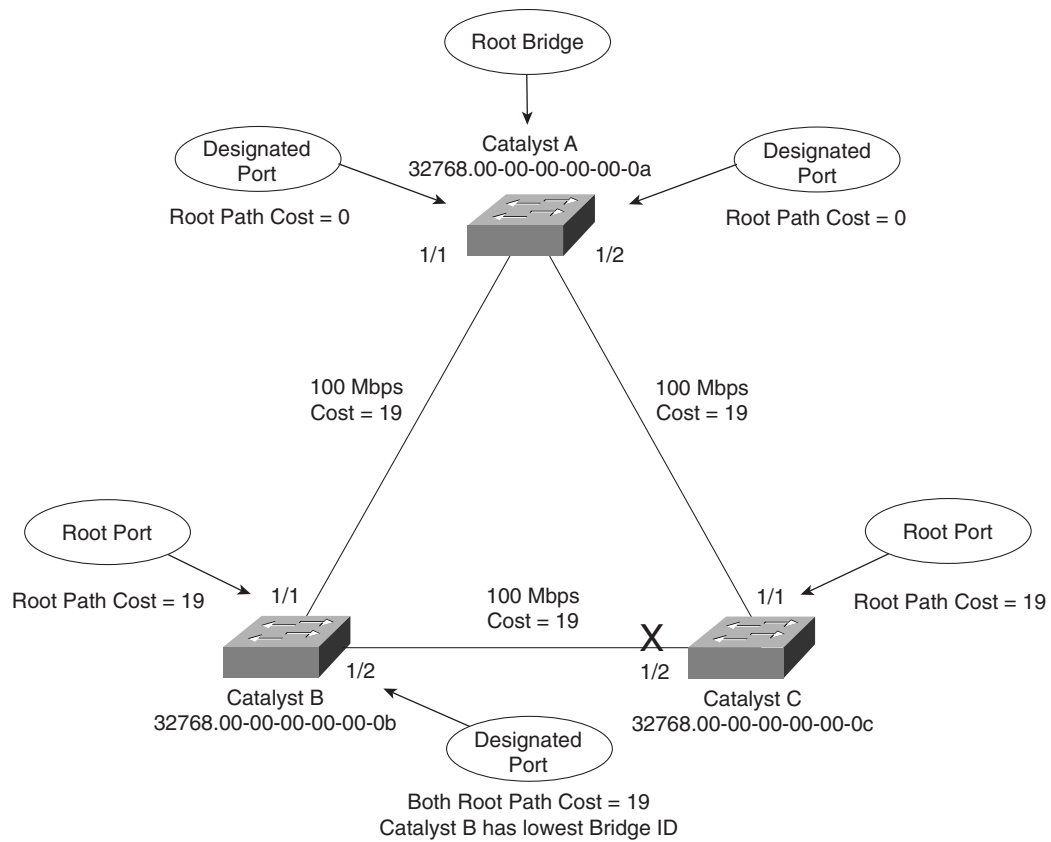
NOTE

In each determination process discussed so far, two or more links to having identical Root Path Costs is possible. This results in a tie condition, unless other factors are considered. In fact, all STP decisions are based on the following sequence of four conditions:

1. Lowest Root Bridge ID
 2. Lowest Root Path Cost to Root Bridge
 3. Lowest Sender Bridge ID
 4. Lowest Port ID
-

Figure 5-6 demonstrates an example of Designated Port selection. This figure is identical to Figure 5-4 and Figure 5-5, with further Spanning Tree development. The only changes shown are the choices of Designated Ports, although seeing all STP decisions shown in one network diagram is handy.

Figure 5-6 Example of Designated Port Selection



The three switches have chosen their Designated Ports (DP) for the following reasons:

- **Catalyst A**—Because this switch is the Root Bridge, all its active ports are Designated Ports by definition. At the Root Bridge, the Root Path Cost of each port is zero.
- **Catalyst B**—Catalyst A port 1/1 is the DP for the Segment A-B because it has the lowest Root Path Cost (0). Catalyst B port 1/2 is the DP for segment B-C. The Root Path Cost for each end of this segment is 19, determined from the incoming BPDU on port 1/1. Because the Root Path Cost is equal on both ports of the segment, the DP must be chosen by the next criteria—the lowest Sender Bridge ID. When Catalyst B sends a BPDU to Catalyst C, it has the lowest MAC address in the Bridge ID. Catalyst C also sends a BPDU to Catalyst B, but its Sender Bridge ID is higher. Therefore, Catalyst B port 1/2 is selected as the DP of the segment.

- **Catalyst C**—Catalyst A port 1/2 is the DP for Segment A-C because it has the lowest Root Path Cost (0). Catalyst B port 1/2 is the DP for Segment B-C. Therefore, Catalyst C port 1/2 will be neither a Root Port nor a Designated Port. As discussed in the next section, any port that is not elected to either position will enter the *blocking state*. Where blocking occurs, bridging loops are broken.

STP States

To participate in STP, each port of a switch must progress through several states. A port begins its life in a Disabled state moving through several passive states and finally into an active state if allowed to forward traffic. The STP port states are as follows:

- **Disabled**—Ports that are administratively shut down by the network administrator or by the system due to a fault condition are in the Disabled state. This state is special and is not part of the normal STP progression for a port.
- **Blocking**—After a port initializes, it begins in the Blocking state so that no bridging loops can form. In the Blocking state, a port cannot receive or transmit data and cannot add MAC addresses to its address table. Instead, a port is only allowed to receive BPDUs so that the switch can hear from other neighboring switches. In addition, ports that are put into standby mode to remove a bridging loop enter the Blocking state.
- **Listening**—The port will be moved from Blocking to Listening if the switch thinks that the port can be selected as a Root Port or Designated Port. In other words, the port is on its way to begin forwarding traffic. In the Listening state, the port still cannot send or receive data frames. However, the port is allowed to receive and send BPDUs so that it can actively participate in the Spanning-Tree topology process. Here the port is finally allowed to become a Root Port or Designated Port because the switch can advertise the port by sending BPDUs to other switches. Should the port lose its Root Port or Designated Port status, it is returned to the Blocking state.
- **Learning**—After a period of time called the *Forward Delay* in the Listening state, the port is allowed to move into the Learning state. The port still sends and receives BPDUs as before. In addition, the switch can now learn new MAC addresses to add into its address table. This gives the port an extra period of silent participation and allows the switch to assemble at least some address table information.
- **Forwarding**—After another Forward Delay period of time in the Learning state, the port is allowed to move into the Forwarding state. The port can now send and receive data frames, collect MAC addresses into its address table, and send and receive BPDUs. The port is now a fully functioning switch port within the Spanning-Tree topology.

NOTE Remember that a switch port is only allowed into the Forwarding state if there are no redundant links (or loops) and if the port has the best path to the root bridge as the Root Port or Designated Port.

Example 5-3 shows the output of a switch as one of its ports progresses through the STP port states.

Example 5-3 *A Port Progressing Through the STP Port States*

```

Console> (enable) set port disable 4/10
This command may disconnect your telnet session.
Do you want to continue (y/n) [n]?y
Port 4/10 disabled.
Console> (enable) set port enable 4/10
Port 4/10 enabled.
Console> (enable) show spant 4/10
Port      Vlan  Port-State   Cost   Priority  Fast-Start
-----
4/10     1    listening    10     32       disabled
Console> (enable) sh spant 4/10
Port      Vlan  Port-State   Cost   Priority  Fast-Start
-----
4/10     1    listening    10     32       disabled
Console> (enable) sh spant 4/10
Port      Vlan  Port-State   Cost   Priority  Fast-Start
-----
4/10     1    listening    10     32       disabled
Console> (enable) sh spant 4/10
Port      Vlan  Port-State   Cost   Priority  Fast-Start
-----
4/10     1    learning     10     32       disabled
Console> (enable) sh spant 4/10
Port      Vlan  Port-State   Cost   Priority  Fast-Start
-----
4/10     1    learning     10     32       disabled
Console> (enable) sh spant 4/10
Port      Vlan  Port-State   Cost   Priority  Fast-Start
-----
4/10     1    learning     10     32       disabled
Console> (enable) sh spant 4/10
Port      Vlan  Port-State   Cost   Priority  Fast-Start
-----
4/10     1    learning     10     32       disabled
Console> (enable) sh spant 4/10
Port      Vlan  Port-State   Cost   Priority  Fast-Start
-----
4/10     1    forwarding   10     32       disabled
Console> (enable)

```

The example begins as the port is administratively disabled from the command line. When the port is enabled, successive **show spantree** *module/port* commands display the Port-State as Listening, Learning, and then Forwarding. Because this port was eligible as a Root Port, the **show** command was never able to execute fast enough to show the port in the Blocking state.

STP Timers

STP operates as switches send BPDUs to each other in an effort to form a loop-free topology. The BPDUs take a finite amount of time to travel from switch to switch. In addition, news of a topology change (such as a link or Root Bridge failure) can suffer from propagation delays as the announcement travels from one side of a network to the other. Because of the possibility of these delays, keeping the Spanning-Tree topology from settling out or converging until all switches have had time to receive accurate information is important.

STP uses three timers to make sure that a network converges properly before a bridging loop can incorrectly form. The timers and their default values are as follows:

- **Hello Time**—The time interval between Configuration BPDUs sent by the Root Bridge. The Hello Time value configured in the Root Bridge switch will determine the Hello Time for all non-root switches because they just relay the Configuration BPDUs as they are received from the root. However, all switches have a locally configured Hello Time that is used to time TCN BPDUs when they are retransmitted. The IEEE 802.1D standard specifies a default Hello Time value of two seconds.
- **Forward Delay**—The time interval that a switch port spends in both the Listening and Learning states. The default value is 15 seconds.
- **Maximum (max) Age**—The time interval that a switch stores a BPDU before discarding it. While executing the STP, each switch port keeps a copy of the “best” BPDU that it has heard. If the source of the BPDU loses contact with the switch port, the switch will notice that a topology change has occurred after the Max Age time elapses and the BPDU is aged out. The default Max Age value is 20 seconds.

The STP timers can be configured or adjusted from the switch command line. However, the timer values should never be changed from the defaults without careful consideration. Then, the values should only be changed on the Root Bridge switch. Recall that the timer values are advertised in fields within the BPDU. The Root Bridge will make sure that the timer values are propagated to all other switches.

NOTE

The default STP timer values are based on some assumptions about the size of the network and the length of the Hello Time. A reference model of a network having a diameter of seven switches is used to derive these values. The diameter is measured from the Root Bridge switch outward, including the Root Bridge. A Hello Time of two seconds is used in this computation.

The network diameter can be configured on the Root Bridge switch to more accurately reflect the true size of the physical network. Making that value more accurate will reduce the total STP convergence time during a topology change. As well, Cisco recommends that if changes need to be made, only the network diameter value should be modified on the Root Bridge switch. When the diameter is changed, the switch will calculate new values for all three timers. This option is discussed in the “Selecting the Root Bridge” section in this chapter.

Topology Changes

To announce a change in the active network topology, switches send a Topology Change Notification (TCN) BPDU. Table 5-5 shows the format of these messages.

Table 5-5 *Topology Change Notification BPDU Message Content*

Field Description	# Bytes
Protocol ID (always 0)	2
Version (always 0)	1
Message Type (Configuration or Topology Change Notification BPDU)	1

A topology change occurs when a switch either moves a port into the Forwarding state or moves a port from Forwarding or Learning into the Blocking state. In other words, a port on an active switch comes up or goes down. The switch sends a TCN BPDU out its Designated Port so that ultimately the Root Bridge will receive news of the topology change. Notice that the TCN BPDU carries no data about the change, but only informs recipients that a change has occurred. Also notice that the switch will not send TCN BPDUs if the port has been configured with PortFast enabled.

The switch will continue sending TCN BPDUs every Hello Time interval until it gets an acknowledgement from an upstream neighbor. As the upstream neighbors receive the TCN BPDU, they will propagate it on toward the Root Bridge. When the Root Bridge receives the BPDU, the Root Bridge also sends out an acknowledgement. However, it also sends out the Topology Change flag in a Configuration BPDU so that all other bridges will shorten their bridge table aging times down from the default (300 seconds) to just the Forward Delay value (default 15 seconds).

This condition causes the learned locations of MAC addresses to be flushed out much sooner than they normally would, easing the bridge table corruption that might occur due to the change in topology. However, any stations that are actively communicating during this time will be kept in the bridge table. This condition lasts for the sum of the Forward Delay and the Max Age (default 15 + 20 seconds).

Spanning-Tree Design

STP and its computations are very predictable. However, other factors exist that may subtly influence STP decisions, making the resulting tree structure neither expected nor ideal. For example, several versions of Spanning Tree exist and are used by various vendors. Interoperability of these versions could be important in a mixed-vendor network.

The network administrator can also make adjustments to the Spanning-Tree operation to control its behavior. The location of the Root Bridge should be determined as part of the design process. As well, redundant links can be used for load balancing in parallel if configured correctly. Spanning Tree can also be configured to converge quickly and predictably in the event of a major topology change.

Types of STP

So far, this chapter has discussed STP in terms of its operation to prevent loops and to recover from topology changes in a timely manner. STP was originally developed to operate in a bridged environment, basically supporting a single LAN (or one VLAN). Implementing STP into a switched environment has required additional consideration and modification to support multiple VLANs. Because of this, the IEEE and Cisco have approached STP differently. This section reviews the three types of STP that are encountered in switched networks and how they relate to one another. There are no specific configuration commands associated with the various types of STP. Rather, you should have a basic understanding of how they interoperate in a network.

Common Spanning Tree (CST)

The IEEE 802.1Q standard specifies how VLANs are to be trunked between switches. As well, it specifies only a single instance of STP for all VLANs. This instance is referred to as the *Common Spanning Tree (CST)* or the *Mono Spanning Tree (MST)*. All BPDUs are transmitted over VLAN 1, the management VLAN.

Having a single STP for many VLANs simplifies switch configuration and reduces switch CPU load during STP calculations. However, the STP can cause limitations, too. Redundant links between switches will be blocked with no capability for load balancing. Conditions can also occur that would cause forwarding on a link that doesn't support all VLANs, while other links would be blocked.

Per-VLAN Spanning Tree (PVST)

Cisco has a proprietary STP that offers more flexibility than the CST method. *Per-VLAN Spanning Tree (PVST)* operates a separate instance of STP for each individual VLAN. This allows the STP on each VLAN to be configured independently, offering better performance and

tuning for specific conditions. Multiple Spanning Trees also make load balancing possible over redundant links when the links are assigned to different VLANs.

Due to its proprietary nature, PVST requires the use of Cisco Inter-Switch Link (ISL) trunking encapsulation between switches. In networks where PVST and CST coexist, interoperability problems occur. Each requires a different trunking method so BPDUs will never be exchanged between STP types.

Per-VLAN Spanning Tree Plus (PVST+)

Cisco has a second proprietary STP that allows devices to interoperate with both PVST and CST. *Per-VLAN Spanning Tree Plus (PVST+)* effectively supports three groups of STP operating in the same campus network: Catalysts running PVST; Catalysts running PVST+; and switches running CST/MST over 802.1Q.

To do this, PVST+ acts as a translator between groups of CST switches and groups of PVST switches. PVST+ can communicate directly with PVST by using ISL trunks. To communicate with CST, however, PVST+ exchanges BPDUs with CST on VLAN 1. BPDUs from other instances of STP (other VLANs) are propagated across the CST portions of the network by tunneling. PVST+ sends these BPDUs by using a unique multicast address so that the CST switches will forward them on to downstream neighbors. Eventually, the tunneled BPDUs will reach other PVST+ switches where they are understood.

STP Configuration

By default, STP is enabled on all ports of a switch. STP should remain enabled in a network to prevent bridging loops from forming. However, if STP has been disabled, it can be re-enabled with the commands documented in the list that follows:

- **STP Configuration on a CLI-Based Switch**—STP can be enabled on either a port, a range of ports, or on all ports and VLANs by using the following command:

```
Switch (enable) set spantree enable [all | module/port]
```

- **STP Configuration on an IOS-Based Switch**—To enable STP on an IOS-based switch, where *vlan-list* is the list of VLANs that should have STP enabled, use the following command:

```
Switch (config)# spantree vlan-list
```

To view the status of STP on either a CLI- or IOS-based switch, use the following command:

```
Switch (enable) show spantree [vlan]
```

The output in Example 5-4 shows the current values for all elections, costs, timers, Bridge IDs, and STP state on each port for a VLAN.

Example 5-4 *show spantree Command Output Displays STP Status on a Switch*

```

Switch (enable) show spantree 10
VLAN 10
Spanning tree enabled
Spanning tree type          ieee
Designated Root             00-50-a2-8d-58-09
Designated Root Priority    32768
Designated Root Cost        0
Designated Root Port        1/0
Root Max Age 20 sec        Hello Time 2 sec    Forward Delay 15 sec
Bridge ID MAC ADDR          00-50-a2-8d-58-09
Bridge ID Priority           32768
Bridge Max Age 20 sec        Hello Time 2 sec    Forward Delay 15 sec
Port      Vlan  Port-State      Cost  Priority  Fast-Start  Group-Method
-----
  1/2      10   forwarding      4     32     disabled
  9/1-2    10   forwarding      14    32     disabled    redundancy
Switch (enable)

```

Root Bridge Placement

While STP is wonderfully automatic with its default values and election processes, the resulting tree structure may perform quite differently than expected. The Root Bridge election is based on the idea that one switch is chosen as a common reference point, and all other switches choose ports that are closest to the Root. The Root Bridge election is also based on the idea that the Root Bridge can become a central hub that interconnects other legs of the network. Therefore, the Root Bridge can be faced with heavy switching loads in its central location.

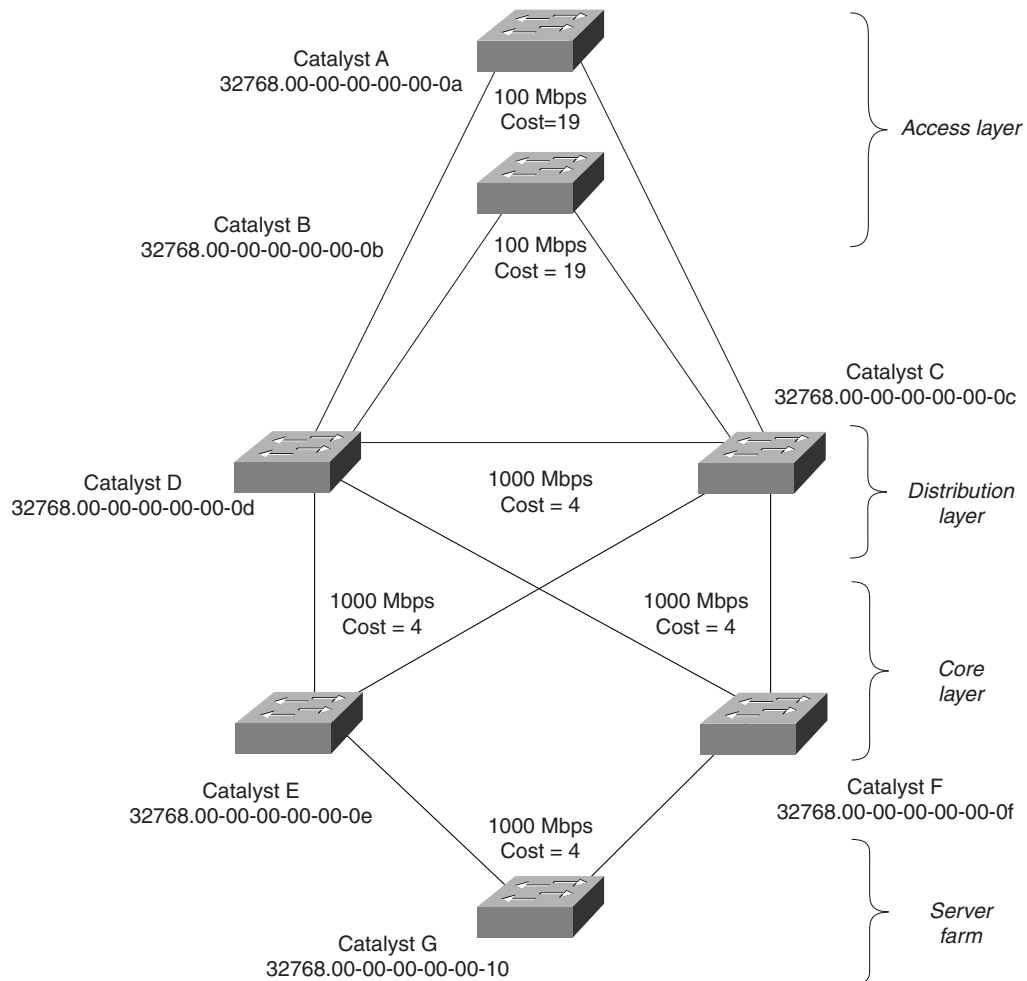
If the Root Bridge election is left to its default state, several things can occur to make a poor choice. For example, the slowest switch (or bridge) can be elected as the Root Bridge. If heavy loads of traffic are expected to pass through the Root Bridge, the slowest switch is not the ideal candidate. Recall that the only criteria for Root Bridge election is the lowest Bridge ID (Bridge Priority and MAC address)—not necessarily the best choice to ensure optimal performance. If the slowest switch has the same Bridge Priority as the others and has the lowest MAC address, the slowest switch will be chosen as the Root.

A second factor to consider relates to redundancy. If all switches are left to their default states, only one Root Bridge will be elected. What will happen if that switch fails? Another Root Bridge election will occur but again the choice might not be the ideal switch or the ideal location.

The final consideration is the location of the Root Bridge switch. As before, an election with default switch values could place the Root Bridge in an unexpected location in the network. More important, a very inefficient Spanning-Tree structure could result causing traffic from a large portion of the network to take a long and winding path just to pass through the Root Bridge.

Figure 5-7 shows a portion of a real-world hierarchical campus network.

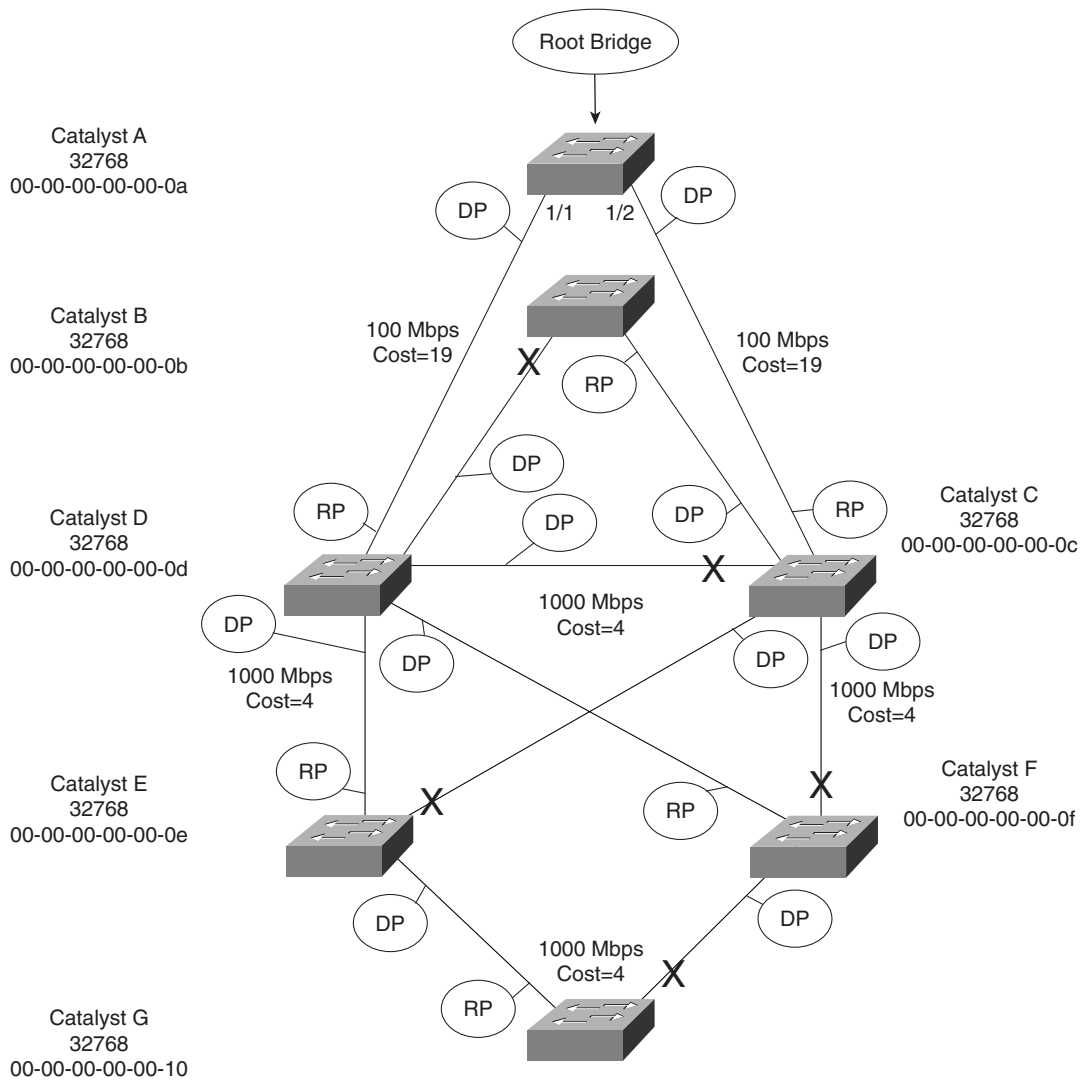
Figure 5-7 A Campus Network with an Inefficient Root Bridge Election



Catalyst switches A and B are two Access layer devices; Catalysts C and D are in the Distribution layer; Catalysts E and F form the Core layer; and Catalyst G connects a server farm into the network core. Notice that all the switches use redundant links to other layers of the hierarchy, as suggested in Chapter 2, “Campus Network Design Models.” As will be seen, Catalyst A will become the Root Bridge due to its low MAC address. The STP process will begin to develop.

Figure 5-8 shows the converged state of STP. For the purposes of this discussion, the Root Ports and Designated Ports are simply shown on the network diagram. (As an exercise, you should work out the Spanning Tree based on the information shown in the figure.)

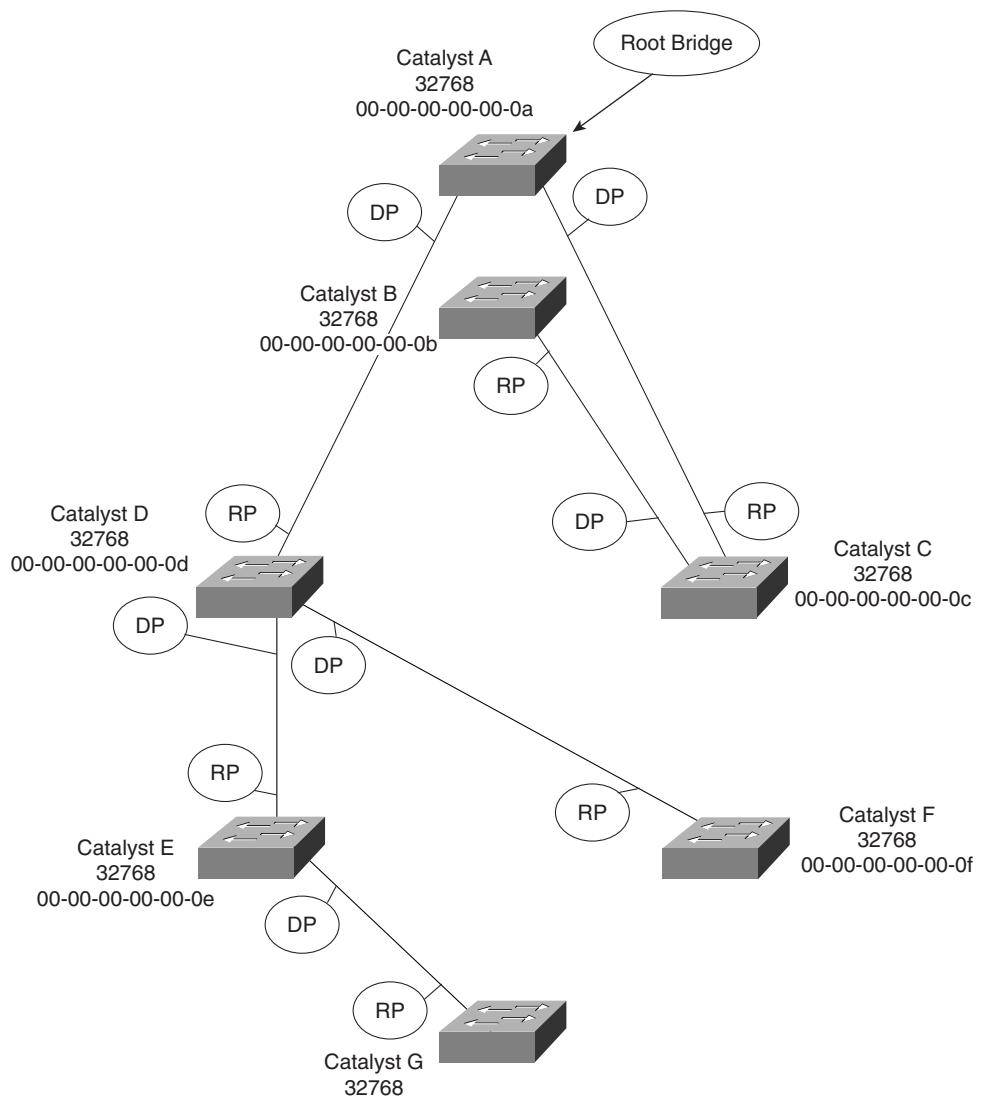
Figure 5-8 A Campus Network with STP Converged



Notice that Catalyst A, one of the Access layer switches, has been elected as the Root Bridge. Also note the location of the X symbols over the ports that are neither Root Ports nor Designated Ports. These ports will enter the Blocking state.

Finally, Figure 5-9 shows the same network with the Blocking links removed. Now you can see the true structure of the final Spanning Tree.

Figure 5-9 Final Spanning Tree Structure for the Campus Network



Catalyst A, an Access layer switch, is the Root Bridge. Workstations on Catalyst A can reach servers on Catalyst G by crossing the Distribution layer and then the Core layer, as expected. However, notice what has happened to the other Access layer switch, Catalyst B. Workstations on this switch must cross into the Distribution layer (Catalyst C), back into the Access layer (Catalyst A), back into the Distribution layer (Catalyst D), through the Core, and finally to the server farm (Catalyst G). This action is obviously very inefficient. In addition, Catalyst A is probably not a high-end switch because it is used in the Access layer. The structure of the Spanning Tree has forced Catalyst A to carry much more than its expected load and it will likely become a bottleneck.

Root Bridge Configuration

To prevent the surprises outlined in the previous section, you should always do two things:

- Set a Root Bridge in a determined fashion.
- Set a secondary Root Bridge in case of primary Root Bridge failure.

As the common reference point, the Root Bridge (and the secondary) should be placed near the center of the network. For example, a switch in the Distribution layer would make a better Root Bridge choice than one in the Access layer because more traffic is expected to pass through the Distribution layer devices. In a flat switched network (no Layer 3 devices), a switch near a server farm would be a more efficient Root Bridge than switches elsewhere. Most traffic will be destined to and from the server farm and will benefit from a predetermined, direct path.

To configure a Catalyst switch to become the Root Bridge, use the commands described in the sections that follow.

Root Bridge Configuration on a CLI-Based Switch

CLI-based Catalysts offer two commands that can be used to modify the Root Bridge selection. The **set spantree priority** command directly modifies the Bridge Priority value so that a switch can be given a lower Bridge ID value to win a Root Bridge election:

```
Switch (enable) set spantree priority bridge-priority [vlan]
```

The *bridge-priority* value defaults to 32768, but can be given a value of 0 to 65535. Although the *vlan* parameter is optional, it should always be specified when using this command. Remember that Catalyst switches run one instance of STP for each VLAN (PVST). Therefore, you should designate an appropriate Root Bridge for each VLAN. If the *vlan* parameter is not used with this command, the Bridge Priority will be modified only for VLAN 1 (the management VLAN).

The other alternative for modifying the Root Bridge selection is the **set spantree root** command. This command is actually a macro on the Catalyst that executes several other commands. The result is a more direct and automatic way to force one switch to become the

Root Bridge. Actual Bridge Priorities are not given in the command. Rather, the switch will modify STP values according to the current values in use within the active network. The command syntax is as follows:

```
Switch (enable) set spantree root [secondary] [vlan-list] [dia diameter] [hello  
hello_time]
```

This command modifies the Bridge Priority value of the switch to become less than the Bridge Priority of the current Root Bridge. If the current root priority is more than 8192, the local switch will have its priority set to 8192. If the current root priority is less than 8192, however, the command will set the local switch priority to some value (1000, 100, or 1) less than the current root. For example, if the current root has a priority of the default 32768, this command will modify the local Bridge Priority to be 8192. If the current root is 8192, the local Bridge Priority might be 7192, and so on.

This command can also be used to configure one or more secondary Root Bridges for specific VLANs. By including the **set spantree root secondary** option, the local switch will receive a Bridge Priority of 16384—a value less than the default 32768 of other switches but greater than 8192 (or less) of the Root Bridge.

The network diameter and the STP Hello Time can also be modified with this command, if needed. This modification is discussed further in the “Tuning Spanning-Tree Convergence” section later in the chapter.

NOTE

The **set spantree root** command will not be shown in a Catalyst switch configuration because the command is actually a macro executing other switch commands. The actual commands and values used by the macro will be shown, however.

Root Bridge Configuration on an IOS-Based Switch

On an IOS-based switch, the following command can be used to modify the Root Bridge selection:

```
Switch (config)# spanning-tree [vlan vlan-list] priority bridge-priority
```

The Bridge Priority for the VLANs specified is modified to the *bridge-priority* value. Catalyst switches default to 32768, but can be given a priority of 0 to 65535. You should always specify the list of VLANs to modify for PVST. Otherwise, the command will modify the Bridge Priority only for VLAN 1.

Spanning-Tree Customization

The most important decision you can make when designing your Spanning-Tree topology is the placement of the Root Bridge. Other decisions, such as the exact loop-free path structure, will

occur automatically as a result of the STA. Occasionally, the path may need additional tuning, but only under special circumstances and after careful consideration.

Recall the sequence of four criteria that STP uses to choose a path:

- 1 Lowest Bridge ID
- 2 Lowest Root Path Cost
- 3 Lowest Sender Bridge ID
- 4 Lowest Port ID

The previous section discussed how to tune the Bridge ID of a switch to place the Root Bridge in a network. This technique can be used to force a switch to have the lowest Bridge ID and also to influence the sending Bridge ID of other switches (Lowest Bridge ID and Lowest Sender Bridge ID). However, only the automatic STP computation has been discussed, using the default switch port costs to make specific path decisions.

Tuning the Root Path Cost

The Root Path Cost for each active port of a switch is determined by the cumulative cost as a BPDU travels along. As a switch receives a BPDU, the port cost of the receiving port is added to the Root Path Cost in the BPDU. The port or path cost is inversely proportional to the port's bandwidth, as listed previously in Table 5-4. If desired, the cost of a port can be modified from the default value.

NOTE

Before modifying the path cost of a switch port, you should always calculate the Root Path costs of other alternate paths through the network. Changing one port's cost may influence STP to choose that port as a Root Port but there could be other paths that are still preferred. You should also calculate a port's existing path cost to determine what the new cost value should be. Careful calculation will ensure that the desired path will indeed be chosen.

Tuning the Root Path Cost on a CLI-Based Switch

The port cost can be modified on a CLI-based switch by using one of the following commands:

```
Switch (enable) set spantree portcost module/port cost  
Switch (enable) set spantree portvlancost module/port [cost cost] [vlan-list]
```

The **set spantree portcost** command modifies the cost for a port, regardless of the VLANs that are assigned to it. However, recall that PVST can create quite different Spanning-Tree topologies for each VLAN. If it is necessary to tune the port cost for a specific VLAN, use the **set spantree portvlancost** command and specify the VLAN number. This command only

applies to trunk ports, where multiple VLANs are in use. The port cost value can range from 0 to 65535 and defaults to 4 if not specified.

Tuning the Root Path Cost on an IOS-Based Switch

IOS-based switches offer a single command that can modify the port cost for individual VLANs. Use the following command:

```
Switch (config-if)# spanning-tree [vlan vlan-list] cost cost
```

If the **vlan** parameter is given, the port cost will be modified only for the specified VLANs. If it is not specified, the cost is modified only for VLAN 1. The cost value ranges from 1 to 65535 and defaults to the standard IEEE values listed previously in Table 5-4.

Tuning the Port ID

The fourth criteria of an STP decision is the Port ID. The Port ID value that a switch uses is actually a 16-bit quantity: 6 bits for the Port Priority and 10 bits for the Port Number. Port Priority is a value from 0 to 63 and defaults to 32 for all ports. The Port Number can range from 0 to 1023 and represents the actual physical mapping of the port. Port Numbers begin with 1 at port 1/1 and increment across each module. (The numbers may not be consecutive because each module is assigned a particular range of numbers.)

Obviously, the Port Number of a switch port is fixed because it is based on hardware location. The Port ID, however, can be modified to influence an STP decision by using the Port Priority.

Tuning the Port ID on a CLI-Based Switch

The Port ID can be tuned on a CLI-based switch by using one of the following commands:

```
Switch (enable) set spantree portpri {module/port} priority  
Switch (enable) set spantree portvlanpri {module/port} priority [vlan]
```

The **set spantree portpri** command is used to set the port priority value on a port for all VLANs. In a PVST environment where multiple instances of STP are running for multiple VLANs, the **set spantree portvlanpri** command should be used. The port priority then can be tailored for each specific VLAN that the port participates. The *priority* value ranges from 0 to 63 and defaults to 32.

Tuning the Port ID on an IOS-Based Switch

Use the following command to modify the Port Priority on an IOS-based switch:

```
Switch(config-if)# spanning-tree [vlan vlan-list] port-priority port-priority
```

The Port Priority can be modified for specific VLANs by using the **vlan** parameter. If this is not specified, the Port Priority is set only for VLAN 1. The value of *port-priority* can range from 0 to 255 and defaults to 128.

NOTE Notice that range of Port Priority values differs between CLI- (0–63) and IOS-based (0–255) switches. IOS-based switches (and routers) adhere to the 802.1D standard, which specifies an 8-bit Port Priority and an 8-bit Port Number field in the Port ID. Higher-end Catalyst switches move the boundary between fields so that the Port Number field has a greater range of values, supporting higher densities of switch ports.

Viewing STP Status

STP costs for a switch port can be viewed by using the **show spantree module/port** command. Example 5-5 demonstrates sample output from this command. Switch port 3/29 is a member of VLAN 53, is in the Forwarding state, has a Port Cost of 19, and a Port Priority of 32. In addition, the port has STP PortFast enabled.

Example 5-5 **show spantree** Command Output Displays STP Status of a Switch Port

Port	Vlan	Port-State	Cost	Priority	Portfast	Channel_id
3/29	53	forwarding	19	32	enabled	0

Tuning Spanning-Tree Convergence

STP uses several timers, a sequence of states that ports must move through, and specific topology change conditions to prevent bridging loops from forming in a complex network. Each of these parameters or requirements is based on certain default values for a typical network size and function. For the majority of cases, the default STP operation is sufficient to keep the network loop-free and enable users to communicate.

However, certain situations occur when the default STP can cause network access to be delayed while timers expire and while preventing loops on links where loops are not possible. It is safe then to make adjustments to the STP convergence process for more efficiency.

Modifying STP Timers

Recall that STP uses three timers to keep track of various port operation states and communication between bridges. The three STP timers can be adjusted by using the commands documented in the sections that follow. Remember that the timers need only be modified on the Root Bridge (and any secondary or backup Root Bridges) because the Root Bridge propagates all three timer values throughout the network as fields in the Configuration BPDU.

Modifying STP Timers on a CLI-Based Switch

Use the following commands to modify STP timers on a CLI-based switch:

```
Switch(enable) set spantree hello interval [vlan]
Switch(enable) set spantree fwddelay delay [vlan]
Switch(enable) set spantree maxage agingtime [vlan]
```

The *Hello Timer* triggers periodic hello messages to be sent to other bridges and sets the interval that a bridge expects to hear a hello from its neighboring bridges. BPDUs are sent every 2 seconds by default. The Hello Timer can be modified per VLAN with the **set spantree hello** command with a range of 1 to 10 seconds.

The *Forward Delay Timer* determines the amount of time a port stays in the Listening state before moving into the Learning state and how long it stays in the Learning state before moving to the Forwarding state. The Forward Delay Timer can be modified per VLAN with the **set spantree fwddelay** command. The default value is 15 seconds and can be set to a value of 4 to 30 seconds. This timer should only be modified under careful consideration because the value is dependent upon the diameter of the network and the propagation of BPDUs across all switches. A value too low will allow loops to form and cripple a network.

The *MaxAge Timer* specifies the lifetime of a stored BPDU that has been received from a neighboring switch with a Designated Port. Suppose BPDUs are being received on a non-Designated switch port every 2 seconds, as expected. Then an *indirect failure*, or one that doesn't involve a physical link going down, occurs that prevents BPDUs from being sent. The receiving switch will wait until the Max Age Timer expires to listen for further BPDUs. If none are received, the non-Designated port will move into the Listening state and Configuration BPDUs will be generated by the receiving switch. This port then becomes the Designated Port to restore connectivity on the segment.

To modify the Max Age Timer on a per-VLAN basis, use the **set spantree maxage** command. The timer value defaults to 20 seconds, but can be set from 6 to 40 seconds.

NOTE

Modifying STP timers can be tricky given the conservative nature of the default values and the calculations needed to derive proper STP operation. Timer values are basically dependent on the Hello Time and the diameter of the switched network, in terms of switch hops. Catalyst CLI-based switches offer a single command that can be used to change the timer values in a more controlled fashion. Although described earlier, the **set spantree root** macro command is a better tool to use than setting the timers with the individual commands:

```
Switch (enable) set spantree root [secondary] [vlan-list] [dia diameter] [hello
hello-time]
```

Here, STP timers will be adjusted exactly according to the formulas specified in the 802.1D standard by giving only the Hello Time and the diameter of the network. Again, this command can be used on a per-VLAN basis to modify the timers for a particular VLAN's Spanning Tree. The network diameter can be a value from one to seven switch hops. Because this command is used to make a switch become the Root Bridge, all the modified timer values resulting from this command will be propagated to other switches through the Configuration BPDU.

Modifying STP Timers on IOS-Based Switches

The following commands can be used on an IOS-based switch to modify STP timers:

```
Switch(config)# spanning-tree [vlan vlan-list] hello-time seconds  
Switch(config)# spanning-tree [vlan vlan-list] forward-time seconds  
Switch(config)# spanning-tree [vlan vlan-list] max-age seconds
```

The Hello Timer defaults to 2 seconds and can be set from 1 to 10 seconds per VLAN. The Forward Timer defaults to 15 seconds and can be set from 4 to 200 seconds per VLAN. The Max Age Timer defaults to 20 seconds and can be set from 6 to 200 seconds per VLAN. If the **vlan** parameter is not specified for any of these commands, VLAN 1 will be assumed.

Redundant Link Convergence

Some additional methods that exist to allow faster STP convergence in the event of a link failure include

- **PortFast**—Enables fast connectivity to be established on access layer switch ports to workstations that are booting up.
- **UplinkFast**—Enables fast uplink failover on an access layer switch when dual uplinks are connected into the distribution layer.
- **BackboneFast**—Enables fast convergence in the network backbone (core) after a Spanning-Tree topology change occurs.

Rather than modifying timer values, these methods work by controlling convergence on specifically located ports within the network hierarchy.

PortFast: Access Layer Nodes

An end-user workstation is usually connected to a switch port in the Access layer. If the workstation is powered off and then turned on, the switch port will not be in a useable state until STP cycles from the Blocking state to the Forwarding state. With the default STP timers, this transition will take at least 30 seconds (15 seconds Listening to Learning and 15 seconds Learning to Forwarding). Therefore, the workstation is unable to transmit or receive any useful data until the Forwarding state is reached on the port.

NOTE

Port initialization delays of up to 50 seconds can be observed. As discussed, 30 of these seconds are due to the STP state transitions. If a switch port is running PAgP to negotiate EtherChannel configuration, an additional 20-second delay can exist.

On switch ports that connect only to single workstations or specific devices, bridging loops will never be possible. Catalyst switches offer the PortFast feature that shortens the Listening and Learning states to a negligible amount of time. The result is that when a workstation link comes up, the switch will immediately move the PortFast port into the Forwarding state. Spanning-Tree loop detection is still in operation, however, and the port will be moved into the Blocking state if a loop is ever detected on the port.

To enable or disable the PortFast feature on a CLI-based switch port, use the following command:

```
Switch(enable) set spantree portfast {module/port} {enable | disable}
```

On an IOS-based switch, use this command:

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

Obviously, you should not enable PortFast on a switch port that is connected to a hub or another switch because bridging loops could possibly form. One other benefit of PortFast is that TCN BPDUs are not sent when a switch port in PortFast mode goes up or down. This simplifies the TCN transmission on a large network when end-user workstations are coming up or shutting down.

To view the PortFast state of switch ports, use the **show spantree** command. Each port is listed, along with the PortFast information under the “Fast-Start” column.

UplinkFast: Access Layer Uplinks

Consider an Access layer switch that has redundant uplink connections to two Distribution layer switches. Normally, one uplink would be in the Forwarding state and the other in the Blocking state. If the primary uplink went down, up to 50 seconds would elapse before the redundant uplink could be used.

The *UplinkFast* feature on Catalyst switches enables leaf-node switches or switches at the ends of the Spanning-Tree branches to have a functioning Root Port while keeping one or more redundant or potential Root Ports in Blocking mode. When the primary Root Port uplink fails, another blocked uplink can be immediately brought up for use.

To enable or disable the UplinkFast feature on a CLI-based switch, use the following command:

```
Switch (enable) set spantree uplinkfast {enable | disable} [rate update-rate] [all-protocols off | on]
```

On an IOS-based switch, use the following command:

```
Switch (config)# spanning-tree uplinkfast [max-update-rate pkts-per-second]
```

When UplinkFast is enabled, it is enabled for the whole switch and all VLANs. UplinkFast works by keeping track of possible paths to the Root Bridge. Therefore, the command is not allowed on the Root Bridge switch. UplinkFast also makes some modifications to the local switch to insure that it does not become the Root Bridge and that the switch is not used as a

transit switch to get to the Root Bridge. First, the Bridge Priority of the switch is raised to 49152, making it very unlikely that the switch will be elected to Root Bridge status. The Port Cost of all local switch ports is incremented by 3000, making the ports undesirable as Root Ports.

The command also includes a **rate** parameter. When an uplink on a switch goes down, UplinkFast makes it easy for the local switch to update its bridging table of MAC addresses to point to the new uplink. However, UplinkFast also provides a mechanism for the local switch to notify other upstream switches that stations downstream (or on toward the Access layer) can now be reached over the newly activated uplink. This action is done by sending multicast frames to all other switches containing the MAC addresses of the stations not learned on the uplink ports. These multicast frames are sent out at a rate specified by the **rate** parameter per 100 milliseconds. The default is 15 per 100 milliseconds. The final **all-protocols** parameter tells the switch whether to generate the multicast updates for all available protocols (IP, IPX, AppleTalk, and Layer 2 packets).

The IOS-based switch command uses a **max-update-rate** parameter to set the rate of multicast updates. The rate can be from 0 to 1000 multicasts per second.

To view the current UplinkFast parameters and ports, use the **show spantree uplinkfast** command.

BackboneFast: Redundant Backbone Paths

In the network backbone, or Core layer, a different method is used to shorten STP convergence. *BackboneFast* works by having a switch actively determine if alternate paths exist to the root bridge in the event that the switch detects an *indirect link failure*. Indirect link failures occur when a link not directly connected to a switch fails. A switch detects an indirect link failure when it receives inferior BPDUs from its Designated Bridge on either its root port or a blocked port. (Inferior BPDUs are sent from a Designated Bridge that has lost its connection to the Root Bridge, making it announce itself as the new Root).

Normally, a switch must wait for the Max Age timer to expire before responding to the inferior BPDUs. However, BackboneFast begins to determine if other alternate paths to the Root Bridge exist according to the type of port that received the inferior BPDU. If the inferior BPDU arrives on a port in the Blocking state, the switch considers the root port and all other blocked ports to be alternate paths to the root bridge. If the inferior BPDU arrives on the root port itself, the switch considers all blocked ports to be alternate paths to the root bridge. If the inferior BPDU arrives on the root port and there are no blocked ports, however, the switch assumes it has lost connectivity with the root bridge. In this case, the switch will assume that it has become the root bridge and BackboneFast will allow it to do so before the Max Age timer expires

Detecting alternate paths to the root bridge also involves an interactive process with other bridges. If the local switch has blocked ports, BackboneFast begins to use the *Root Link Query (RLQ)* protocol to see if there are upstream switches that have stable connections to the Root

Bridge. RLQ Requests are sent out. If a switch receives an RLQ Request and is either the Root Bridge or has lost connection to the Root, it sends an RLQ Reply. Otherwise, the RLQ Request is propagated on to other switches until an RLQ Reply can be generated. On the local switch, if an RLQ Reply is received on its current Root Port, the path to the Root Bridge is intact and stable. If it is received on a non-Root Port, an alternate Root Path must be chosen. The Max Age Timer is immediately expired so that a new Root Port can be found.

BackboneFast is simple to configure and operates by short-circuiting the Max Age Timer when needed. Although this function shortens the time a switch waits to detect a Root Path failure, ports still must go through full-length Forward Delay Timer intervals during the Listening and Learning states. Where PortFast and UplinkFast enabled immediate transitions, BackboneFast can only reduce the maximum convergence delay from 50 to 30 seconds.

To configure BackboneFast, use the following command:

```
Switch (enable) set spantree backbonefast {enable | disable}
```

When used, BackboneFast should be enabled on *all* switches in the network because BackboneFast requires the use of the RLQ request and reply mechanism to inform switches of Root Path stability. The RLQ protocol is only active when BackboneFast is enabled on a switch. By default, BackboneFast is disabled.

Foundation Summary

The Foundation Summary is a collection of tables and figures that provides a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final preparation before the exam, these tables and figures will hopefully be a convenient way to review the day before the exam.

Table 5-6 *Frame Distribution on a Two-Link EtherChannel*

Binary Addresses	Two-Link EtherChannel XOR and Link Number
Addr1: ... xxxxxxx0	
Addr2: ... xxxxxxx0	... xxxxxxx0: Link 0
Addr1: ... xxxxxxx0	
Addr2: ... xxxxxxx1	... xxxxxxx1: Link 1
Addr1: ... xxxxxxx1	
Addr2: ... xxxxxxx0	... xxxxxxx1: Link 1
Addr1: ... xxxxxxx1	
Addr2: ... xxxxxxx1	... xxxxxxx0: Link 0

Table 5-7 *EtherChannel Configuration Commands*

Task	CLI-Based Command	IOS-Based Command
Assign EtherChannel Distribution Method	set port channel all distribution {ip mac}[source destination both]	port group group-number [distribution {source destination}]
View EtherChannel Capabilities	show port capabilities [module/port]	
Configure EtherChannel	set port channel module/port-range mode {on off desirable auto}	port group group-number [distribution {source destination}]
Show EtherChannel	show port channel [mod[/port]]	show port group [group-number]

Table 5-8 *Configuration BPDU Message Content*

Field Description	Number of Bytes
Protocol ID (always 0)	2
Version (always 0)	1
Message Type (Configuration or Topology Change Notification BPDU)	1
Flags	1
Root Bridge ID	8
Root Path Cost	4
Sender Bridge ID	8
Port ID	2
Message Age (in 256ths of a second)	2
Maximum Age (in 256ths of a second)	2
Hello Time (in 256ths of a second)	2
Forward Delay (in 256ths of a second)	2

Table 5-9 *Topology Change Notification BPDU Message Content*

Field Description	Number of Bytes
Protocol ID (always 0)	2
Version (always 0)	1
Message Type (Configuration or Topology Change Notification BPDU)	1

Table 5-10 *Basic Spanning Tree Operation*

Task	Procedure
Elect Root Bridge	Lowest Bridge ID.
Select Root Port (one per switch)	Lowest Root Path Cost.
Select Designated Port (one per segment)	Lowest Root Path Cost.
Block ports with loops	Block ports that are non-Root and non-Designated Ports.

Table 5-11 *Spanning Tree Tie Breaker Criteria*

Sequence	Criteria
1	Lowest Root Bridge ID
2	Lowest Root Path Cost
3	Lowest Sender Bridge ID
4	Lowest Port ID

Table 5-12 *STP Path Cost*

Link Bandwidth	Old STP Cost	New STP Cost
4 Mbps	250	250
10 Mbps	100	100
16 Mbps	63	62
45 Mbps	22	39
100 Mbps	10	19
155 Mbps	6	14
622 Mbps	2	6
1 Gbps	1	4
10 Gbps	0	2

Table 5-13 *STP Port States*

State	Activity	Duration
Disabled	Administratively shut down; no STP activity.	N/A
Blocking	Receive BPDUs; no data transmitted or received.	Indefinite if loop has formed
Listening	Receive BPDUs; transmit BPDUs; can become Root or Designated Port; no data transmitted or received.	Forward Delay timer (15 seconds)
Learning	Receive BPDUs; transmit BPDUs; data received; MAC addresses learned; no data transmitted.	Forward Delay timer (15 seconds)
Forwarding	Receive and transmit BPDUs; receive and transmit data; MAC addresses learned.	Indefinite as long as port is up and loop not formed

Table 5-14 *STP Timers*

Timer	Function	Default Value
Hello	Interval between Configuration BPDUs.	2 seconds
Forward Delay	Time spent in Listening and Learning states before transitioning toward Forwarding state.	15 seconds
Max Age	Maximum length of time a BPDU can be stored without receiving an update; timer expiration signals an indirect failure with Designated or Root Bridge.	20 seconds

Table 5-15 *Types of STP*

Type of STP	Function
CST/MST	One instance of STP, over VLAN 1; 802.1Q-based.
PVST	One instance of STP per VLAN; Cisco ISL-based.
PVST+	Provides interoperability between CST and PVST; operates over both 802.1Q and ISL.

Table 5-16 *STP Configuration Commands*

Task	CLI-Based Command	IOS-Based Command
Enable STP	set spantree enable [all module/port]	spantree vlan-list
View STP	show spantree [vlan] show spantree module/port	show spantree [vlan] show spantree module/port
Set Bridge Priority	set spantree priority bridge-priority [vlan]	spanning-tree [vlan vlan-list] priority bridge-priority
Set Root Bridge (macro)	set spantree root [secondary] [vlan-list] [dia diameter] [hello hello_time]	
Set Port Cost	set spantree portcost module/port cost set spantree portvlancost module/port [cost cost] [vlan-list]	spanning-tree [vlan vlan-list] cost cost
Set Port Priority	set spantree portpri {module/port} priority set spantree portvlanpri {module/port} priority [vlans]	spanning-tree [vlan vlan-list] port-priority port-priority

continues

Table 5-16 STP Configuration Commands (Continued)

Set STP Timers	set spantree hello <i>interval</i> [<i>vlan</i>] set spantree fwddelay <i>delay</i> [<i>vlan</i>] set spantree maxage <i>agingtime</i> [<i>vlan</i>] set spantree root [<i>secondary</i>] [<i>vlan-list</i>] [<i>dia diameter</i>] [hello <i>hello-time</i>]	spanning-tree [<i>vlan vlan-list</i>] hello-time <i>seconds</i> spanning-tree [<i>vlan vlan-list</i>] forward-time <i>seconds</i> spanning-tree [<i>vlan vlan-list</i>] max-age <i>seconds</i>
Set PortFast	set spantree portfast { <i>module/port</i> } { enable disable }	spanning-tree portfast
Set UplinkFast	set spantree uplinkfast { enable disable } [<i>rate update-rate</i>] [all-protocols off on]	spanning-tree uplinkfast [max- update-rate <i>pkts-per-second</i>]
Set BackboneFast	set spantree backbonefast { enable disable }	

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however, they are designed to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to these questions can be found in Appendix A, on page 477.

- 1 What is EtherChannel? What types of switch links can it be used with?

- 2 How does an EtherChannel distribute broadcasts and multicasts?

- 3 How is traffic distributed over an EtherChannel?

- 4 What CLI-based switch command could be used to configure a 4-port EtherChannel on switch ports 3/1, 3/2, 3/3, and 3/4? The switch should use PAgP to actively negotiate the EtherChannel.

5 What is PAgP used for?

6 What happens if one port of an EtherChannel is unplugged or goes dead? What happens when that port is reconnected?

7 What is a bridging loop? Why is it bad?

8 Put the following STP port states in chronological order:

- a. Learning
- b. Forwarding
- c. Listening
- d. Blocking

9 Name two types of STP messages used to communicate between bridges.

10 What commands can be used to configure a CLI-based switch as the Root Bridge on VLAN 10, assuming that the other switches are using the default STP values?

- 11 Using your Root Bridge answer from question 10, what commands can be used to configure another CLI-based switch as a secondary or backup Root Bridge on VLAN 10?

- 12 What criteria are used to select the following:

- a. Root Bridge
 - b. Root Port
 - c. Designated Port
 - d. Redundant (or secondary) Root Bridges
- 13 Which of the following switches will become the Root Bridge, given the information in the table below? Which switch will become the secondary Root Bridge in the event that the Root Bridge fails?

Switch Name	Bridge Priority	MAC Address	Port Costs
Catalyst A	32768	00-d0-10-34-26-a0	All are 19
Catalyst B	32768	00-d0-10-34-24-a0	All are 4
Catalyst C	32767	00-d0-10-34-27-a0	All are 19
Catalyst D	32769	00-d0-10-34-24-a1	All are 19

- 14 What conditions cause an STP topology change? What effect does this have on STP and the network?

- 15 A Root Bridge has been elected in a switched network. Suppose a new switch is installed with a lower Bridge ID than the existing Root Bridge. What will happen?

16 What is the single-most important design decision to be made in a network running STP?

17 Where should the Root Bridge be located in a switched network?

18 Suppose a switch receives Configuration BPDUs on two of its ports. Both ports are assigned to the same VLAN. Each of the BPDUs announces Catalyst A as the Root Bridge. Can the switch use both of these ports as Root Ports? Why?

19 What happens to a port that is neither a Root Port nor a Designated Port?

20 Suppose you need to troubleshoot your Spanning-Tree topology and operation. What commands and information can you use on a switch to find information about the current STP topology?

21 How is the Root Path Cost calculated for a switch port?

- 22** What conditions can cause ports on the Root Bridge of a network to move into the Blocking state? (Assume that all switch connections are to other switches. No crossover cables are used to connect two ports together on the same switch.)

- 23** What is the maximum number of Root Ports that a Catalyst switch can have?

- 24** What mechanism is used to set STP timer values for all switches in a network?

- 25** What parameters can be tuned to influence the selection of a port as a Root or Designated Port?

- 26** What CLI-based command can be used to enable fast STP convergence for a single workstation on switch port 3/7?

- 27** What technology can be useful to decrease the amount of time STP keeps an end user's workstation in the Blocking state when it powers up?

- 28** What happens if the STP Hello Time is decreased to one second in an effort to speed up STP convergence? What happens if the Hello Time is increased to ten seconds?

- 29** Where should the UplinkFast feature be used in a switched network?


- 30** What CLI-based switch command can be used to safely adjust STP timers on the Root Bridge in VLAN 7? Assume that the network consists of Catalysts A, B, and C, all connected to each other in a triangle fashion.

Scenarios

Scenario 5-1: Spanning-Tree Protocol Operation

Given the network diagram shown in Figure 5-10, complete the exercises that follow:

Figure 5-10 Network Diagram for Scenario 5-1



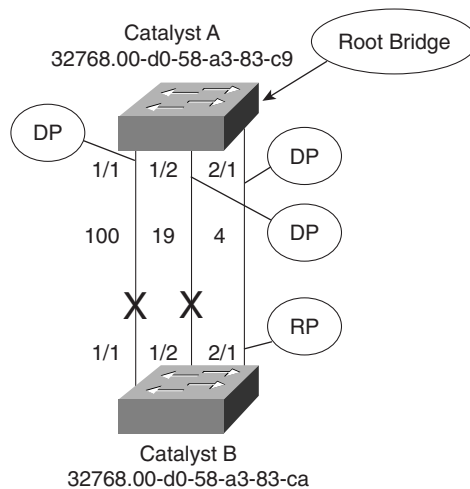
- 1 Manually compute the Spanning-Tree topology. Note which switch is the Root Bridge, which ports are Root Ports and Designated Ports, and which ports are in the Blocking state.
- 2 If the 100-Mbps link (ports 1/2) is disconnected, what will happen with the STP?
- 3 If the 1000-Mbps link (ports 2/1) is disconnected, how much time will elapse before the two switches can communicate again? (Assume both switches are using the default STP timer values and no additional features for faster convergence.)
- 4 Assume that for some reason the physical 1000-Mbps link (ports 2/1) stays up and active, but BPDUs are not allowed to pass (that is, an access list filter is blocking BPDUs). What will happen and when?

Scenario Answers

Scenario 5-1 Answers: Spanning-Tree Protocol Operation

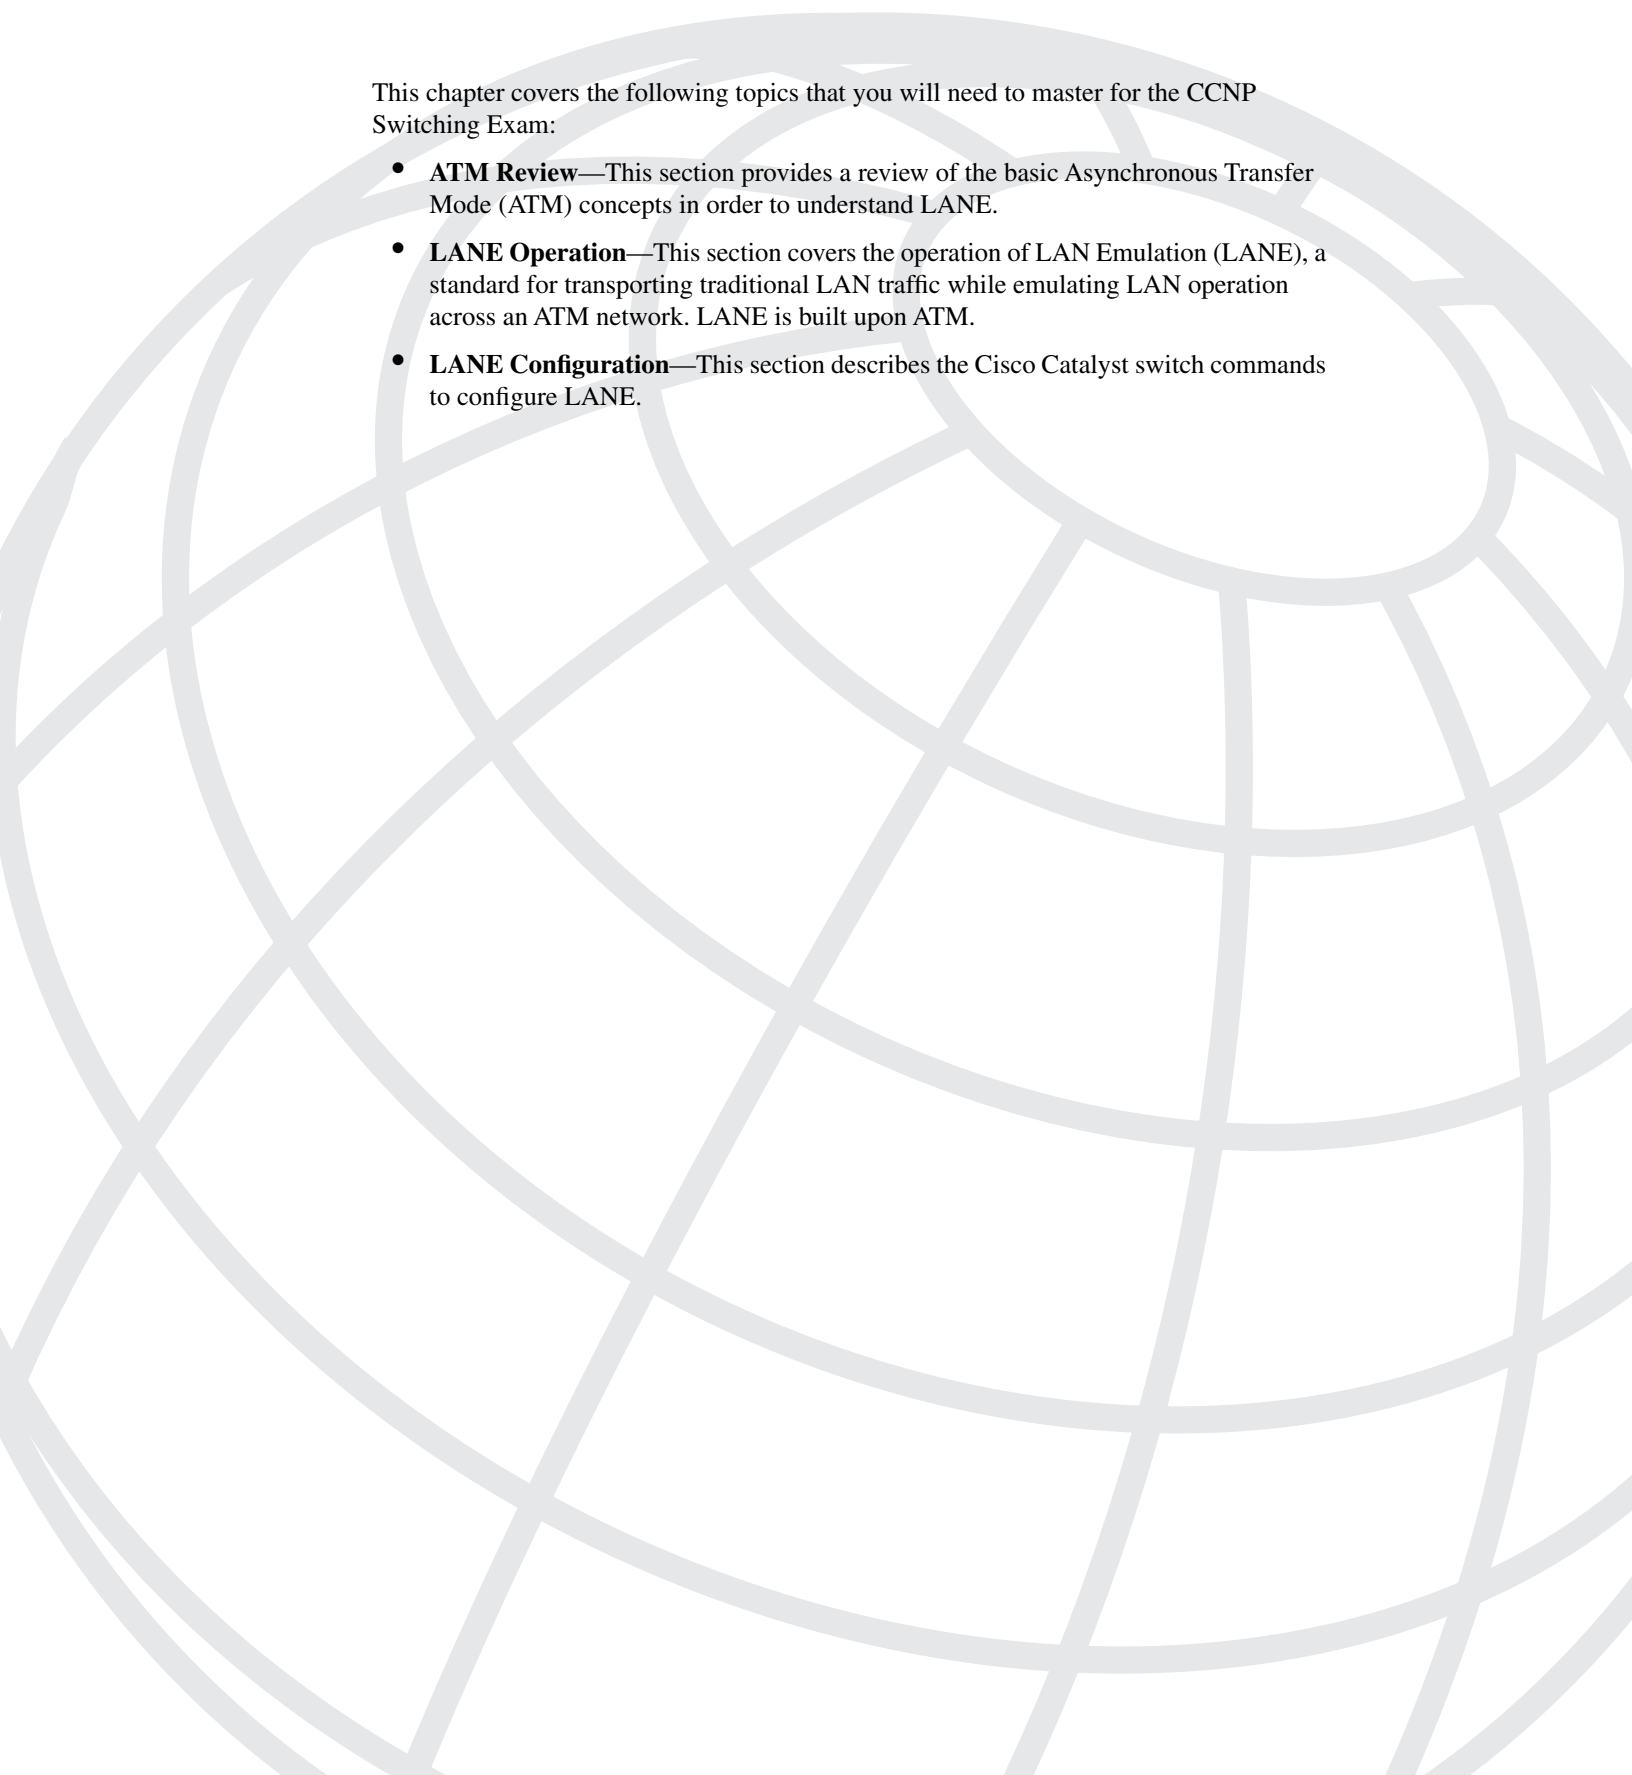

- 1 The Spanning-Tree topology should look like the diagram in Figure 5-11. Catalyst A is the Root Bridge and only the 1000-Mbps link is Forwarding. The Root Ports (RP) and Designated Ports (DP) are labeled on the diagram.

Figure 5-11 Resulting Spanning-Tree Topology for Scenario 5-1



- 2 Because the 100-Mbps link is in the Blocking state on Catalyst B, no major change in the topology will occur. Effectively, this link was already “disconnected.” However, after the physical link status goes down, both Catalyst A and Catalyst B will sense the change and begin sending TCN BPDUs to notify each other of the topology change. Because Catalyst A is the Root Bridge, it will acknowledge the TCN to Catalyst B. Both switches will age out their MAC address tables in Forward Delay seconds.
- 3 Disconnecting the 1000-Mbps link will cause Catalyst B to immediately find another Root Port. Ports 1/1 and 1/2 will go into the Listening state waiting to receive BPDUs. Port 1/2 with a cost of 19, will become the next Root Port, as soon as Catalyst B computes the Root Path Cost (0+19) for it. Port 1/2 will stay in the Listening state for Forward Delay (15 seconds) and then in the Learning state for Forward Delay (15 seconds). Port 1/2 will move into the Forwarding state, restoring connectivity in 30 seconds. (If PAgP is operating on the port, an additional delay of 20 seconds will be experienced.)

- 4 Because the link status of the 1000-Mbps link stays up, neither Catalyst will detect a link failure. Therefore, there will be no immediate attempt to find another Root Port. Instead, Catalyst B will not receive BPDUs from Catalyst A over link 2/1 because they are being filtered out. After the MaxAge Timer expires (20 seconds), Catalyst B will age out the stored BPDU for Catalyst A on port 2/1. Catalyst B will move ports 1/1 and 1/2 into the Listening state to determine a new Root Port. As in step 3, port 1/2 will become the Root Port with a lower Root Path Cost than port 1/1. The port will move through the Listening (15 seconds) and Learning (15 seconds) states and into the Forwarding state. The total time that has elapsed before connectivity is restored is $20+15+15=50$ seconds. (Again, if PAgP is active on the port, an additional 20 seconds can be added to the delay.)



This chapter covers the following topics that you will need to master for the CCNP Switching Exam:

- **ATM Review**—This section provides a review of the basic Asynchronous Transfer Mode (ATM) concepts in order to understand LANE.
- **LANE Operation**—This section covers the operation of LAN Emulation (LANE), a standard for transporting traditional LAN traffic while emulating LAN operation across an ATM network. LANE is built upon ATM.
- **LANE Configuration**—This section describes the Cisco Catalyst switch commands to configure LANE.

Trunking with ATM LANE

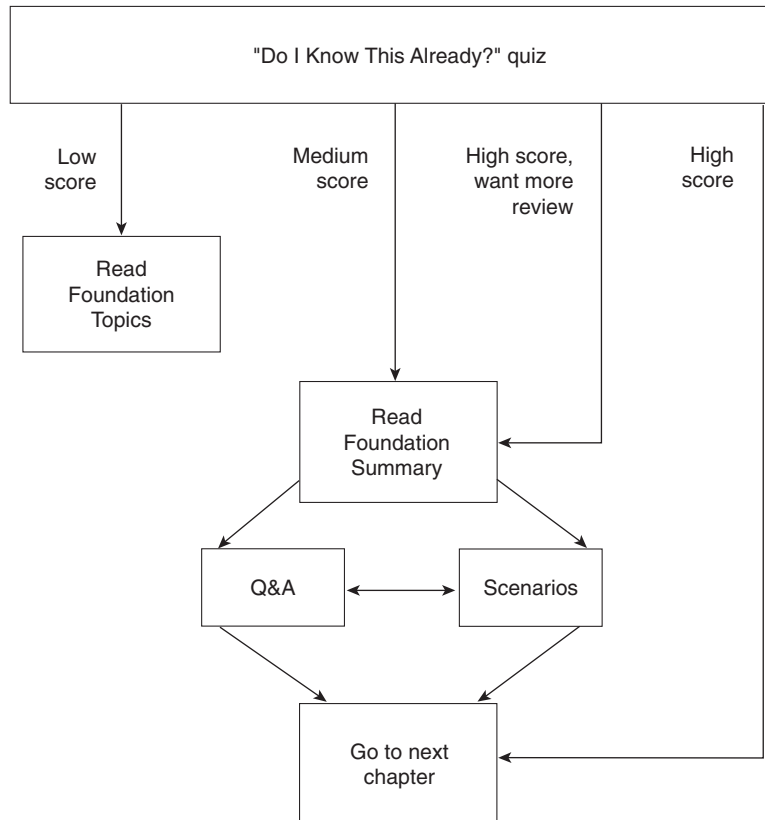
Chapter 4, “VLANs and Trunking,” provided a discussion of trunking with various forms of Ethernet links. Trunking, or carrying multiple VLANs over a single link, can also be accomplished using Asynchronous Transfer Mode (ATM). ATM by itself is a connection-oriented technology built upon relaying cells of data. Therefore, it cannot inherently trunk VLANs. However, the LAN Emulation (LANE) protocol uses ATM as a means to mimic traditional LAN media and can provide the trunking function. Multiprotocol over ATM (MPOA) is another protocol that extends LANE to offer more efficient path selection through an ATM network. (MPOA is not covered in this book.)

This chapter presents a review of ATM and focuses on the use of LANE technology for trunking. While ATM is a very complex technology, it is presented only briefly to set the foundation for a more detailed discussion of LANE.

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and the answers for all your work with this book in one place, for easy reference.
- Take the “Do I Know This Already?” quiz and write down facts and concepts (even if you never look at the information again).
- Use the diagram in Figure 6-1 to guide you to the next step.

Figure 6-1 *How to Use This Chapter*

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The quiz helps you make good choices of how to spend your limited study time. The quiz is sectioned into five smaller “quizlets,” which correspond to the five major headings in the Foundation Topics section of the chapter. Although your answer may differ somewhat from the answers given, finding out if you have the basic understanding of what is presented in this chapter is more important. You will find that these questions are open-ended, rather than multiple choice as found on the exams. You will be able to focus more on understanding the subject matter than on memorizing details.

Figure 6-1 outlines suggestions on how to spend your time in this chapter. Use the scoresheet in Table 6-1 to record your score.

Table 6-1 *Scoresheet for Quiz and Quizlets*

Quizlet Number	Foundation Topics Sections Covering These Questions	Questions	Score
1	ATM Review	1–4	
2	LANE Operation	5–9	
3	LANE Configuration	10–13	
All questions		1–13	

1 What is the basic unit of ATM data? What is its basic format (header, payload, and so forth)?

2 What is an ATM edge device? What Cisco devices can be used?

3 What type of addressing is used to identify ATM devices?

4 What information is carried within each ATM data unit to specify how to get from the source to the destination?

5 What is LANE used for?

6 What are the functional components of LANE?

7 What is the difference between a VLAN and an ELAN?

8 When is an LE_ARP used?

9 For each of the LANE components, how many are necessary for LANE operation?

10 Which LANE component provides connectivity between a VLAN and an ELAN on a switch?

11 What NSAP addresses must be configured into the LECS database?

12 What is SSRP? Which LANE components can be configured for SSRP?

- 13 What Catalyst switch command can be used to view the current status of each LANE component (LECS, LES, BUS, and LEC)?

The answers to the quiz are found in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections,” on page 477. The suggested choices for your next step are as follows:

- **6 or less overall score**—Read the entire chapter. This reading includes the “Foundation Topics” and “Foundation Summary” sections, the Q&A section, and the scenarios at the end of the chapter.
- **7–9 overall score**—Begin with the “Foundation Summary” section and then go to the Q&A section and the scenarios at the end of the chapter.
- **10 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” section. Then, go to the Q&A section and the scenarios at the end of the chapter. Otherwise, move on to the next chapter.

Foundation Topics

ATM Review

This section presents a brief review of ATM concepts and operation. Although it is not necessary to know the intimate details of ATM, how to configure an ATM switch, or how the underlying ATM protocols work, you should understand ATM well enough to effectively design and configure LANE.

ATM is designed to provide multiple service types (voice, video, and data) over single pipelines very efficiently. All traffic is transported as small fixed-size cells. As well, traffic is not moved based on cell-by-cell decisions (as in packet-based routing) but upon connections built between end points.


Networks built on ATM typically have ATM switches (such as the Cisco Lightstream 1010) at the core or within the “cloud.” ATM switches do nothing but build connections between each other and relay native ATM cells across the connections. At the edge of the ATM network, many types of traditional networking devices can be connected. For example, LAN switches, routers, workstations, and video CODECs can all be ATM *edge devices* or *endpoints*. These devices provide the conversion between other network media formats and ATM cells and interface with the ATM switches using ATM protocols.

ATM defines these two types of interfaces as the *User-Network Interface (UNI)* (or the connection between ATM endpoints and ATM switches) and the *Network-to-Network Interface (NNI)* (or the connection between two ATM switches).

Figure 6-2 shows the structure and components of an ATM network. The ATM cloud consists of ATM switches, with NNI connections to each other. At the edge of the network, many sources of data can be connected to an ATM switch through UNI connections. Notice that each source starts with a different type of data element, but the ATM switch uses only cells.

NOTE

For further information about the many standards, protocols, and acronyms related to ATM, refer to the ATM Forum at www.atmforum.com. The Cisco Press title, *Cisco ATM Solutions*, also provides a wealth of information on ATM.

Figure 6-2 *An ATM Network*

Cells and SAR

All types of traffic are transported over ATM as small cells. Using cells of an optimal fixed size allows the following benefits:

- **Low latency, high throughput**—Small cells can be moved very quickly from switch to switch with a low propagation delay for the short serialized data from each cell. Fixed-size cells then can be relayed at a predictable rate. ATM switches also use hardware-based switching and reduced addressing and decision spaces to speed cell relay.
- **Multiservice traffic**—Traffic from many sources can be converted into the same fixed-length cells going into and out of the ATM network. ATM cells can also be relayed at a predictable rate, providing Quality of Service (QoS) or guaranteed timely delivery of real-time data streams for services like voice and video.

ATM cells have been standardized to a 53-byte length: a 5-byte header and a 48-byte payload. On this basis, using some method of fragmenting larger pieces of data (for example, an IP packet) into 53-byte cells should be obvious. Part of the ATM cell generation process is concerned with *segmentation and reassembly (SAR)*. On one ATM edge device, larger packets of higher-layer protocols are segmented or fragmented into 48-byte units. Each payload unit is attached to a 5-byte header to make a cell and then relayed to an ATM switch.

On the far end of a connection at another ATM edge device, cells are received from an ATM switch. ATM headers are removed, and the 48-byte payloads are reassembled into the original larger packets. Using this process, any type of data can be segmented, transported, and reassembled using the same 53-byte ATM cells.

ATM Model

The ATM standard uses a reference model to describe the hierarchy of its various operations, similar to the 7-layer OSI model. Table 6-2 shows a comparison between the layers of the OSI and ATM reference models. Notice that ATM operations are only performed in the data link and physical layers. The higher layers are still concerned with passing upper-layer protocols downward to the data link layer to be processed into ATM cells.

Table 6-2 *ATM Reference Model*

OSI Reference Model	ATM Reference Model
Layer 7 (application)	Higher Layers
Layer 6 (presentation)	
Layer 5 (session)	
Layer 4 (transport)	
Layer 3 (network)	
Layer 2 (data link)	ATM Adaptation Layers (AAL)
	ATM Layer
Layer 1 (physical)	ATM Physical Layer

The process of converting higher-layer data blocks into ATM cells is performed in the ATM Adaptation Layer (AAL). This layer contains several different processes, each tailored for the segmentation and re-assembly of a different higher-layer data type. The standardized AAL processes are described as follows:

- **AAL1**—Supports connection-oriented services that can emulate leased line circuits. AAL1 requires timing synchronization between source and destination, allowing the transport of real-time streams like voice and video. Data is sampled and put into cell payloads, along with sequence information before transmission.
- **AAL2**—Supports connection-oriented services for variable bit rate (VBR) applications.
- **AAL3/4**—Supports both connection-oriented and connectionless services. AAL3/4 was designed for use with Switched Multimegabit Data Service (SMDS). Data is segmented into cell payloads, while cyclic redundancy check (CRC) error information is added.

- **AAL5**—Supports both connection-oriented and connectionless services. Non-SMDS data, like TCP/IP over ATM and LANE, is transported in sequence with simple segmentation. A CRC value is added to the end of the pre-segmented frame.

The ATM Layer is used to establish connections to other ATM devices. Then cells are sent across the ATM network using the header contents of each ATM cell.

Although ATM is designed to be media independent, it does have a physical layer that converts cells into bitstreams for the appropriate media. As well, all timing information and error control are handled in this layer. Common types of physical media include DS-3/E3, 155 Mbps over multimode and single-mode fiber (OC3), 155 Mbps over shielded twisted-pair cabling, 622 Mbps over fiber (OC12), and Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH).

Virtual Circuits

ATM relies on connections to be built across the ATM network in order to relay cells end-to-end. Before cells are sent, connections must be negotiated and set up. *Permanent virtual circuits (PVCs)* are manually built to support a predetermined path through an ATM cloud. As the name implies, these permanent, static connections remain in place until you manually remove them.

Switched virtual circuits (SVCs) are just the opposite—they are dynamically built and torn down by ATM switches as they are needed. When one ATM edge device requires a new circuit, it informs an ATM switch that an SVC needs to be built to the destination.

Virtual circuits (VCs) can also be built in two ways: *Point-to-point*, where one device talks to one device; and *Point-to-multipoint*, where one device communicates in one direction to many end devices.

ATM uses specific terminology to identify the hierarchical arrangement of its VCs. A *virtual channel Connection (VCC)* is the connection that is set up across the ATM network between endpoints. A *virtual path (VP)* is a bundle of virtual channels that share a common path through the switched network. When a VP is used, you do not need to switch the individual VCs contained in it—just switch the overall VP bundle. Lastly, a *transmission path* is a bundle of virtual paths.

ATM Addressing

Two types of addresses exist within ATM: virtual path and virtual channel identifiers (VPI/VCI), and network service access point (NSAP) addresses. Each type of address is used for a specific purpose and is described in the sections that follow.

VPI/VCI Addresses

Recall that an ATM cell contains a 5-byte header and a 48-byte payload. A 5-byte space is not sufficient to contain a large address space for the source or destination endpoints. In packet- or frame-based networks, placing the source and destination addresses in each frame is necessary so that routers and switches can forward frames correctly. When ATM transports a cell from source to destination, however, only the specific VC needs to be known.

ATM uses the *VPI/VCI* identifier combination to deliver cells through a network. An ATM edge device places the VPI/VCI values in a cell before presenting the cell to an ATM switch. The switch then references a forwarding table, relating the VPI and VCI values to outbound switch ports.

For comparison purposes, a UNI cell contains an 8-bit VPI and a 16-bit VCI address. An NNI cell expands the VPI to 12 bits and has a 16-bit VCI value. These values are only locally significant to an ATM switch and do not have to be globally unique.

NSAP Addresses

To identify individual ATM devices and endpoints, 20-byte *network service access point (NSAP)* addresses are used. Each device is required to have a unique NSAP address. Typically these addresses are written out as groups of 4 hex digits, separated by dots (for example, 47.0091.8100.0000.00c0.1004.d94e.0000.3750.a011.00). The rightmost two hex digits are usually grouped by themselves, with the first dot separator to the left. (This grouping will be important when using LANE.) NSAP ATM addresses are composed of the following:

- **Prefix**—A 13-byte field that uniquely identifies every ATM switch in the network. Cisco ATM switches have a predefined 7-byte value of 47.0091.8100.0000, followed by a 6-byte unique MAC address preconfigured on each switch.
- **End-System Identifier (ESI)**—A 6-byte field that uniquely identifies every device attached to an ATM switch. A 6-byte MAC address is usually used in this field.
- **Selector**—A 1-byte field that identifies a process running on an ATM device. Cisco devices usually use the selector value to identify an ATM subinterface number.

NOTE

NSAP addresses are only used to build SVCs within an ATM network. For PVCs, only VPI/VCI values are needed. Once an SVC or PVC has been built, the VPI/VCI pair is used to relay cells from switch to switch.

Obviously, the size of the NSAP address prohibits it from being carried within ATM cells. ATM devices only communicate NSAP addresses within signaling protocols, when SVCs are being built or torn down.

Inherent ATM Protocols

ATM uses two protocols to automate several functions, simplifying network configuration and operation. *Integrated Local Management Interface (ILMI)* is a protocol that provides an automatic way for an ATM device to learn about its neighbors. For example, an ATM edge device can learn the NSAP prefix address from an ATM switch, rather than requiring an administrator to manually configure the address. As well, an ATM switch can learn the ESI portion of the NSAP address of an ATM edge device through ILMI. This concept will be important when applied to LANE.

Private Network-to-Network Interface (PNNI) is an ATM protocol that is used to build and dynamically tear down SVCs. PNNI is only used between ATM switches because it involves SVC administration. PNNI is sometimes termed a Layer 2 routing protocol because of its use in determining paths through an ATM network. In addition, PNNI enables ATM switches to load balance traffic across multiple paths and parallel links, and provides redundancy in case of path failure.

LAN Emulation (LANE)

This section discusses how ATM can be used to provide LAN connectivity and to trunk multiple VLANs across an ATM network. Because ATM is so dissimilar to traditional LAN technologies, it must be made to emulate a LAN (hence the name *LAN Emulation*) through the use of several processes and components. LANE is an ATM Forum standard that specifically provides mechanisms to emulate IEEE 802.3 Ethernet and IEEE 802.5 Token Ring LANs. These mechanisms are described in more detail in the following sections.

Note that LANE is used only where LAN functionality is needed at the edges of an ATM network. Although LANE makes the ATM cloud appear to be a LAN, only native ATM is used on the ATM switches. LANE is built as a layer on top of ATM so that LANE operation is transparent to a standard ATM network.

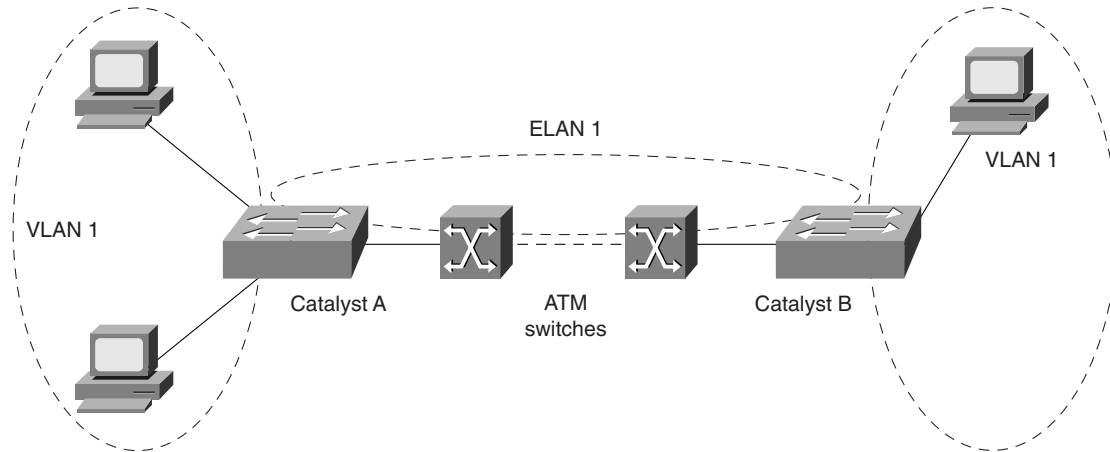
LANE Components

To understand how LANE works within a larger switched network, you should first understand how *emulated* LANs (ELANs) differ from *virtual* LANs (VLANs). Recall that virtual LANs are used on Catalyst switches as a means to segment traffic into logical networks. VLANs can also be trunked between switches. Similarly, ELANs are used by ATM devices to segment traffic into logical networks. However, *ELANs exist only within the ATM domain*. ELANs and VLANs remain separate except where they are physically bridged in Catalyst switches that support both.

This difference is illustrated in Figure 6-3 as a network consisting of two Catalyst switches connected by two ATM switches. Each Catalyst switch has ports assigned to VLAN 1. Note that each VLAN 1 is significant only to the local switch, and that some other means must be used to

logically connect the two VLANs. An ELAN exists on each Catalyst switch as ELAN 1. LANE makes logically connecting the ELAN between switches across an ATM network possible.

Figure 6-3 An Example of VLAN and ELAN Functionality



Because LANE must emulate the functions of a LAN, consider the following characteristics of a LAN:

- **Layer 2 MAC addresses**—All stations on a LAN use unique 6-byte MAC addresses.
- **Physical or logical connectivity**—All stations on a LAN are physically connected through hubs or switches, and logically connected through VLANs and trunking on switches.
- **Broadcasts and multicasts**—LANs allow one-to-many transmission of frames in the form of broadcasts and multicasts.

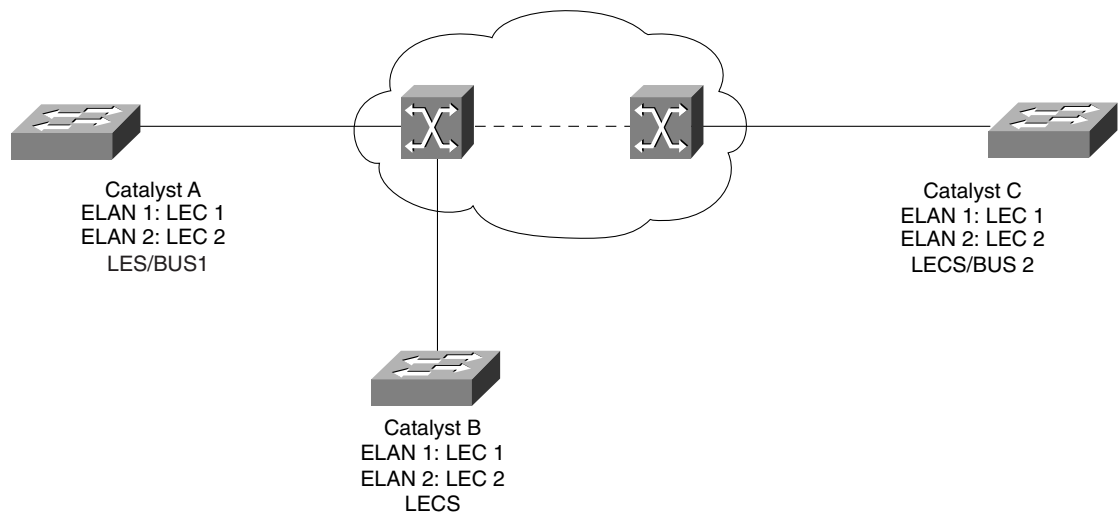
LANE has been designed with discrete components to address these and other LAN requirements. For *each ELAN* that exists in a network, the following LANE components are used:

- **LAN Emulation Client (LEC)**—The LEC provides the basic ELAN function on each ATM device where an ELAN connection is needed. The LEC emulates an interface to a traditional LAN and provides data forwarding, address resolution, and MAC address registration with other LANE components. Each LEC can communicate with other LECs in the network over ATM virtual channel connections (VCCs) emulating physical or logical LAN connectivity.

- **LAN Emulation Server (LES)**—The LES (one per ELAN) acts as a central control point for all LECs in an ELAN. Each LEC must register its MAC addresses with the LES so that the LES can provide MAC address to NSAP address translation. The LES also maintains connections to each LEC to provide control information to all ELAN members.
- **Broadcast and Unknown Server (BUS)**—Because ATM does not support broadcast functions, the BUS is used to emulate LAN broadcasts. Any LEC that needs to send a broadcast sends it directly to the BUS. The BUS is then able to forward the broadcast to all other LECs (and all stations in each LEC) through a point-to-multipoint VCC. The BUS takes care of queuing and sequencing broadcast and multicast frames over the point-to-multipoint VCC so that the frames arrive in the correct order at the LECs.
- **LAN Emulation Configuration Server (LECS)**—The LECS (only one per administrative domain) acts as the central administrative control point for all ELANs in a domain. The LECS maintains a database of ELANs and the ATM addresses of the LESs that control each ELAN. Each LEC must query the LECS to ask for membership in an ELAN and to request the NSAP address of the respective LES.

Figure 6-4 shows a typical LANE network. Three Catalyst switches are ATM edge devices; each is connected to the ATM network. LANE is configured to support two ELANs associated with two VLANs present on each switch. Therefore, each Catalyst is configured with two LEC components—one for each ELAN/VLAN. The remaining LANE components can be configured arbitrarily. In this example, the LECS is configured on Catalyst B. For ELAN1, a LES/BUS pair is configured on Catalyst A, and a LES/BUS pair for ELAN2 is configured on Catalyst C. Notice that both LES/BUS pairs could have been configured on a single Catalyst along with the LECS.

Figure 6-4 An Example LANE Network with Component Placement



LANE Operation

With LAN technology, it is usually not necessary to knowing exactly how the network operates is not usually important. It is enough to know that when a workstation is plugged into hub or switch it is able to talk to other workstations on that LAN. LANE also appears to be plug-and-play to end stations because it is emulating a LAN. You should have a good grasp on how LANE accomplishes this because you will be configuring the LANE components and making sure they communicate with one another.

This section describes the inner workings of LANE from the viewpoint of a newly created LEC. Before end user stations can use the emulated LAN, the local LEC must first become a member of the ELAN. The steps to LANE membership are described in detail in the sections that follow.

Step 1: Contacting the LECS

A newly created LEC must first contact the LECS so that the LEC can be pointed to the LES of its specific ELAN. Remember that the LECS is the keeper of the LES/ELAN database for all ELANs in the network. A separate LES is available for each ELAN.

To begin, an LEC needs to find the ATM address of the LECS. This address can be obtained in one of the following ways:

- **ILMI**—The NSAP address of the LECS must be manually configured into every ATM switch in the network. The LEC then queries a local switch through ILMI for the LECS address. This method is preferred due to its simplicity and automatic nature.
- **Manual Configuration**—The LECS NSAP address must be manually configured into every LEC.
- **Well-known VPI/VCI**—LECs can contact the LECS over the well-known VPI/VCI value of 0/17. This action assumes a PVC has already been built using 0/17.
- **Well-known NSAP**—Using the well-known NSAP of 47.0079000000000000000000000000.00A03E000001.00, a LEC can find the nearest LECS.

Once the LECS address has been found, the LEC can contact the LECS address directly with a configuration request over a *Configuration Direct VC*. The LECS looks up the desired ELAN in its database. Once the ELAN is found, the LECS responds with the ATM address of the respective LES, the type of LAN being emulated, the ELAN name, and the MTU of the ELAN.

Step 2: Contacting the LES

Now the LEC is ready to check in to become a member of the ELAN. The LEC opens a direct SVC (a *Control Direct VC*) with the LES, and registers its own ATM address and MAC address with the LES database. As well, the LEC can optionally register any other MAC addresses

directly connected to the LEC. For these devices, the LEC will become a proxy. Remember that the LES will provide MAC-NSAP address translation for the ELAN.

The LES also keeps control over members of the ELAN. A point-to-multipoint connection (the *Control Distribute VC*) is maintained from the LES to all known LECs. Any ELAN control information is passed over this one-to-many VCC. Therefore, the LES adds the LEC as a leaf node on the control connection.

Step 3: Contacting the BUS

The BUS is the keystone to emulating a LAN, by providing the means to send a broadcast (or multicast) to all stations in the ELAN. Instead of requiring each LEC to maintain a connection to every other LEC for broadcast traffic to be flooded over, each LEC needs only to point broadcasts to the BUS. The BUS in turn maintains a point-to-multipoint connection to every LEC in the ELAN for broadcast purposes.

To find the BUS, the LEC must learn its ATM address by sending the LES an ARP request (or an *LE_ARP_REQUEST* in LANE terms). The LEC will attempt to query the BUS ATM address by using the broadcast MAC address of 0xFFFFFFFF. How then does the LES know the ATM address of the BUS? Cisco requires that the LES and BUS be located on the same device. (Their NSAP addresses will differ slightly because the LES and BUS are run as separate processes on the device.)

Now the LEC can contact the BUS directly by building a VCC for sending broadcast and multicast traffic. The BUS replies to the LEC and adds the LEC as a leaf node on its broadcast point-to-multipoint connection.

NOTE Once a LEC sends a broadcast to the BUS, on behalf of a connected device on the LEC, the BUS distributes the broadcast to all LECs in the ELAN over its multipoint connection. Each LEC is responsible for flooding the received broadcast to its own LAN- or VLAN-connected stations.

Step 4: Communicating Between LECs

Once the previous three steps have been completed, the LEC has joined the ELAN and has formed the required ATM connections to the LANE components. The LEC can now communicate with other LECs in the ELAN. Building Data Direct VCs, or direct connections, from one LEC to another completes this communication. In certain cases, one LEC can build a Data Direct VC to another LEC, while the other LEC is building a Data Direct VC back to the first LEC. Two VCs are not needed so only the VC that was requested by the lowest NSAP address is used.

Address Resolution

Two types of address resolution can occur within ATM LANE: IP ARP and LE_ARP (or LAN Emulation ARP). IP ARP is associated with resolving between MAC addresses and IP addresses. LE_ARP occurs when a LANE component needs to resolve between an NSAP address and a MAC address. These two processes are summarized in the sections that follow, as they apply to LANE operation.

Address Resolution Scenario 1: Using IP ARP to Resolve MAC Addresses

Workstation A needs to contact Workstation B, but only knows its IP address. IP ARP is used to find Workstation B's MAC address:

- Step 1** A workstation generates an ARP request broadcast on its local LAN (switch port) to find a MAC address.
- Step 2** The switch floods the broadcast out all VLAN ports, as well as to the ELAN associated with the VLAN. This flooding occurs on the switch's ATM module.
- Step 3** The LEC contacts the BUS with a broadcast frame to be delivered.
- Step 4** The BUS sends the broadcast to all LECs in the ELAN over its multipoint connection.
- Step 5** Each LEC in the ELAN receives the broadcast from the BUS and floods the broadcast out all local VLAN ports.
- Step 6** The destination station receives the ARP request and sends an ARP reply frame. Because the reply is not a broadcast, it is returned to the source via a Data Direct VC between destination and source LECs.

Address Resolution Scenario 2: Using LE_ARP to Resolve NSAP Addresses

Workstation A needs to contact Workstation B, and Workstation A already knows both its IP and MAC addresses. Once Workstation A's LEC receives a unicast frame, it must find Workstation B's LEC via its NSAP address so that a Data Direct VC can be built between the LECs. LE_ARP is used to resolve the NSAP address:

- Step 1** Workstation A sends a frame to Workstation B's MAC address. Workstation A's switch has an entry for the MAC address on its ATM module, pointing toward an ELAN.
- Step 2** Workstation A's LEC does not know the NSAP address for Workstation B's LEC. Therefore, Workstation A's LEC sends an LE_ARP request to the LES.

- Step 3** The LES looks up Workstation B's MAC address in its MAC/NSAP address table. If the LES finds the NSAP entry, it replies to Workstation A's LEC with the address.
- Step 4** If the NSAP entry is not in the table, the LES forwards the LE_ARP request on to all ELAN LECs over its multipoint control connection.
- Step 5** The LEC where Workstation B is attached has the MAC address in its bridging table. Therefore, Workstation B's LEC sends an LE_ARP reply back to the LES with its NSAP address.
- Step 6** Workstation A's LEC can now build a Data Direct VC to Workstation B's LEC for data transfer.

Design of LANE Components

Several factors are to be considered when implementing LANE. Because the various LANE components can be placed anywhere in the network, some thought should be put into each component's location. As well, some options are available to provide redundancy in the LANE components. The following sections discuss these topics.

LANE Component Placement

LANE is built upon several independent components: the LECS, the LES, the BUS, and one or more LECs. These components can be configured on any device in the network that supports LANE. For instance, any of the LANE components can be provided by Catalyst switch LANE modules, Cisco routers with ATM Interface Processors (AIPs), and Cisco LightStream 1010 ATM switches.

Some simple guidelines will help you choose which LANE components to place on which Cisco devices:

- **LECS**—The LECS (because it is only consulted when a LEC is initializing and looking for the appropriate LES) is not very CPU-intensive. Any available Cisco LANE device can be used for the LECS. However, "available" is the key word. The LECS is the central LANE database for all ELANs. Therefore, the LECS should be placed on a device that is highly available to all LECs in the network. This LEC could be a centrally located ATM switch or a distribution layer Catalyst switch LANE module.
- **LES/BUS**—The LES is used to provide MAC/NSAP address resolution. Most any LANE device can provide this table lookup function without much burden. However, Cisco requires that the LES and BUS be configured as a single unit on the same device. The BUS is saddled with accepting and sending all broadcast traffic for the ELAN—a computational burden that increases as the number of LECs and workstations grow. The

BUS should always be configured on the most robust Catalyst switch in the network, so that its function doesn't hamper other switching duties of the switch.

- **LEC**—Obviously, a LEC must be configured on every LANE device where ELAN connectivity is needed. Each ATM edge device will need a LEC. Although an ATM switch does not require a LEC because it switches native ATM traffic, it will need its own LEC to process any management traffic (for example an IP address assignment, syslog, or SNMP support).

LANE Component Redundancy (SSRP)

Although LANE can be implemented as separate components dispersed throughout a network, each component is still a single point of failure. The ATM Forum's LANE 1.0 standard only allows a single LECS for a network, and a single LES and BUS for each ELAN. However, Cisco has implemented a redundancy protocol for LANE that allows multiple LECS, LES, and BUS components. *Simple Server Redundancy Protocol (SSRP)* provides communication between the primary active component and one or more standby components so that the standby can take over if the primary fails. Under this protocol, only one active LANE component is allowed at any time. The *Fast Simple Server Redundancy Protocol (FSSRP)* allows multiple LES/BUS pairs to be defined and active at any time. Recovery from a failed LES causes no noticeable delay.

Redundant LECS components are provided by configuring the list of multiple LECS NSAP addresses on all ATM switches. The switches then provide the address of the first active LECS to LANE devices when the switches request the LECS address via ILMI.

Likewise, a list of redundant LES/BUS components can be configured in the LECS database for an ELAN. Under normal SSRP, the LECS will provide the next available address in the list when a LEC requests it. FSSRP allows up to four LES/BUS pairs to be active at one time. LECs that are FSSRP-aware build VCs to every LES/BUS pair automatically.

LANE Configuration

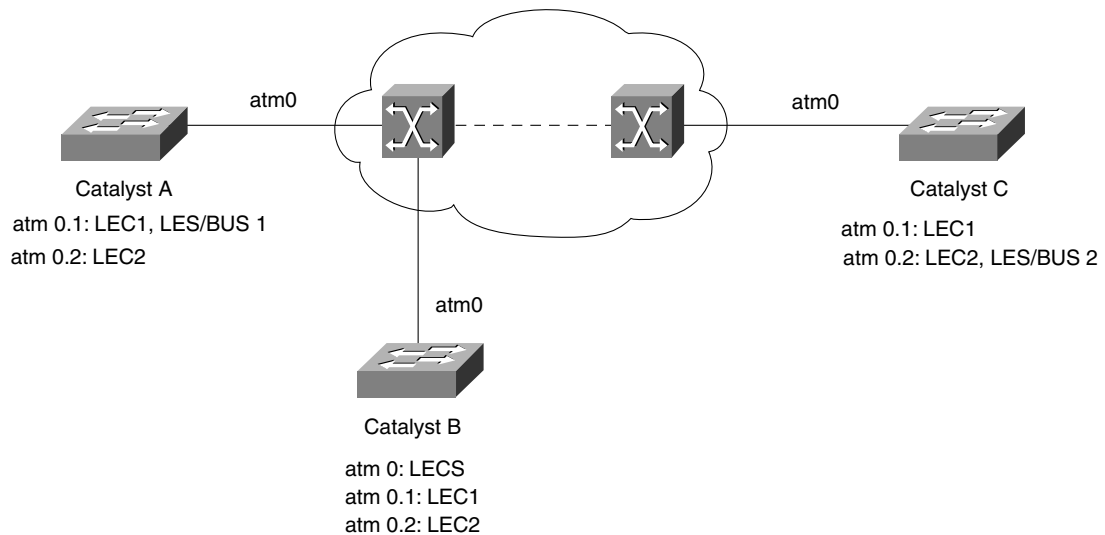
This section discusses the procedures for configuring the various LANE components on Cisco Catalyst switches. The order that the components are configured is important because each component is dependent upon another.

On Cisco ATM devices, ELANs are configured on ATM subinterfaces. This configuration makes it possible to support many ELANs over a single ATM link. As well, the LANE components necessary for a specific ELAN must be configured on the respective subinterface for that ELAN (ATM 0.1, ATM 0.2, and so on). The LECS, because it exists for *all* ELANs, must be configured on the major ATM interface (ATM 0).

The example from Figure 6-4 is shown again in Figure 6-5, along with ATM interface and subinterface numbers. Notice that each LEC must be configured on a different subinterface. The

LES/BUS pairs must be configured on the subinterfaces where their respective ELANs are present. The LECS must be configured on an ATM major interface because it keeps a database for all ELANs.

Figure 6-5 An Example LANE Network Showing ATM Interfaces



NSAP addresses on Cisco devices can be either manually configured or automatically generated. Automatic generation is most often used, because it offers easy configuration and use. Recall that an NSAP address is made up of a 13-byte prefix, a 6-byte ESI, and a 1-byte selector. The ATM switch provides the prefix value. The ESI value comes from the MAC address assigned to the particular LANE module and ATM interface on the Catalyst switch. The selector byte comes from the ATM edge device itself and is equal to the ATM subinterface number.

The LANE components are also given automatic NSAP addresses, according to the scheme shown in Table 6-3. You should thoroughly understand how these addresses are generated. You might find it useful in some situations to be able to work backwards and locate a LANE component given its NSAP address.

Table 6-3 Automatic NSAP Address Generation for LANE components

LANE Component	Prefix	ESI	Selector
LEC	From ATM switch	MAC address	ATM subinterface
LES	From ATM switch	MAC address + 1	ATM subinterface
BUS	From ATM switch	MAC address + 2	ATM subinterface
LECS	From ATM switch	MAC address + 3	.00

Even before configuration, seeing a listing of the automatically generated NSAP addresses on any LANE-capable switch module is possible. For example, because configuring the complete LECS NSAP address into the ATM switches is necessary, you can find that address by using the **show lane default** command on the LANE module where the LECS will be configured. Refer to the **show lane default** command output in Example 6-1 and notice how the NSAP addresses for the LANE components follow the scheme in Table 6-3. (All LANE component addresses are shown with this command, even if none of the components are configured on the switch. You can now see what the automatically generated NSAP addresses would be if the components were configured.)

Example 6-1 **show lane default** *Output Displays NSAP Addresses for LANE Components*

```
CatalystA_ATM#show lane default
interface ATM0:
LANE Client:          47.00001606288000300000000F.0050A28D5880.**
LANE Server:         47.00001606288000300000000F.0050A28D5881.**
LANE Bus:            47.00001606288000300000000F.0050A28D5882.**
LANE Config Server: 47.00001606288000300000000F.0050A28D5883.00
note: ** is the subinterface number byte in hex

CatalystA_ATM#
```

Notice that all LANE components are listed with their automatically generated NSAP addresses. To provide human readability, the addresses are output with dots separating the NSAP components. A dot appears after the leftmost byte of the prefix, at the end of the prefix, and at the end of the ESI (MAC address). Although an address is a long string of hex digits, it is easily broken down into prefix, ESI, and selector portions.

NOTE The LANE module on a Cisco Catalyst switch is not accessible from the normal switch command-line interface (CLI)-based user interface. Instead, the LANE module provides an IOS-based interface that must be accessed through the **session** switch command. Think of this command as a Telnet session that moves you from the switch CLI to another device's IOS interface. The LANE module and the switch supervisor share a common backplane but have independent user interfaces.

NOTE The **session** command requires an argument that specifies the module number where the LANE module is located in the chassis. For example, if a LANE module occupies slot 2 in a Catalyst 5500 switch, **session 2** would open a session to the LANE module's IOS. From that point on, IOS commands are used to display, debug, and configure the LANE module. Once configuration changes have been made, you must copy the running configuration to NVRAM (**copy run start**).

Configuring the LES and BUS

The LES and BUS for an ELAN must be located on the same device and must use the same ATM subinterface. To configure both LES and BUS components for an ELAN, use the following commands:

```
ATM(Config)# interface atm number.subint multipoint
ATM(Config-subif)# lane server-bus ethernet elan-name
```

The subinterface number used can be arbitrarily chosen. Remember that each subinterface (or each ELAN) is segmented from the others. Therefore, you can configure a different LES/BUS pair on one or more subinterfaces. Each pair will operate only for its assigned ELAN. The *elan-name* parameter is a text string (case sensitive) that identifies the name of the ELAN. This name must be defined the same in all LANE components.

To implement SSRP for redundant LES/BUS components in an ELAN, use the preceding commands on other LANE modules to create the redundant pieces. All of the redundant LES/BUS pairs will be configured into the LECS database.

Configuring the LECS

The LECS is configured on a major ATM interface, not on a subinterface. First you must build the LECS database of ELANs and their associated LES NSAP addresses. Configure the LECS database with the following commands:

```
ATM(Config)# lane database database-name
ATM(lane-config-database)# name elan1-name server-atm-address les1-nsap-address
ATM(lane-config-database)# name elan2-name server-atm-address les2-nsap-address
ATM(lane-config-database)# name ...
```

The *database-name* argument is a text string that identifies the LECS database as a whole. Several LECS databases can be defined, each with a unique name string, and applied to LECS components on individual ATM major interfaces. Usually, one LECS and one database are sufficient on a LANE module.

Each ELAN in the LANE network must be defined with a single **name** database command using the NSAP address of that ELAN's LES. Remember that you can find the NSAP address of the LES on a switch using the **show lane default** command.

To implement SSRP for redundant LES/BUS entities, use the same commands as above, with multiple NSAP addresses for each ELAN. For example, you can use the following commands:

```
ATM(Config)# lane database database-name
ATM(lane-config-database)# name elan1-name server-atm-address les1-nsap-address
ATM(lane-config-database)# name elan1-name server-atm-address les2-nsap-address
ATM(lane-config-database)# name elan2-name server-atm-address les1-nsap-address
ATM(lane-config-database)# name elan2-name server-atm-address les2-nsap-address
ATM(lane-config-database)# name ...
```

After the database has been built, the LECS must be enabled on the ATM major interface. This is done with the following commands:

```
ATM(Config)# interface atm number
ATM(Config-if)# lane config database database-name
ATM(Config-if)# lane config auto-config-atm-address
```

The LECS database that was first built is now referenced by its name and bound to the LECS process. The **auto-config-atm-address** option is used to tell the LANE module to use the automatically generated NSAP addresses for the LECS component. The LECS NSAP address must now be configured into the ATM switches so that other LANE components can automatically retrieve the LECS address from the switches via ILMI.

To implement SSRP for redundant LECS entities, first create two or more LECS components on different LANE modules. Because the LECS database specifies the order of other redundant components to use, the LECS database must be configured identically on all LECS machines. Then the NSAP addresses of the redundant LECS components can be configured into the ATM switch.

Configuring Each LEC

You must configure a LEC on each device where required. One LEC is necessary for each ELAN that a device participates in. The LEC configuration also specifies which VLAN the ELAN will be bridged to on the switch.

Each LEC is configured on a different ATM subinterface, using the following commands:

```
ATM(Config)# interface atm number.subint multipoint
ATM(Config-subif)# lane client ethernet vlan-num elan-name
```

The *vlan-num* argument references an existing VLAN number on the local switch. The *elan-name* argument references the name of an existing ELAN on the local LANE module. The two are then bridged on the LANE module and become a single broadcast domain.

Viewing the LANE Configuration

The configuration of each individual LANE component is fairly straightforward. However, because each component can be placed on a separate device, determining if all of the components are communicating and operating properly can be difficult. Catalyst LANE modules offer a number of commands that can be used to display and debug LANE configurations and status.

Viewing Default NSAP Addresses

To view the default NSAP addresses for the local LANE module, use the **show lane default** command as demonstrated previously in Example 6-1.

Viewing LES Status

To view the status of an LES, use the **show lane server** command on the LES machine. Example 6-2 demonstrates this command for a switch that is the LECS for two ELANs:

Example 6-2 **show lane server** Output Displays LES Status

```
CatalystA_ATM#show lane server
LE Server ATM0.1 ELAN name: lan1 Admin: up State: operational
type: ethernet Max Frame Size: 1516
ATM address: 47.00001606288000300000000F.0050A28D5881.01
LECS used: 47.00001606288000300000000F.0050A28D5883.00 connected, vcd 1460
control distribute: vcd 1485, 2 members, 479995 packets

proxy/ (ST: Init, Conn, Waiting, Adding, Joined, Operational, Reject, Term)
lecid ST vcd pkts Hardware Addr ATM Address
  1P 0 1474 479994 0050.a28d.5880 47.00001606288000300000000F.0050A28D5880.0B
  2P 0 1505 3 0090.6f7a.1c80 47.00001606288000100000000F.00906F7A1C80.01

LE Server ATM0.2 ELAN name: lan2 Admin: up State: operational
type: ethernet Max Frame Size: 1516
ATM address: 47.00001606288000300000000F.0050A28D5881.02
LECS used: 47.00001606288000300000000F.0050A28D5883.00 connected, vcd 1415
control distribute: vcd 1427, 2 members, 2 packets

proxy/ (ST: Init, Conn, Waiting, Adding, Joined, Operational, Reject, Term)
lecid ST vcd pkts Hardware Addr ATM Address
  1P 0 1397 2 0050.a28d.5880 47.00001606288000300000000F.0050A28D5880.0C
  2P 0 1511 2 0090.6f7a.1c80 47.00001606288000100000000F.00906F7A1C80.02
```

The shaded lines in Example 6-2 show two LES components, one on interface ATM 0.1 and one on ATM 0.2. Both LESs are shown to be fully functional, listed in the operational state. Also note that the ELAN name is given for each LES.

Viewing BUS Status

To view the status of a BUS, use the **show lane bus** command. Example 6-3 demonstrates this command on the same switch used in Example 6-2. This switch is configured as the LES/BUS for two ELANs:

Example 6-3 **show lane bus** Output Displays BUS Status

```
CatalystA_ATM#show lane bus
LE BUS ATM0.1 ELAN name: lan1 Admin: up State: operational
type: ethernet Max Frame Size: 1516
ATM address: 47.00001606288000300000000F.0050A28D5882.01
data forward: vcd 1499, 2 members, 523298 packets, 0 unicasts
```

continues

Example 6-3 *show lane bus* Output Displays BUS Status (Continued)

lecid	vcd	pkts	ATM Address
1	1489	505572	47.00001606288000300000000F.0050A28D5880.0B
2	1506	17726	47.00001606288000100000000F.00906F7A1C80.01
LE BUS ATM0.2 ELAN name: lan2 Admin: up State: operational			
type: ethernet Max Frame Size: 1516			
ATM address: 47.00001606288000300000000F.0050A28D5882.02			
data forward: vcd 1451, 2 members, 0 packets, 0 unicasts			
lecid	vcd	pkts	ATM Address
1	1433	0	47.00001606288000300000000F.0050A28D5880.0C
2	1512	0	47.00001606288000100000000F.00906F7A1C80.02

The shaded lines in Example 6-3 show two BUS components on interfaces ATM 0.1 and ATM 0.2. Both BUSs are in the operational state and are listed with their respective ELAN names. Each BUS is also shown with a breakdown of packet forwarding activity to individual LECs within the ELAN.

Viewing the LECS Database

To view the current LECS database, use the **show lane database** command on the LECS machine. Example 6-4 demonstrates the output generated by this command.

Example 6-4 *show lane database* Output Displays the LECS Database Contents

CatalystA_ATM#show lane database	
LANE Config Server database table 'company_db' bound to interface/s: ATM0	
no default elan	
elan 'lan1': un-restricted	
server 47.00001606288000300000000F.0050A28D5881.01 (prio 0) active	
elan 'lan2': un-restricted	
server 47.00001606288000300000000F.0050A28D5881.02 (prio 0) active	

The shaded lines in the output in Example 6-4 show that the LECS has a database called `company_db`, which contains a number of LES and ELAN entries. The LECS is assigned to interface ATM0 (the major physical interface). The contents of the database are then listed by ELAN, each containing a single LES. All LESs are in the active state, and are listed with their complete NSAP addresses.

Viewing LEC Status

To view the status of a LANE Client, use the **show lane client** command as demonstrated in Example 6-5 for a switch with two LECs (two ELANs).

Example 6-5 show lane client Output Displays the LEC Status

```

CatalystA_ATM#show lane client
LE Client ATM0.11  ELAN name: lan1  Admin: up  State: operational
Client ID: 1          LEC up for 11 days 2 hours 43 minutes 9 seconds
Join Attempt: 20
HW Address: 0050.a28d.5880  Type: ethernet  Max Frame Size: 1516
VLANID: 1
ATM Address: 47.00001606288000300000000F.0050A28D5880.0B

VCD  rxFrames  txFrames  Type          ATM Address
0    0          0         0  configure
47.00001606288000300000000F.0050A28D5883.00

1477      1      480139  direct  47.00001606288000300000000F.0050A28D5881.01

1481     480140      0  distribute
47.00001606288000300000000F.0050A28D5881.01

1491      0      505675  send    47.00001606288000300000000F.0050A28D5882.01

1495     17729      0  forward  47.00001606288000300000000F.0050A28D5882.01

LE Client ATM0.12  ELAN name: lan2  Admin: up  State: operational
Client ID: 1          LEC up for 11 days 2 hours 43 minutes 12 seconds
Join Attempt: 19
HW Address: 0050.a28d.5880  Type: ethernet  Max Frame Size: 1516
VLANID: 2
ATM Address: 47.00001606288000300000000F.0050A28D5880.0C

VCD  rxFrames  txFrames  Type          ATM Address
0    0          0         0  configure
47.00001606288000300000000F.0050A28D5883.00

1403      1          2  direct
47.00001606288000300000000F.0050A28D5881.02

1421      2          0  distribute
47.00001606288000300000000F.0050A28D5881.02

1439      0          0  send
47.00001606288000300000000F.0050A28D5882.02

1445      0          0  forward
47.00001606288000300000000F.0050A28D5882.02

```

Each LEC that is present on the local switch is listed with the show lane client command. The shaded lines show that there are two LECs on this switch—one assigned to interface ATM 0.11 and one to ATM 0.12. Both are up and operational. The most important information given here are the LEC uptimes and join attempts. Here, both LECs have been up for a little more than 11 days. When a LEC is not operating properly, this command will show that the LEC is down and how many attempts the LEC has made to join the ELAN. A timer value will also be shown that tells when the LEC will try to join the ELAN again.

Foundation Summary

The Foundation Summary is a collection of tables and figures that provide a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final prep before the exam, these tables and figures will hopefully be a convenient way to review the day before the exam.

Table 6-4 *ATM Reference Model*

OSI Reference Model	ATM Reference Model
Layer 7 (application)	Higher Layers
Layer 6 (presentation)	
Layer 5 (session)	
Layer 4 (transport)	
Layer 3 (network)	
Layer 2 (data link)	ATM Adaptation Layers (AAL)
	ATM Layer
Layer 1 (physical)	ATM Physical Layer

Table 6-5 *ATM Adaptation Layers*

ATM	
Adaptation Layer	Function
AAL1	Supports connection-oriented services that can emulate leased line circuits. AAL1 requires timing synchronization between source and destination, allowing the transport of real-time streams like voice and video. Data is sampled and put into cell payloads, along with sequence information, before transmission.
AAL2	Supports connection-oriented services for variable bit rate (VBR) applications.
AAL3/4	Supports both connection-oriented and connectionless services. AAL3/4 was designed for use with Switched Multimegabit Data Service (SMDS). Data is segmented into cell payloads, while CRC error information is added.
AAL5	Supports both connection-oriented and connectionless services. Non-SMDS data, like TCP/IP over ATM and LANE, is transported in sequence with simple segmentation. A CRC value is added to the end of the pre-segmented frame.

Table 6-6 *Types of ATM Circuits*

ATM Circuit Type	Function
PVC	<i>Permanent Virtual Circuit</i> —manually configured between endpoints.
SVC	<i>Switched Virtual Circuit</i> —dynamically configured between endpoints.
Point-to-point	Circuit between exactly two endpoints.
Point-to-multipoint	Circuit from one endpoint to multiple endpoints.
VCC	<i>Virtual Channel Connection</i> —ATM circuit built between two endpoints.
VP	<i>Virtual Path</i> —a bundle of virtual channels.
Transmission Path	A bundle of virtual paths.

Table 6-7 *Components of a 20-Byte ATM NSAP Address*

NSAP		
Component	Size	Function
Prefix	13 bytes	Uniquely identifies each ATM switch.
ESI	6 bytes	Uniquely identifies device attached to ATM switch (MAC address).
Selector	1 byte	Identifies process on ATM device (ATM subinterface number).

Table 6-8 *LANE Components*

LANE	
Component	Function
LEC	Provides basic ELAN function on ATM device; bridges ELAN to VLAN.
LES	Provides MAC-to-NSAP address resolution for one ELAN.
BUS	Provides broadcast functionality for all clients of one ELAN.
LECS	Provides LES addresses and ELAN membership for all clients of all ELANs.

Table 6-9 *Sequence of Communication as a LEC Initializes*

Step	Activity
1	LEC finds LECS address (ILMI, etc); LEC contacts LECS for membership and LES NSAP address.
2	LEC contacts LES; LEC registers its own MAC and NSAP addresses with LES.
3	LEC contacts BUS via broadcast address (0xFFFFFFFFFFFF); BUS adds LEC as a leaf node on multipoint VC.
4	LEC builds Data Direct VCCs to other LECs as needed.

Table 6-10 *Automatic NSAP Address Generation for LANE Components*

LANE Component	Prefix	ESI	Selector
LEC	From ATM Switch	MAC Address	ATM Subinterface
LES	From ATM Switch	MAC Address + 1	ATM Subinterface
BUS	From ATM Switch	MAC Address + 2	ATM Subinterface
LECS	From ATM Switch	MAC Address + 3	.00

Table 6-11 *Summary of Catalyst LANE Configuration Commands*

Catalyst Command	Function
<code>show lane default</code>	Display default or autoconfigured NSAP addresses for all LANE components on the switch.
<code>session module-num</code>	Open a user interface session with the switch module in slot <i>module-num</i> . This is needed to begin a session with a LANE module.
<code>lane server-bus ethernet elan-name</code>	Configure LES/BUS pair on an ATM subinterface (subinterface config mode).
<code>lane database db-name</code> <code>name elan-name server-atm-address</code> <code>les-nsap-address</code> <code>name ...</code>	Configure LECS database.
<code>lane config-database db-name</code> <code>lane config-auto-atm-address</code>	Apply LECS database to an ATM major interface (interface config mode); Configure automatic ATM NSAP addressing for LANE components.
<code>lane client ethernet vlan-num elan-name</code>	Configure LEC for VLAN number and ELAN name on ATM subinterface (subinterface config mode).
<code>show lane server</code>	Display status of LES on the switch.
<code>show lane bus</code>	Display status of BUS on the switch.
<code>show lane database</code>	Display status of LECS on the switch.
<code>show lane client</code>	Display status of LEC on the switch.

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however they are designed to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to these questions can be found in Appendix A, on page 477.

- 1 What is the basic unit of ATM data? What is its basic format (header, payload, etc.)?

- 2 What process allows an IP packet to be transported within ATM cells?

- 3 What is an ATM edge device? What Cisco devices can be used?

- 4 What types of VCs can be built with ATM?

- 5 List the hierarchy of ATM VCs.

6 What type of addressing is used to identify ATM devices?

7 What are the three basic components of a NSAP address?

8 What information is carried within each ATM data unit to specify how to get from the source to the destination?

9 Name two inherent ATM protocols used by ATM switches to communicate with other ATM devices.

10 What is LANE used for?

11 Where should the LECS NSAP address be configured in a LANE network? Where should the LES NSAP address be configured?

12 What are the functional components of LANE?

13 What is the effect of a failed LECS on a LANE network? A failed LES? A failed BUS? A failed LEC?

14 When a LEC initializes, what LANE component does it contact first? Why?

15 What is the difference between a VLAN and an ELAN?

16 When is an LE_ARP used?

17 The following NSAP addresses were obtained from **show lane default**. Match the LANE components to their NSAP addresses:

47.00001606288000300000000F.0064A28D5EA0.**

47.00001606288000300000000F.0064A28D5EA1.**

47.00001606288000300000000F.0064A28D5EA2.**

47.00001606288000300000000F.0064A28D5EA3.00

18 How many of each LANE component are necessary for LANE operation?

19 A network consists of two Catalyst switches connected by an ATM switch. Four VLANs exist on each Catalyst and are trunked between all Catalyst switches. How many LECS components are present? How many LES components? How many BUS components? How many LECs?

20 Which LANE component performs bridging between a VLAN and an ELAN?

21 Given an NSAP address of 47.00001606288000300000000F.0064A28D5EA1.0A, what LANE component does this represent? On which subinterface of ATM 0 is this assigned?

22 What NSAP addresses must be configured into the LECS database?

23 A LANE database (mylane-db) has already been created on a Catalyst switch. The switch contains a single ATM interface (ATM 0) and subinterfaces ATM 0.1 and ATM 0.2. What commands can be used to enable the LECS with this database on the ATM link?

24 If a switch supports four VLANs, how many LECs will be needed? How many ATM subinterfaces will be required?

25 What is SSRP? Which LANE components can be configured for SSRP?

26 How many ATM subinterfaces are required to place both an LES and a BUS on a Catalyst switch?

27 What Catalyst switch command can be used to view the current status of each LANE component (LECS, LES, BUS, and LEC)?

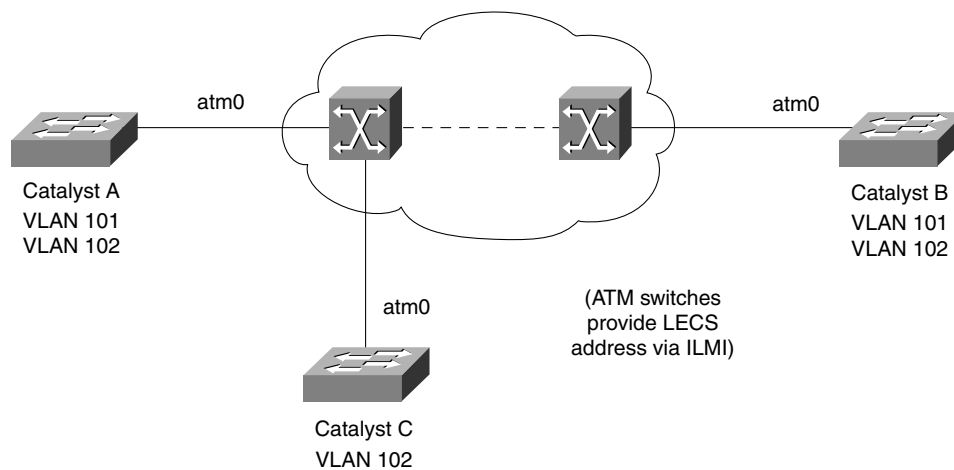
28 Is it possible to place all LANE components (LECS, LES, BUS, and LEC) on a single Catalyst switch? If only one VLAN is required, what is the minimum number of interfaces required?

Scenarios

Scenario 6-1

Given the network diagram shown in Figure 6-6, perform the following exercises:

Figure 6-6 Network Diagram for Scenario 6-1



- 1 Choose locations where each of the LANE components can be configured, so that VLANs 101 and 102 have connectivity across the ATM cloud. (Ignore SSRP for this scenario.)
- 2 Assume the ELANs in this network are named “ELAN101” and “ELAN102”. What commands are needed to configure a LEC for ELAN 101 on Catalyst A?
- 3 With the network shown in Figure 6-5, can a host on Catalyst A VLAN 101 communicate with a host on Catalyst B VLAN 101? With a host on Catalyst C VLAN 102?
- 4 A station on Catalyst C VLAN 102 sends a broadcast. Which Catalysts and VLANs will see the broadcast?
- 5 Suppose Catalyst A is configured with the LECS and Catalyst B is configured with the LES for ELAN 102. On Catalyst C, **show lane client** shows that the LEC for ELAN 102 has not joined ELAN 102. What are some possible causes?
- 6 Suppose Catalyst A is the LECS, Catalyst C is the LES/BUS for ELAN 102, and Catalyst B has a LEC for ELAN 102. A station on Catalyst B tries to ping a station on Catalyst A. Where is an ARP request broadcast seen first?

- 7 Now suppose that when the station on Catalyst A replies, the LEC for ELAN 102 on Catalyst A only knows Catalyst B's LEC 102 MAC address. How does Catalyst A's LEC 102 find the NSAP address for Catalyst B's LEC 102?
- 8 You now need to add an additional VLAN (103) to only Catalysts B and C. What steps should you take to accomplish this?

Scenarios Answers

Scenario 6-1 Answers

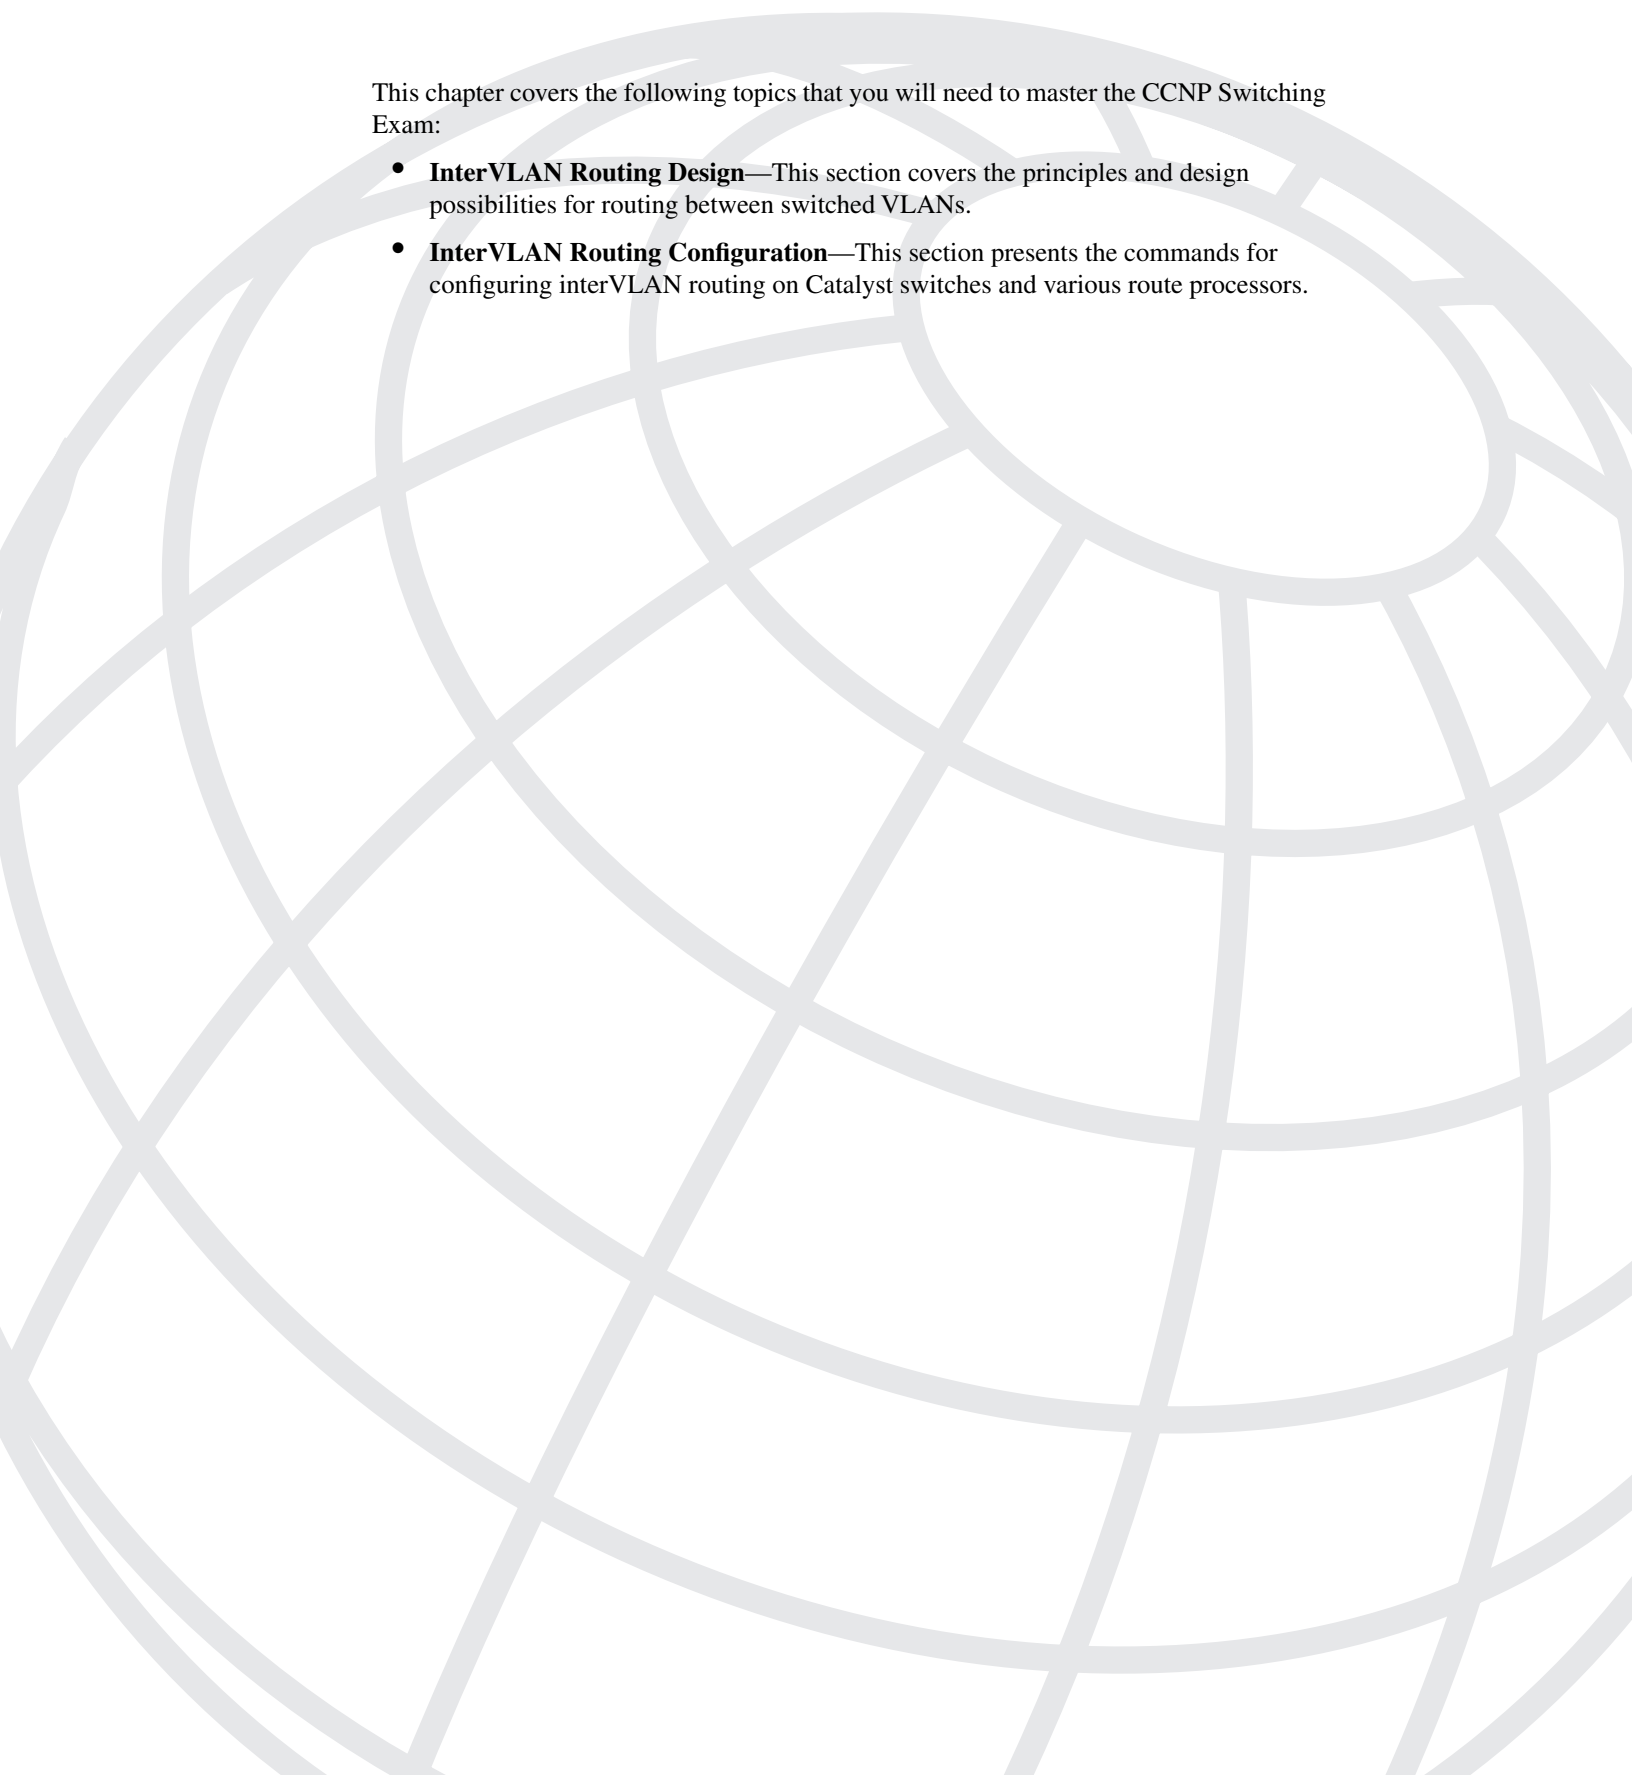

- 1 The LECS could be placed on either Catalyst A, B, or C. The LES/BUS for ELAN 101 could be placed on either A or B, and the LES/BUS for ELAN 102 could be placed on either A, B, or C. A LEC for ELAN 101 must be placed on Catalyst A *and* B. A LEC for ELAN 102 must be placed on all three Catalysts.
- 2 The commands needed are as follows:

```
interface atm 0.1 multipoint
lane client Ethernet 101 ELAN101
```
- 3 A host on Catalyst A VLAN 101 can indeed communicate with a host on Catalyst B VLAN 101. Because VLAN 101 is trunked across the ATM cloud using LANE, end-to-end communication is possible. The same host on Catalyst A cannot communicate with a host on Catalyst C VLAN 102, however. VLANs 101 and 102 are segmented, as are ELANs 101 and 102. Without a router, hosts cannot talk across the VLANs.
- 4 When the station sends the broadcast, the LEC for ELAN 102 on Catalyst C receives it first. The broadcast is sent on to the BUS (wherever it is configured) and then on to all three Catalysts because VLAN/ELAN 102 is configured on all switches. VLAN 102 will be the only VLAN to receive the broadcast because VLAN 102 is the boundary of the broadcast domain.
- 5 The LEC will not be able to join the ELAN if the following conditions occur:

The LES NSAP address has been misconfigured in the LECS database. In that case, the LEC will contact the LECS and receive the wrong LES address. Because the appropriate LES can't be reached, the LEC will not be able to join the ELAN.

If the LECS NSAP address is misconfigured in the ATM switch, the LEC will receive an incorrect address via ILMI. Then the LEC will never be able to query a LECS for the address of the LES.
- 6 The ARP request is received by the LEC for ELAN 102 on Catalyst B first because it is the bridge between VLAN 102 (where the station is located) and ELAN 102. The LEC then forwards the broadcast on to the BUS on Catalyst C.
- 7 Because the LEC on Catalyst A doesn't know the NSAP address for the LEC on Catalyst B, it must query the LES for ELAN 102 with an LE_ARP request. The LES (located on Catalyst C) will resolve the MAC address and return the LEC's NSAP address.

- 8** First, define VLAN 103 on Catalysts B and C, either manually or by using VTP. Then add an entry in the LECS database for a LES/BUS in ELAN 103. A LES and BUS for ELAN 103 must be configured, either on Catalyst B or C. Finally, a LEC bridging VLAN 103 to ELAN 103 must be added to both Catalyst B and C where the VLAN exists.



This chapter covers the following topics that you will need to master the CCNP Switching Exam:

- **InterVLAN Routing Design**—This section covers the principles and design possibilities for routing between switched VLANs.
- **InterVLAN Routing Configuration**—This section presents the commands for configuring interVLAN routing on Catalyst switches and various route processors.

InterVLAN Routing

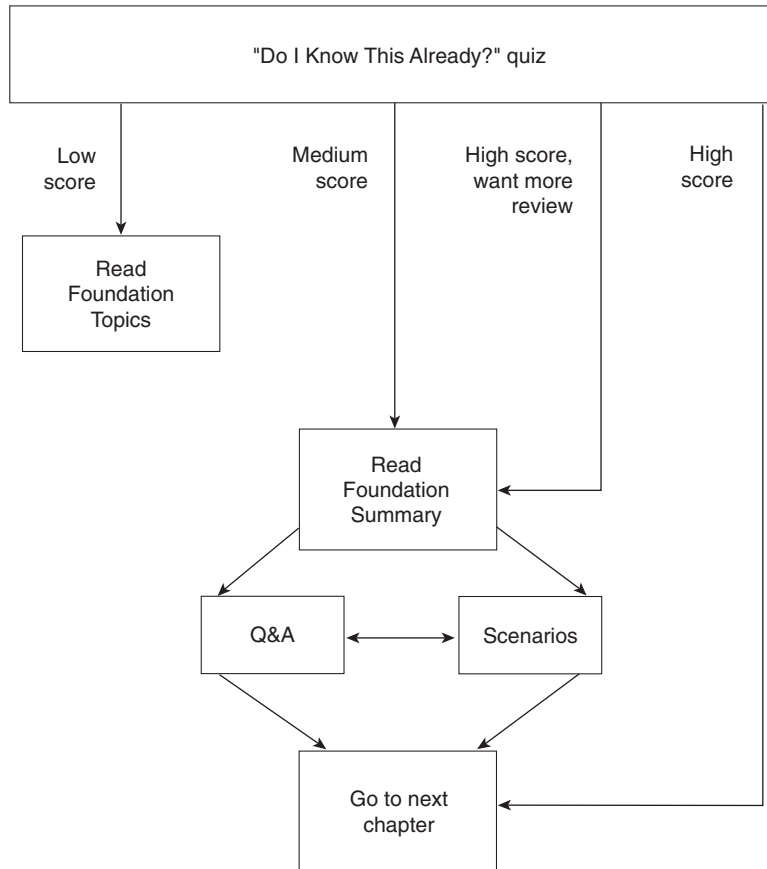
One of the main features of VLANs is network segmentation. Multiple VLANs can exist on a switch and be transported or trunked between switches. VLANs, however, remain isolated from each other. A station on one VLAN cannot communicate with a station on another VLAN unless some assistance is there from an additional device. Two VLANs can be bridged together, or they can be joined by a Layer 3 routing function.

This chapter discusses routing between VLANs to provide complete connectivity across the switched network. Several design methodologies are presented, along with Cisco Catalyst and router configuration procedures for interVLAN routing.

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and the answers for all your work with this book in one place, for easy reference.
- Take the “Do I Know This Already?” quiz, and write down facts and concepts (even if you never look at the information again).
- Use the diagram in Figure 7-1 to guide you to the next step.

Figure 7-1 *How to Use This Chapter*

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The quiz helps you make good choices of how to spend your limited study time. The quiz is sectioned into two smaller “quizlets,” which correspond to the two major headings in the Foundation Topics section of the chapter. Although your answer may differ somewhat from the answers given, finding out if you have the basic understanding presented in this chapter is more important. You will find that these questions are open-ended, rather than multiple choice as found on the exams. This will enable you to focus more on understanding the subject matter than on memorizing details.

Use the scoresheet in Table 7-1 to record your score.

Table 7-1 *Scoresheet for Quiz and Quizlets*

Quizlet Number	Foundation Topics Sections Covering These Questions	Questions	Score
1	InterVLAN Routing Design	1–4	
2	InterVLAN Routing Configuration	5–8	
All questions		1–8	

1 Where can a router be placed in relation to switches for interVLAN routing?

2 What types of links can be used to interconnect switches and an external router? How many VLANs can be carried on each?

3 What trunking methods can a router support?

4 What is the difference between interVLAN routing and multilayer switching (MLS)?

5 What Catalyst commands can be used to locate and then connect to an internal route processor?

- 6 What should be configured on a route processor to dynamically determine routing paths to remote networks?

- 7 If a router is used to route between VLANs, what additional information is needed so that traffic will actually be routed?

- 8 Suppose a router connects four VLANs of a switched network, providing interVLAN routing. If the router is then configured for bridging to support nonroutable protocols, how would the network be affected?

The answers to the quiz are found in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections,” on page 477. The suggested choices for your next step are as follows:

- **4 or less overall score**—Read the entire chapter. This reading includes the “Foundation Topics” and “Foundation Summary” sections, and the Q&A section.
- **5–6 overall score**—Begin with the “Foundation Summary” section. Then follow with the Q&A section at the end of the chapter.
- **7 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” section. Then go to the Q&A section appearing at the end of the chapter. Otherwise, move on to the next chapter.

Foundation Topics

InterVLAN Routing Background

Several chapters have dealt with the foundations of network design, VLANs, and trunking. However, VLANs have been presented as isolated broadcast domains. The next two chapters discuss how to transport traffic between VLANs using Layer 3 functions.

VLANs are typically configured on Layer 2 switches to form broadcast domains. VLANs can exist in one or more switches through the use of trunking. As well, VLANs usually represent subnetworks of Layer 3 protocols. Since Layer 2 switches do not use Layer 3 addressing to make switching decisions, a Layer 3 decision is needed to move packets between VLANs.

Traditionally, routers have performed this function. Routers have rich feature sets that include intelligent, dynamic routing protocols for packet transport and many packet filtering capabilities. Additionally, many types of connectivity enhancements like DHCP relaying, Network Address Translation, Quality of Service, and policy-based routing add to the functionality. Routers also offer connectivity between LAN technologies and Wide Area Network technologies, connecting a broad range of network media.

Routing in a switch network can take on three forms: interVLAN routing, multilayer switching (MLS), and Cisco Express Forwarding (CEF). InterVLAN routing is based on adding a route processor somewhere in the switched network to provide Layer 3 routing. *Every* packet destined from one VLAN to another must pass through the router. MLS, on the other hand, is based on the principle that the router only sees the *first* packet of a conversation. Switching paths are then set up so that subsequent packets bypass the router and are switched by a more efficient “shortcut” path. CEF is a distributed switching mechanism keeping copies of route cache information in several different forms to be used for efficient switching. Catalyst switches can hand off packets to a CEF-capable router for processing, as in interVLAN routing. Some Catalyst platforms implement CEF directly in hardware.

The following sections discuss interVLAN routing at length, using a variety of designs and configurations. MLS is discussed in Chapter 8, “Multilayer Switching.” Information about CEF can be found at www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/switch_c/xcp2/xcdcef.htm.


InterVLAN Routing Design

Several options are available when placing a route processor within a switch campus network. This section discusses the techniques, which are primarily based on the type of connectivity between the switches and the router, and location of the route processor.

Routing with Multiple Physical Links

The simplest and most straightforward method of routing between VLANs is to use several physical links between switches and an external router. Each link is configured for a single VLAN, so that a link is there for each VLAN to be routed. This approach is shown in Figure 7-2. Catalyst A is configured for three VLANs (1, 2, and 3). The switch is connected to the router using three separate links, each assigned to one of the three VLANs. Catalysts B and C each have only one VLAN (4 and 5, respectively). Each requires only a single link to the router. Therefore, five VLANs exist on this network and are interconnected by a single router with five network links.

Figure 7-2 *InterVLAN Routing Using Multiple Physical Links*



To illustrate the routing operation, suppose a station connected to VLAN 1 on Catalyst A needs to communicate with a station connected to VLAN 5 on Catalyst B. Because the destination has an IP address that is off the local network, the source station will send the packets to the router. These packets will be sent out the VLAN 1 link on Catalyst A. The packets will be examined by the router and sent over the VLAN 5 link to Catalyst B. In this fashion, every packet travels across a link reserved for traffic on a specific VLAN and passes through the router.

Using one VLAN per link offers an intuitive approach to routing between VLANs. Routers naturally associate each physical link with a subnetwork (some Layer 3 protocol and address range), and transport packets between links. Each link is also inherently segmented from the others, unless bridging arrangements are made within the router. This method is useful when the switches and router are already available and can be quickly connected, using a small number of VLANs. No special configuration is needed, other than the usual interface addressing used on the router.

As the network grows, this method can quickly get out of hand. Every additional VLAN requires an additional physical link to the router. Clearly, the expense and logistics of adding a large number of connections to the router outweigh the simplicity of the design.

Routing over Trunk Links


A more robust and cost effective approach uses trunk links between the switches and routers, instead of multiple physical links. Because trunk links transport multiple VLANs over a single link, only one link to an external router is required. A router connected to a switch by a single trunk link is usually referred to as a *router on a stick*, or a *one-armed router*. However, a router can also connect to several switches using trunk links. This connection provides end-to-end Layer 3 connectivity between blocks of switches.

This section discusses three types of trunk links that can be used: IEEE 802.1Q, ISL, and ATM LANE.

802.1Q and ISL Trunks

Both IEEE 802.1Q and Inter-Switch Link (ISL) trunks can be used to transport multiple VLANs to a router. Both encapsulation methods use Fast Ethernet or Gigabit Ethernet as the physical media for trunking. Figure 7-3 shows a typical application using a router and ISL trunk links for interVLAN routing. Catalyst A is connected to the router using a trunk link that transports VLANs 1, 2, and 3.

Recall from Chapter 4, “VLANs and Trunking,” that both 802.1Q and ISL trunk links identify each frame with a VLAN number. As a frame leaves a switch, the frame is encapsulated and identified with its VLAN. When the router receives a frame over a trunk link, the router unencapsulates the frame and associates it with an interface assigned to the VLAN number that identified the frame. Recall also that frames from the native VLAN of an 802.1Q trunk are not tagged with the VLAN number.

Figure 7-3 *InterVLAN Routing Using Trunk Links*

As in the previous example, suppose a workstation attached to VLAN 1 on Catalyst A needs to send a packet to a station on VLAN 5 on Catalyst B. The station sends the frame to the router, using the VLAN connected to that station. The packet is sent out over the ISL trunk link, identified with VLAN 1. The router receives the packet and unencapsulates it, examining the identifier. The packet is inbound on VLAN 1 and destined for a station on VLAN 5. The router makes a routing decision and forwards the packet toward Catalyst B, after encapsulating it in the ISL trunk with VLAN identifier 5.

On a router, 802.1Q and ISL trunks are connected by either a Fast or Gigabit Ethernet interface. (EtherChannel is also an option that can be used.) To support various VLANs, individual subinterfaces are configured with either 802.1Q or ISL encapsulation and a VLAN number.

802.1Q and ISL trunks offer the advantage of scalability because a single link can transport many VLANs; however, some CPU overhead is involved as the router processes the encapsulation. Therefore, the router cannot use its most efficient packet switching method for packet forwarding. Trunking encapsulations also require some link bandwidth overhead, as either a 30-byte header (ISL) or a 4-byte header (802.1Q) increases each frame.

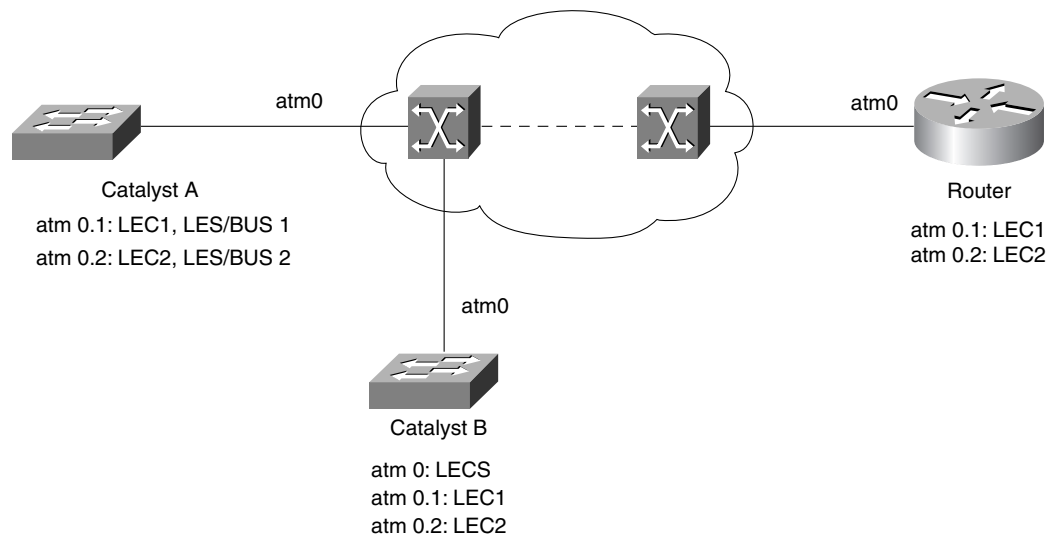
ATM LANE

As discussed in Chapter 6, “Trunking with ATM LANE,” switches can trunk VLANs to other switches over ATM links using the LAN Emulation (LANE) standard. Switches with LANE modules can participate by operating local LANE Clients (LECs) for each VLAN-to-emulated LAN (ELAN) bridge. The LANE modules must also communicate with the other LANE components (LANE Configuration Server [LECS], LANE Server [LES], and BUS) in the network to become active members in the emulated LANs.

If only switches are involved in LANE, each ELAN remains isolated from other ELANs. This isolation follows the principle of VLANs, which require a Layer 3 routing function to move data between VLANs. Routing can be performed between ELANs trunked over LANE by adding an ATM interface to an external router.

The router can use a single ATM link to connect to an ATM switch—not directly to an ATM link on a Catalyst switch. Figure 7-4 shows an example network using a LANE-attached router to route between ELANs. The router interacts with LANE just as a switch does: it must find the LECS; operate a LEC for each ELAN that it supports; locate and join the LES and BUS for each ELAN; and build Data Direct virtual circuits (VCs) to other LECs in the network.

Figure 7-4 *InterELAN Routing Using ATM LANE*



A major ATM interface on the router is broken up into logical subinterfaces. The major interface can be configured with an LECS, if required. The individual subinterfaces represent single ELANs and can be configured with LES/BUS pairs or a LEC for the ELAN. Packets coming in over a LEC on one subinterface are processed by the router and forwarded out another LEC on a different subinterface, depending on the Layer 3 addresses and other routing decisions.

Routing with an Integrated Router

Early approaches to interVLAN routing used external routers and physical links. Further developments have moved the route processor *inside* the switch, for convenience and tighter integration of the Layer 2 and 3 components.

Cisco offers several types of integrated route processors in its Catalyst switch family.

- **Route Switch Module (RSM) and Route Switch Feature Card (RSFC)**—Route processor modules for the Catalyst 5000 family. The RSM is based on the Cisco 7500 Route Switch Processor and runs the Cisco IOS software. The RSFC is based on the 7200 series Route Switch Processor, runs IOS, and operates with the Catalyst 5000 Supervisor IIG and IIIG.
- **Multilayer Switch Feature Card (MSFC)**—An IOS-based route processor for the Catalyst 6000/6500 families that provides MLS. Rather than provide interVLAN routing, the MSFC works in conjunction with the Catalyst switching engine to set up and shortcut switch Layer 3 and higher protocols. As noted below, MLS and the MSFC are covered in Chapter 8.

Each of these route processors can also participate in MLS, where the switch creates shortcut paths with the assistance of a route processor. Because MLS is covered in the following chapter, this section only discusses the application of interVLAN routing with these modules.

The RSM interfaces to the Catalyst 5000/5500 backplane through two direct memory access (DMA) connections. The channels are referred to as *Channel0* and *Channel1*, each providing 200-Mbps throughput. VLANs are supported through the use of virtual VLAN interfaces (that is, VLAN 1, VLAN 2, and so on).

InterVLAN Routing Configuration

This section presents a more detailed look at configuring interVLAN routing using both external and integrated Cisco route processors. Routing, in itself, is a complex and extensive topic. Only the basic routing protocol configuration is presented here, to provide simple but functional routing between VLANs. If you need further information on the various routing protocols, refer to the Interconnecting Cisco Network Devices (ICND) and Building Scalable Cisco Networks (BSCN) courses or course books and the *Cisco CCNP Routing Exam Certification Guide*.

Accessing the Route Processor

To begin interVLAN routing configuration, the route processor must first be accessed. On an external router, a terminal emulator program can be used to connect directly with the console port. If some IP connectivity is already available on the router, a telnet session can be opened to the router.

An integrated or internal route processor must first be located in the switch chassis. Use the **show module** command on a Catalyst switch to get a listing of the installed modules (see Example 7-1). The route processor is the module with *Route Switch* in the Module-Type field. For example, an RSM is installed in slot 8 in this Catalyst switch.

Example 7-1 `show module` Command Output Displays a Listing of Installed Modules on a Catalyst Switch

Switch (enable) <code>show module</code>						
Mod	Module-Name	Ports	Module-Type	Model	Serial-Num	Status
1		2	1000BaseSX Supervisor	WS-X5530	010849466	ok
2		24	10/100BaseTX Ethernet	WS-X5224	010445883	ok
3		24	10/100BaseTX Ethernet	WS-X5224	010433970	ok
4		24	10/100BaseTX Ethernet	WS-X5224	010435124	ok
5		24	10/100BaseTX Ethernet	WS-X5224	010432035	ok
7		24	10/100BaseTX Ethernet	WS-X5224	010909044	ok
8		1	Route Switch	WS-X5302	012766975	ok
9		2	SM OC-3 Dual-Phy ATM	WS-X5157	010995049	ok

Other integrated route processors can be located in the same fashion. An RSFC module in a Catalyst 5000 switch will be listed as *Route Switch Feature Card*. A MSM module (now obsolete) in a Catalyst 6000 switch will be listed as *Multilayer Switch Module*.

To establish a terminal session with the integrated route processor, use the `session` Catalyst switch command with the module number as an argument. For example, if an RSM is located in slot 5, then `session 5` will bring up a new session with a user interface to the RSM. Note that the integrated route processors run Cisco IOS. Therefore, the user interface and command set may be different from that of the host switch. The `session` command essentially starts a Telnet session with the route processor. By using the `exit` IOS command, the router session is terminated and the switch session is resumed.

For future identification and readability, you should assign a hostname to the route processor at this point. The `hostname name` command assigns the string *name* as the router's hostname and also as its command line prompt.

Establishing VLAN Connectivity

Next, the route processor will need to have its interfaces configured to support connectivity to the necessary VLANs. This is accomplished using interfaces and commands that are unique to the route processor hardware. The following sections present this information for each type of link between the route processor and switches.

Establishing VLAN Connectivity with Physical Interfaces

External routers are connected to switches using traditional LAN media links. For example, Ethernet, Fast Ethernet, Gigabit Ethernet, or Token Ring could be used. Here, individual physical router interfaces are configured for a single network each and connected to non-trunk switch ports configured for VLAN membership. By way of the physical connection, the router interface inherits the VLAN identity of the switch port.

To configure a physical interface, enter the configuration mode and the interface configuration mode. Then assign a network address to the interface and make sure the interface is in operation. The commands necessary for these operations are as follows:

```
Router# configure terminal
Router(config)# interface media module/port
Router(config-if)# description description-string
Router(config-if)# ip address ip-addr subnet-mask
Router(config-if)# no shutdown
```

The interface is identified by its media type, module, and port number. For example, a Fast Ethernet interface on module slot 0 and port 3 would be referenced as *interface fastethernet 0/3*. A text string can be given as a comment or description to further identify the interface. A network address, or an IP address in this case, is assigned to the interface. Finally, the interface is forced to be active using the **no shutdown** command.

Establishing VLAN Connectivity with Trunk Links

When an external router is connected to a switch by a trunk link, the trunk must be configured. The physical interface on the router must be Fast Ethernet or Gigabit Ethernet to support trunking and VLAN encapsulation. In addition, the physical interface is identified with a slot number and a major interface number. For example, a Fast Ethernet interface could be located in slot 3, port 2 of the router—**interface fastethernet 3/2**.

Once trunking is enabled on the interface, each VLAN in the trunk is represented by a subinterface number. These numbers can be arbitrarily chosen, but must be unique within the major interface number. Subinterface 1 in the previous example would be identified as **interface fastethernet 3/2.1**.

For each VLAN to be connected, trunking and VLAN encapsulation must be configured on the respective subinterface. Then the subinterface is assigned a network address. The router commands necessary for these steps are as follows:

```
Router(config)# interface module/port.subinterface
Router(config-if)# encapsulation [isl | dot1q] vlan-number
Router(config-if)# ip address ip-address subnet-mask
```

Once the subinterface has been identified, VLAN configuration takes place with the **encapsulation** command. Both IEEE 802.1Q and ISL trunking encapsulations are supported. Notice that the subinterface is assigned to a specific VLAN with this command's *vlan-number* field. In this fashion, each subinterface can be associated with a different VLAN number. Because a trunk link is being used, the VLAN numbers will match on both the router and switch ends.

Establishing VLAN Connectivity with LANE

In some cases, an external router must be connected to campus switches via an ATM network. LANE must be used, and the router must be equipped with an ATM Interface Processor, or module, that supports LANE. Both routers and Catalyst switch LANE modules run Cisco IOS, so the commands used to configure LANE are identical. Chapter 6 can be referenced for further information on LANE.

LANE trunking is similar to Ethernet trunking in two ways: a major interface is used for connectivity, and individual subinterfaces correspond to single VLANs. On the major interface, ATM is configured. Each ATM subinterface is configured with a LEC per ELAN. Recall that the LANE module on a Catalyst switch performs bridging between the VLAN on the LAN side and the ELAN on the ATM side. A router only needs to route packets between ELANs directly.

Assuming that the LECS, LES, and BUS reside on other network devices, only LECs need to be configured on the router. Configuration begins with the ATM interface:

```
Router(config)# interface atm module/port
Router(config-if)# no ip address
Router(config-if)# atm pvc 1 0 5 qsaal
Router(config-if)# atm pvc 2 0 16 ilmi
Router(config-if)# interface atm module/port.subinterface multipoint
Router(config-if)# ip address ip-address subnet-mask
Router(config-if)# lane client ethernet elan-name
Router(config-if)# interface atm module/port.subinterface multipoint
Router(config-if)# ip address ip-address subnet-mask
Router(config-if)# lane client ethernet elan-name
Router(config-if)# ...
```

The first command line begins configuration of the ATM major interface. Notice that this interface does not receive a network address because subinterfaces provide all network access. The **atm pvc** lines configure permanent virtual circuits (PVCs) for various overhead processes. The first PVC is for the Q Signaling ATM Adaptation Layer (QSAAL)—(VPI/VCI 0/5)—which provides the signaling protocol responsible for setting up and tearing down switched virtual circuits (SVCs). The second PVC is for Integrated Local Management Interface (ILMI) (VPI/VCI 0/16), which provides communication between the ATM edge device and the ATM switches.

Then ATM subinterfaces are created one per ELAN. The subinterface number is arbitrary, but must be unique. Each subinterface is assigned a network address to participate in the ELAN. In this example, IP addresses are used. In reality, any routed protocol network address (for example, Internetwork Packet Exchange [IPX] and AppleTalk) can be used. Lastly, each subinterface is assigned to the appropriate ELAN by referencing the ELAN's name.

NOTE

For further information about configuring the remaining LANE components, see Chapter 6. The IOS commands necessary for LANE configuration can be found in the *Cisco IOS Switching Services Configuration Guide*, under “LAN Emulation”: www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/switch_c/xcprt8/index.htm.

Establishing VLAN Connectivity with Integrated Routing Processors

Route processors internal to a Catalyst switch typically have no physical interfaces to connect and configure. Rather, these modules use internal connections into the switching backplane or fabric. The type of connection is related to the specific module being used. However, the configuration is basically the same as an external router—interfaces are configured with network addresses and some reference to VLANs.

To begin configuration, the route processor module must be located and a terminal session opened. Use the **show module** command to list the installed modules. Then use the **session module-number** command to open a Telnet session to the routing module located at slot *module-number*. This command will open a session to the IOS command line running on the route processor module.

The RSM and RSFC modules in a Catalyst 5000 use virtual VLAN interfaces to provide connectivity to individual VLANs. These interfaces are configured with network addresses (IP, IPX, AppleTalk). The RSM and RSFC both make connections to the configured VLANs over their internal trunk ports to the switch backplane. Use the following commands to configure the VLAN interfaces:

```
Switch(enable) session module-number
...
Router# configure terminal
Router(config)# interface vlan vlan-number
Router(config-if)# ip address ip-address subnet-mask
Router(config-if)# no shutdown
Router(config-if)# interface vlan vlan-number
Router(config-if)# ip address ip-address subnet-mask
Router(config-if)# no shutdown
...
```

The VLAN numbers used in the **interface vlan** commands correspond to the VLAN numbers assigned to the various switch ports. After a VLAN interface is configured, it must be enabled with the **no shutdown** command. Even though the interfaces are virtual and have no physical link status, they are shut down by default after initial configuration.

Configure Routing Processes

Once connectivity has been configured between the switch and a route processor, you must also configure routing. Routes are paths to distant networks known on the local route processor, along with metrics for path costs and the addresses of next-hop route processors. In this fashion, a router hands off packets destined for a remote network to a neighboring router who is closer to the destination. Routers are used by end-user devices when the destination is not attached to the local network (VLAN).

A route processor keeps a local table of known routes, metrics, interfaces, and neighboring routers. The table entries can be derived from *static route* entries that are manually configured or from *dynamic routing protocols* that run on the router. Dynamic routing protocols

communicate with other routers running the same protocols so that optimal routes can be determined and advertised in real-time.

NOTE Routing and routing protocols are only briefly discussed in this section. For more detailed information, refer to the Interconnecting Cisco Network Devices (ICND) and Building Scalable Cisco Networks (BSCN) courses or course books, in addition to the *Cisco CCNP Routing Exam Certification Guide* from Cisco Press.

To configure dynamic routing on a route processor, use the following IOS commands:

```
Router(config)# ip routing
Router(config)# router ip-routing-protocol
Router(config-router)# network ip-network-number
Router(config-router)# network ip-network-number
...
```

The preceding sequence of commands configures IP routing. First, IP routing must be enabled with the **ip routing** command. This command is usually enabled by default. A routing protocol process must be configured with the **router** command. Several routing protocols are available for IP—Interior Gateway Routing Protocol (IGRP), Enhanced IGRP (EIGRP), Routing Information Protocol (RIP), and Open Shortest Path First (OSPF), for example. The routing protocol configured on a router must be identical to routing protocols configured on neighboring routers routes to be exchanged and determined. Some routing protocols require additional arguments to define process numbers or autonomous system numbers.

Further configuration is necessary to tell the routing protocol which networks to advertise and which interfaces should be used. This procedure is done using the **network** command. Classful (Class A, B, or C) networks are specified so that routing advertisements can be sent out concerning only these directly connected networks. As well, the routing protocol will become active on all interfaces assigned to each specified network. Advertisements will be sent and received on these interfaces.

Additional InterVLAN Routing Configurations

Once a route processor has been configured for interVLAN routing, end-user stations can use the processor. Normally, an end-user device knows only about its local subnet and can communicate only with stations on the local network or VLAN. To reach another station on a different VLAN, packets must be forwarded to a router. Therefore, each end-user device should be configured with the router's IP address on the local VLAN. This configuration is known as a *default gateway*.

In addition, a switch also needs to have a router's address configured. Consider the case when an administrator needs to establish a Telnet session with a remote switch. Perhaps the switch

has been configured to send logging or Simple Network Management Protocol (SNMP) trap information to a distant server. Unless the switch has the router's address, the switch will be unable to forward management traffic off its local management VLAN.

To configure a default gateway on an IOS-based switch, use the **ip default-gateway** *ip-address* command. On a CLI-based switch, a default route must be configured, using the following command:

```
Switch(enable) set ip route default gateway
```

Although an IP route is configured, the switch does not run or participate in a routing protocol. This is a single static route that points to a default gateway. The *gateway* field is the IP address of the router on the local management VLAN. This command can be used several times to include more routes, if necessary. The **default** keyword can be replaced by a *destination*, which is an IP network or host address to be reached through a specific gateway address.

Foundation Summary

The Foundation Summary is a collection of tables and figures that provide a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final prep before the exam, these tables and figures will hopefully be a convenient way to review the day before the exam.

Table 7-2 *InterVLAN Routing Designs*

Design	Features
Multiple Physical Links	One physical link from switch to external router per VLAN.
Trunk Link	Links between switch and external router use trunking (ISL or ATM LANE) to support multiple VLANs.
Integrated Route Processor	Internal links between integrated router and switch backplane; supports multiple VLANs.

Table 7-3 *Types of Integrated Route Processors*

Route Processor	Features
RSM	First generation Catalyst 5000 route processor (full switch module), runs IOS, features internal trunk links to switch backplane, and uses virtual VLAN interfaces.
MSFC	Second generation Catalyst 6000/6500 route processor (daughter card on Supervisor).

Table 7-4 *Catalyst Switch Commands for InterVLAN Routing Configuration*

Command	Function
show module	Lists installed switch modules to show route processor (if available).
session <i>slot-number</i>	Opens a Telnet session to the route processor in switch slot <i>slot-number</i> .
interface <i>media module/port</i> ip address <i>ip-address subnet-mask</i> no shutdown	Configures and enables physical interface (external router).
interface <i>media module/port.subinterface</i> encapsulation [isl dot1q] <i>vlan-number</i> ip address <i>ip-address subnet-mask</i>	Configures subinterfaces and VLANs on an Ethernet trunk interface (external router).

continues

Table 7-4 Catalyst Switch Commands for InterVLAN Routing Configuration (Continued)

Command	Function
interface atm <i>module/port.subinterface</i> multipoint ip address <i>ip-address subnet-mask</i> lane client ethernet <i>elan-name</i>	Configures subinterfaces, LECs, and ELANs on an ATM LANE trunk interface (external router).
interface vlan <i>vlan-number</i> ip address <i>ip-address subnet-mask</i> no shutdown	Configures and enables virtual VLAN interface on RSM/RSFC.
ip routing router <i>ip-routing-protocol</i> network <i>ip-network-number</i>	Configures routing protocol process on route processor.
set ip route default <i>gateway</i>	Configures default gateway on Catalyst switch (management traffic only).

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however, the questions are designed to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to these questions can be found in Appendix A, on page 477.

- 1 Which of the following modules performs interVLAN routing?
 - a. Catalyst 5000 Supervisor III
 - b. Catalyst 5000 RSFC
 - c. Catalyst 5000 LANE Module
 - d. Catalyst 5000 Gigabit EtherChannel Module
 - e. Catalyst 5000 RSM
- 2 Where can a router be placed in relation to switches for interVLAN routing?

- 3 How many links are needed to connect a router to four VLANs on a switch?

- 4 What types of links can be used to interconnect switches and an external router? How many VLANs can be carried on each?

5 Which Catalyst route processor module uses four internal Gigabit Ethernet links to interface with the switch backplane?

6 What trunking methods can a router support?

7 Which is better: one link per VLAN or a single trunk link supporting all VLANs? Why?

8 What is the difference between interVLAN routing and multilayer switching (MLS)?

9 What commands are needed to assign IP address 10.1.2.247 255.255.0.0 to an RSM interface for VLAN 72?

10 What Catalyst commands can be used to locate and then connect to an internal route processor?

- 11** What commands can be used to configure IP address 10.10.10.1 255.0.0.0 on an external router's VLAN 101 interface, assuming FastEthernet 0/0 is being used as a trunk link?

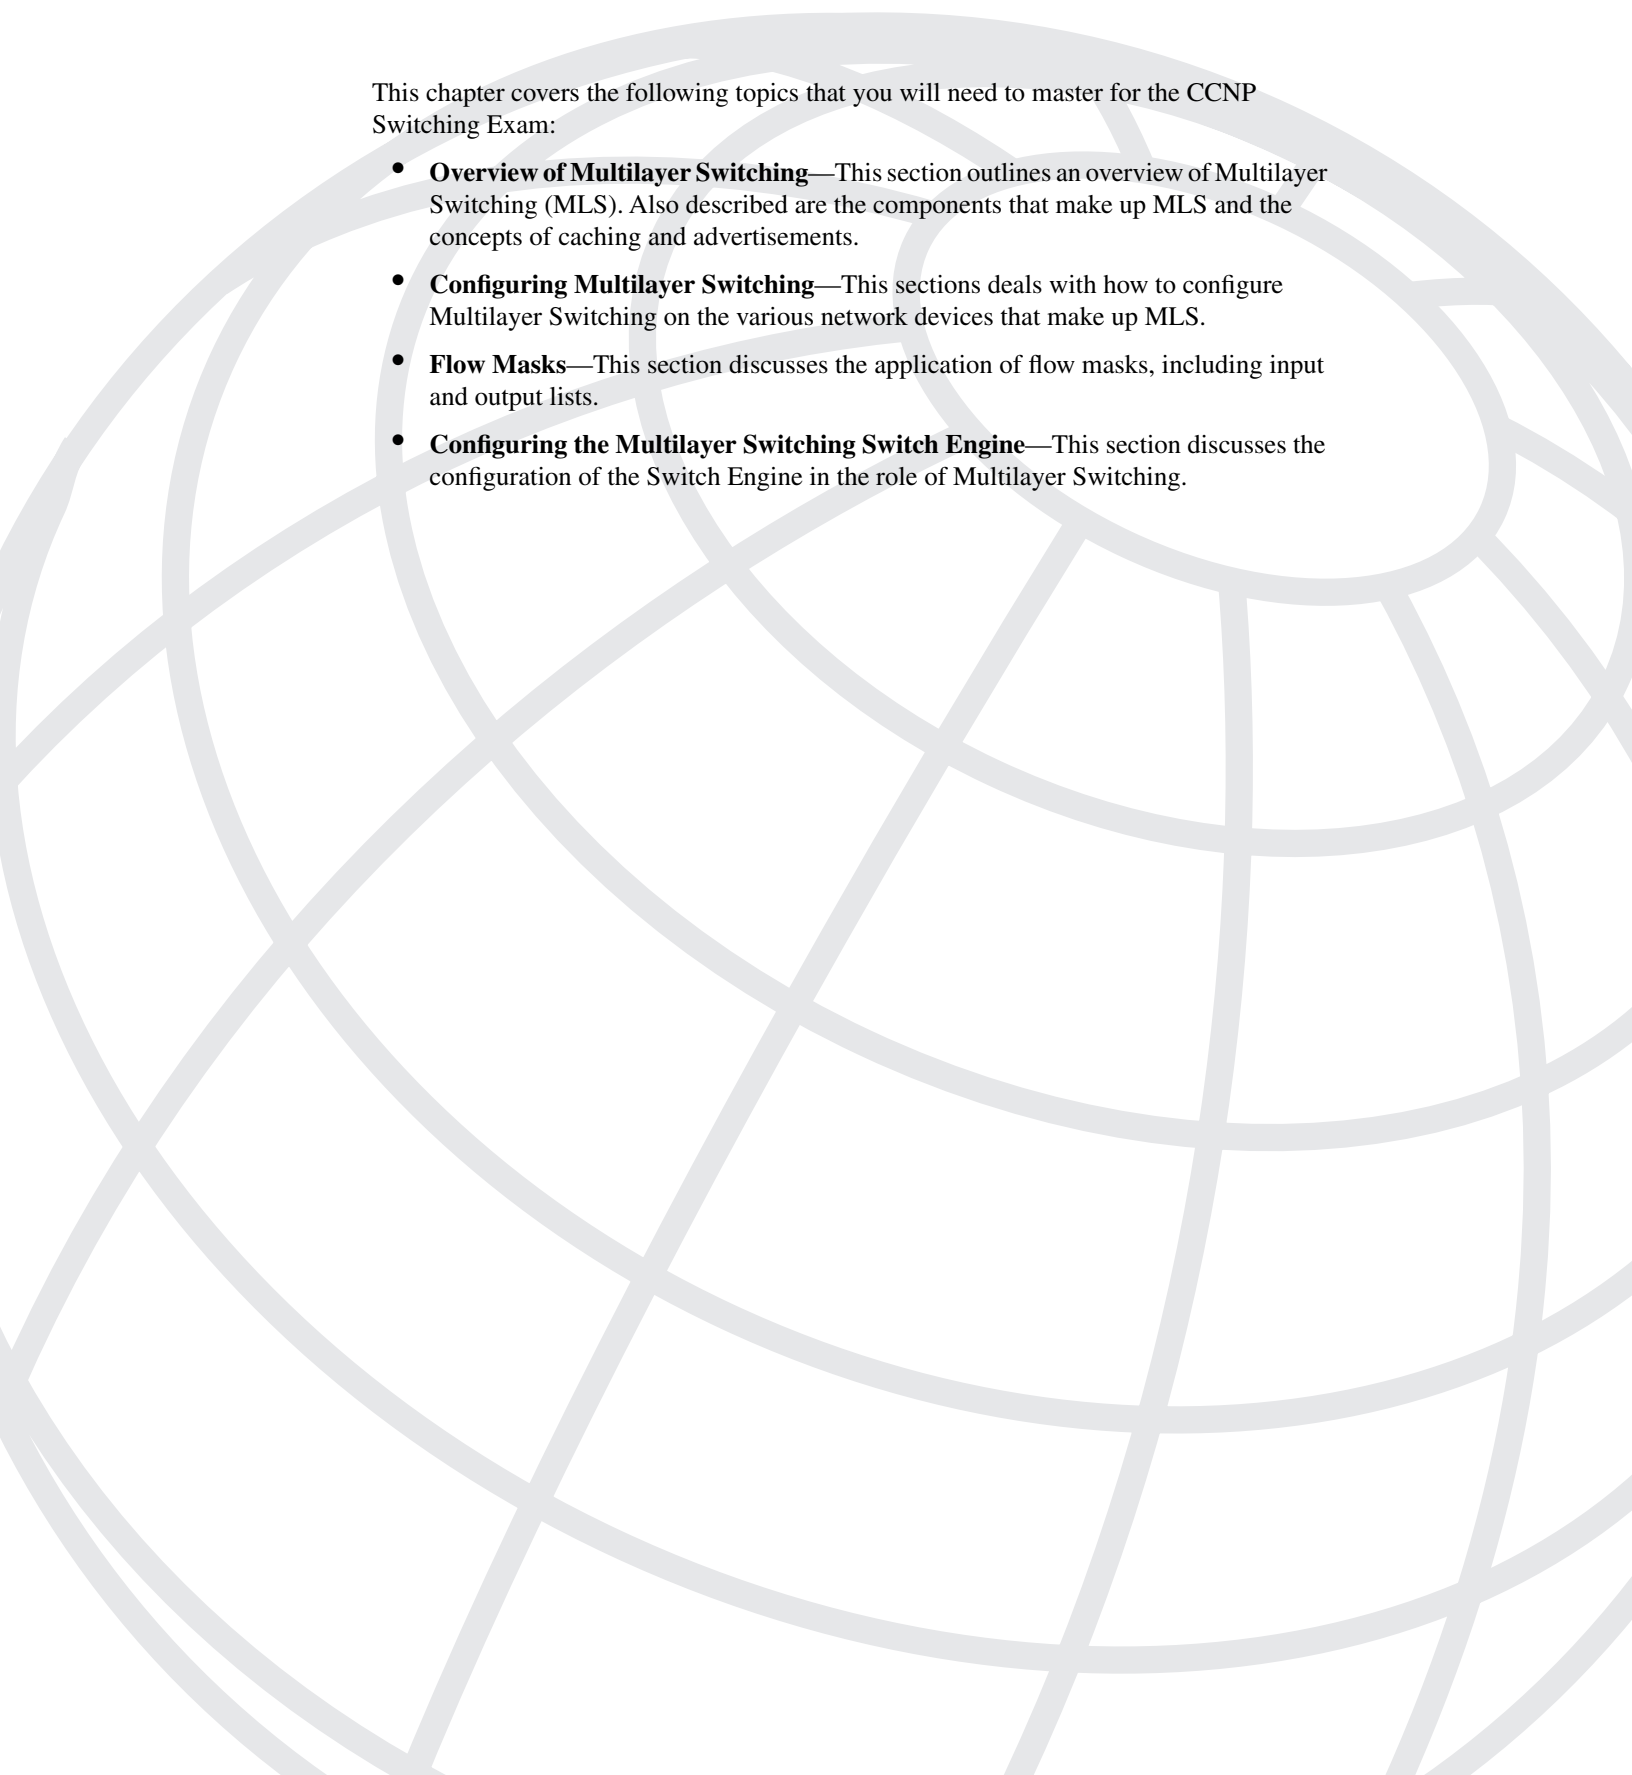

- 12** What should be configured on a route processor to dynamically determine routing paths to remote networks?

- 13** A router with an ATM LANE module is used to route traffic between six VLANs located on each of 10 Catalyst switches. Each switch has a LANE module configured for the matching VLANs and ELANs. All switches are connected by an ATM switch. What LANE components are needed on the router to support interVLAN routing? How many of each component should be configured?

- 14** If a router is used to route between VLANs, what additional information is needed so that traffic will actually be routed?

- 15** Suppose a router connects four VLANs of a switched network, providing interVLAN routing. If the router is then configured for bridging to support non-routable protocols, how would the network be affected?

- 16** A switch has its sc0 interface configured with an IP address of 10.10.1.17 255.255.0.0 on VLAN 10. A router is connected to the switched network and has interfaces at 10.10.1.1 and 10.11.1.1. What must be configured to allow someone connected to the switch console to ping host 10.11.1.100 on VLAN 11?



This chapter covers the following topics that you will need to master for the CCNP Switching Exam:

- **Overview of Multilayer Switching**—This section outlines an overview of Multilayer Switching (MLS). Also described are the components that make up MLS and the concepts of caching and advertisements.
- **Configuring Multilayer Switching**—This section deals with how to configure Multilayer Switching on the various network devices that make up MLS.
- **Flow Masks**—This section discusses the application of flow masks, including input and output lists.
- **Configuring the Multilayer Switching Switch Engine**—This section discusses the configuration of the Switch Engine in the role of Multilayer Switching.

Multilayer Switching

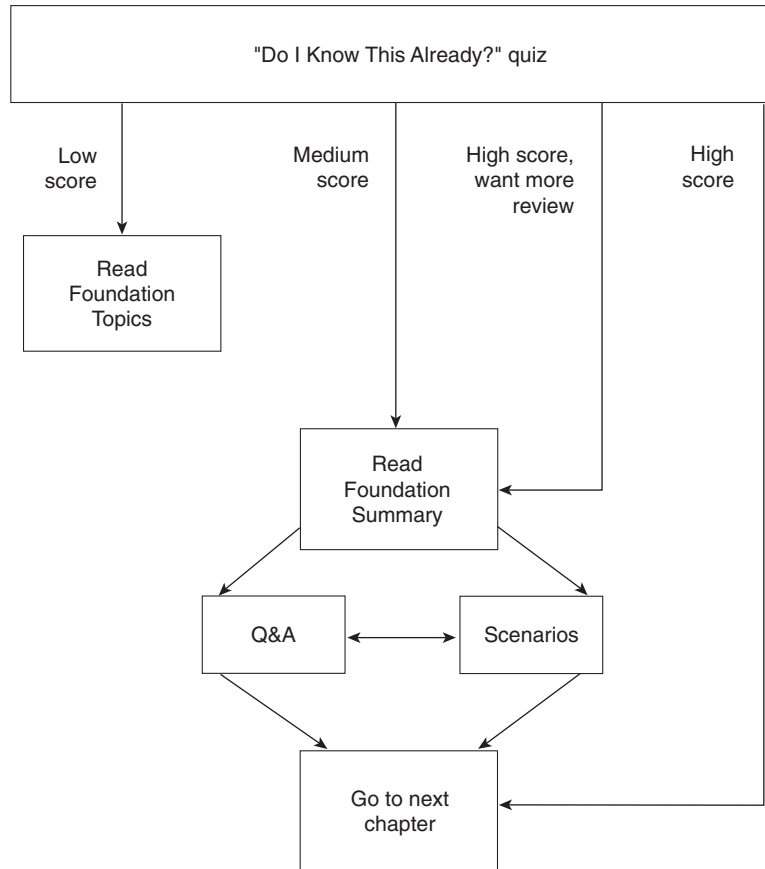
Switching technologies have matured over the years and now are a standard part of the campus network. Switching has solved a couple of problems, namely a lack of bandwidth and the inability to have disparate physical groups logically connected. Recently we've taken switching to a higher level, incorporating a routing function within the switch itself. Add some new software that allows true Layer 3 switching, and you have a recipe for success in the campus network. The performance levels are unprecedented and the ability to scale is quite different than even a few years ago. The Internet is quickly driving this industry as web server farms and their associated switching fabrics are popping up around the world.

CCNP's are expected to be able to understand what multilayer switching can do for a campus network.

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and answers for all your work with this book in one place for easy reference.
- Take the "Do I Know This Already?" quiz and write down your answers. Studies show retention is significantly increased through writing facts and concepts down, even if you never look at the information again.
- Use the diagram in Figure 8-1 to guide you to the next step.

Figure 8-1 *How To Use This Chapter*

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The ten-question quiz helps you make good choices of how to spend your limited study time. The quiz is sectioned into four areas that correspond to the four major headings in the Foundation Topics section of the chapter. Use the scoresheet in Table 8-1 to record your score.

Table 8-1 *Scoresheet for Quiz*

Quizlet Number	Foundation Topics Section Covering These Questions	Questions	Score
1	Overview of Multilayer Switching	1–3	
2	Configuring Multilayer Switching	4–6	
3	Flow Masks	7–8	
4	Configuring the MLS-SE	9–10	
All questions		1–10	

1 What devices make up the basis for Layer 3 switching as it relates in a Cisco environment?

2 What device is the definition of a Multilayer Switching Switch Engine (MLS-SE)?

3 What devices can be used as a Multilayer Switch Route Processor (MLS-RP)?

4 What is the command for enabling MLS on an RP?

5 What two things are required to make an interface on an RP MLS-enabled?

6 What command is used to verify the MLS configuration for an MLS-RP ?

7 What are the three types of flow masks modes supported on a MLS-SE?

8 What is the command to add an input access list to a MLS flow?

9 When using an external RP to a switch, is this configured automatically or manually?

10 What is the command to enable multilayer switching for a Catalyst switch?

The answers to the “Do I Know This Already?” quiz are found in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections,” on page 477. The suggested choices for your next step are as follows:

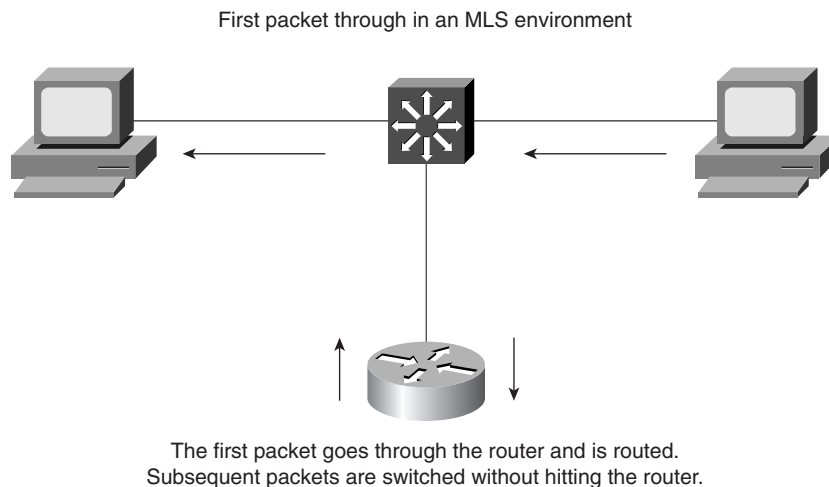
- **6 or fewer overall score**—Read the chapter. This includes the “Foundation Topics”, the “Foundation Summary”, Q&A, and scenarios at the end of the chapter.
- **7–8 overall score**—Begin with the “Foundation Summary” and then follow with the Q&A and scenarios at the end of the chapter.
- **9 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary”, and then go to the Q&A and scenarios at the end of the chapter. Otherwise, move to the next chapter.

Foundation Topics

Overview of Multilayer Switching

Catalyst switches are the basis for Layer 3 switching in the Cisco environment. Multilayer Switching (MLS) performs IP data (also IPX and IP multicast) packet flows at a much higher level of performance than traditional routing. This preserves the CPU of an upstream router without compromising functionality. Figure 8-2 shows that the first packet through enters and exits the router illustrated. Subsequent packets would be switched.

Figure 8-2 *Multilayer Switching Flow: First Packet Through*



Strictly defined, a flow is a specific conversation, consisting of multiple packets, between a network source and destination within a specific time sequence. Let's take a user that is pulling down a web page from a specific web server. This example would be one flow. The same user could be performing a File Transfer Protocol (FTP) file transfer at the same time from an FTP server. This example would be a completely different flow. Two different applications—two different protocols—two different flows; however, only one host is performing two flows. In terms of flows, there is no distinction between unicasts or multicasts.

MLS was conceived in an effort to increase the performance of a router by combining the functionality in hardware with a switch. The frame forwarding and the rewrite function is moved to hardware and then Layer 3 switching takes over the task formerly done by the router.

MLS should not be confused with NetFlow switching supported by Cisco routers. MLS uses the Route Switch Module (RSM), a directly attached external router, and the engine. With MLS,

you are not required to use NetFlow switching on the RSM or directly attached external router; any switching path on the RSM or directly attached external router will work.

MLS can be implemented by using a Layer 3 switch or an external router topology. The Layer 3 switch contains an RSM and the NetFlow Feature Card (NFFC). MLS requires the following software and hardware:

- Catalyst 2926G, 5000, or 6000 series switch with Supervisor Engine software Release 4.1(1) or later.
- Cisco IOS Release 11.3(2)WA4(4) or later.
- Supervisor Engine III or III F with the NFFC II, or Supervisor Engine II G or III G.
- Route Switch Feature Card (RSFC).
- Multilayer Switch Feature Card (MSFC).

MLS is also supported on the following software and hardware:

- Catalyst 5000 series switch with Supervisor Engine software Release 4.1(1) or later.
- Cisco IOS Release 12.0W5 or later.
- Supervisor Engine IIG or IIIG with an RSFC daughter card.

You can also implement MLS with an external router and Catalyst switch combination. The following equipment is necessary when implementing MLS with an external router and Catalyst switch combination:

- Catalyst 2926G, 5000, or 6000 series switch with Supervisor Engine software Release 4.1(1) or later.
- Supervisor Engine III or III F with the NFFC II, or Supervisor Engine II G or III G.
- Cisco high-end routers, such as Cisco 7500, 7200, 4500, 4700, or 8500 series.
- Cisco IOS Release 11.3(2)WA4(4) or later.

The connection between the external router and the switch can be multiple Ethernet links or Fast Ethernet with the Inter-Switch Link (ISL), 802.1Q, or ATM LANE.

Multilayer Switching Components

The Cisco MLS implementation includes the following components:

- **Multilayer Switching Switch Engine (MLS-SE)**—The switching entity that handles the function of moving and rewriting the packets. The MLS-SE is an NFFC residing on a Supervisor Engine III card in a Catalyst switch. It can also be a Supervisor I and the PFC on the 6000 series.

- **Multilayer Switching Route Processor (MLS-RP)**—An RSM, RSFC, MSFC, or an externally connected Cisco 7500, 7200, 4500, 4700, or 8500 series router with software that supports multilayer switching. The MLS-RP sends MLS configuration information and updates, such as the router Media Access Control (MAC) address, virtual LAN (VLAN) number flow mask, and routing and access list changes.
- **Multilayer Switching Protocol (MLSP)**—This protocol operates between the MLS-SE and MLS-RP to enable multilayer switching. MLSP is the method in which the RSM or router advertises routing changes and the VLANs or MAC addresses of the interfaces that are participating in MLS.

MLS-RP Advertisements

As soon as an MLS-RP is enabled in the campus network, MLS-RP advertisements begin. The MLS-RP sends out multicast Hello messages every 15 seconds to all switches in the network. The advertisement message consists of the following:

- The MAC addresses used by the MLS-RP on its interfaces that are participating in MLS.
- Access list information.
- Additions and deletions of routes.

MLSP uses the Cisco Group Management Protocol (CGMP) multicast address as the destination address of the Hello message. This address ensures interoperability with the Cisco switches in the network. Although this address is the same as that used by CGMP, the message contains a different protocol type so the switch can distinguish these messages from other multicast packets.

Hello Messages

All switches in the network receive the Hello message. Only Layer 3 switches actually process the message. Any switches that are not Layer 3 capable simply pass the frames through to any downstream switches.

When an MLS-SE receives the frame, the device extracts all the MAC addresses received in the frame, along with the associated interface or VLAN ID for that address. The MLS-SE records the addresses of the MLS-RPs in the MLS-SE content-addressable memory (CAM) table.

XTAGs

XTAGs are assigned by the MLS-SE to each and every MLS-RP attached to a switch. The XTAG is a one-byte value attached to the MAC address of each attached MLS-RP. These values are instrumental in differentiating between MLS-RPs when there are more than one MLS-RP available.

The XTAG is useful for deleting a specific set of Layer 3 entries from the Layer 3 table when an MLS-RP fails or exits the network.

MLS Caching

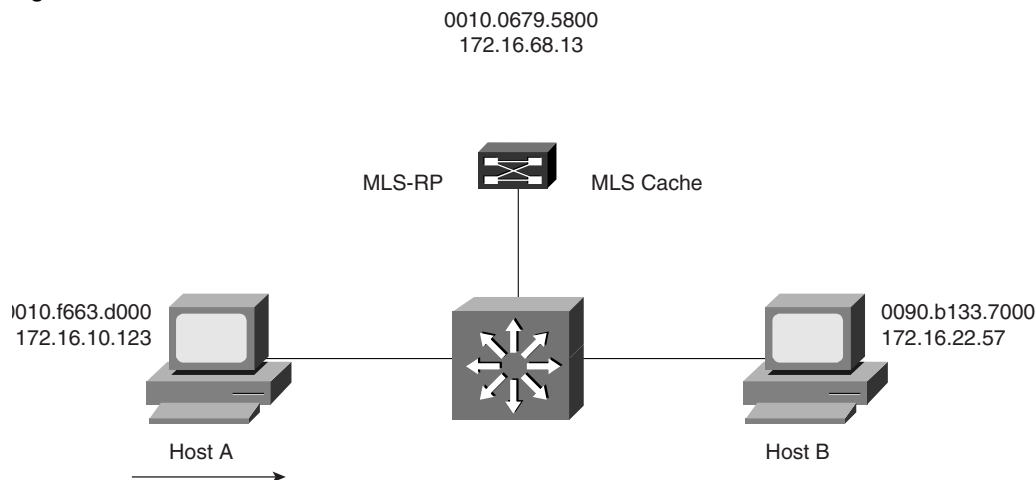
MLS caching is a process that occurs based on individual flows. In this section, we will walk through the process, step by step, in order to gain an intimate knowledge of just what occurs.

The Switching Engine (SE) is involved in the process to maintain the cache for MLS flows. Packets in a flow are compared to the cache.

Cache entries are based on one-way flows. In other words, a flow from Host A to Host B would be one flow and a flow in the reverse direction would be another flow. This action would yield two cache entries.

Here's the part of the equation that yields the payoff. In the event that the cache has an entry that is a match for the packet, the SE switches the packet instead of passing it to the router. If it does not match an entry in the cache, a process occurs that goes on to make an entry in the cache. This concept is illustrated in Figure 8-3.

Figure 8-3 *MLS Cache*



Host A sends a frame to Host B. If there is a match in the MLS cache, the packet would never go to the router but simply be switched using the sequence that follows.

- Step 1** The switch receives an incoming frame and looks at the destination MAC address in the frame.

- Step 2** The switch recognizes the destination MAC address of the frame as the address of the MLS-RP because the switch initially received this destination MAC address in a Layer 3 Hello message and programmed that destination MAC address in the CAM table.
- Step 3** The MLS-SE then checks the MLS cache to determine if an MLS flow is already established for this flow. If the frame is the first in a flow, there will not be an entry in the cache. Because the frame contained a route processor destination address, the switch recognizes the potential for Layer 3 switching for that frame.
- Step 4** On the initial packet, the switch does not have all the information for a Layer 3 switch for the frame. The switch, therefore, forwards the frame to the addressed route processor. This process of sending the frame to the addressed route processor creates a “candidate” entry in the MLS cache.
- Step 5** The route processor receives the frame and consults the routing table to determine if, in fact, the route processor has knowledge of a route for the destination address.
- Step 6** If the route processor finds the destination address in the routing table, the route processor constructs a new Layer 2 header, which now contains the route processor’s own MAC address as the source MAC address.
- The route processor also enters the MAC address of the destination host or next-hop route processor in the destination MAC address field of the Layer 2 frame.
- Step 7** The route processor then forwards the frame back to the MLS-SE.

When the switch receives the frame, the switch knows which port needs to forward the frame, based on the CAM table (displayed in Example 8-1). Moreover, the switch also recognizes the MAC address in the source field and knows that that this destination belongs to the route processor.

Example 8-1 *Displaying the CAM Table*

```

Console> (enable) show cam 00-10-29-8a-4c-00
* = Static Entry. + = Permanent Entry. # = System Entry. R = Router Entry.
VLAN Dest MAC/Route Des Destination Ports or VCs / [Protocol Type]
-----
10 00-10-29-8a-4c-00R 9/1 IP
51 00-10-29-8a-4c-00R 9/1 IP
52 00-10-29-8a-4c-00R 9/1 IP
53 00-10-29-8a-4c-00# 9/1 IP
54 00-10-29-8a-4c-00# 9/1 IP
Total Matching CAM Entries Displayed = 5
Console> (enable)

```

This recognition triggers the process of checking the MLS cache to see if there is an entry for this route processor. The switch compares the XTAGs for both the candidate entry in the MLS cache and the returned frame. If the two XTAGs match, the frame came from the same route processor for the same flow.

The switch records the information from the returned frame in the MLS cache. The switch forwards the frame out the appropriate port using the destination MAC address. This second frame becomes the “enable” entry in MLS cache and the partial entry for that flow is completed.

Remembering that the MLS-SE must see both sides of the flow going from the source to the destination in order to perform Layer 3 switching is important. In other words, you can’t do Layer 3 switching by just knowing the source or destination.

When the switch receives subsequent packets in the flow, the switch recognizes that the frames contain the MAC address of the route processor. The switch checks the MLS cache and finds the entry matching the flow in question.

The switch rewrites the Layer 2 frame header, changing the destination MAC address to the MAC address of Host B and the source MAC address to the MAC address of the MLS-RP. The Layer 3 IP addresses remain the same, but the IP header Time to Live (TTL) is decremented and the checksum is recomputed. The MLS-SE rewrites the switched Layer 3 packets so that they appear to have been routed by a route processor.

The switch rewrites the frame to look exactly as if the route processor processed the frame. The final destination sees the frame exactly as if the router processed the frame.

After the MLS-SE performs the packet rewrite, the switch forwards the rewritten frame to the destination MAC address.

The state and identity of the flow are maintained while traffic is active; when traffic for a flow ceases, the entry ages out. Partial, or candidate, entries will remain in the cache for five seconds with no enabled entry before timing out. Cache entries that are complete, where the switch captures both the candidate and the enabling packet, will remain in the cache as long as packets in that flow are detected.

Disabling MLS

Actually the title of this section should read, “What not to do if you want your MLS to keep running.” Believe it or not, there are a few commands that, if entered, will have the undesirable effect of disabling MLS.

The basic guideline to follow is that if you enter any command that forces the router to examine the packet, MLS will be disabled. That includes a whole host of commands, but I thought I’d list a few of the most common here:

- ip tcp header-compression
- no ip routing
- ip security

Configuring Multilayer Switching

The basic tasks for configuring multilayer switching include the following:

- 1 Enabling MLSP.
- 2 Assigning a VLAN ID to a route processor interface.
- 3 Adding the interfaces to the same VLAN Trunking Protocol (VTP) domain as the switch.
- 4 Enabling MLS on every interface.
- 5 Configuring the MLS Management interface.
- 6 Verifying MLS on an MLS-RP.

Before you can configure MLS for a specific VLAN or interface, you must globally enable the MLSP that operates between the route processor and the switch.

To enable MLSP on the route processor, enter the following command in global configuration mode:

```
Router(config)#mls rp ip
```

Example 8-2 states that the MLS-RP is configured to multilayer switch routed IP packets using MLSP. As of 12.0, MLS also routes Internetwork Packet Exchange (IPX) packets.

Example 8-2 *Determining the MLS-RP Is Configured*

```
Router#show run
Building configuration...
Current configuration:
!
version 11.3
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Router
!
!
mls rp ip
!
```

To disable MLS on the route processor/RSM, enter the **no mls rp ip** command in global configuration mode.

In Cisco's MLS implementation, Layer 3 switches IP IPX, and IP multicast packets. Any other packets are routed as in a non-Layer 3 switched network.

MLS is interVLAN routing. Multilayer switches make forwarding decisions based upon which ports are configured for which VLANs. Internal route processors and ISL-configured links inherently use VLAN IDs to identify interfaces. External route processor interfaces have

knowledge regarding subnets but not VLANs. Therefore, MLS requires that each external route processor interface have a VLAN ID assigned to it.

To assign a VLAN ID to a route processor interface, enter the following commands in interface configuration mode:

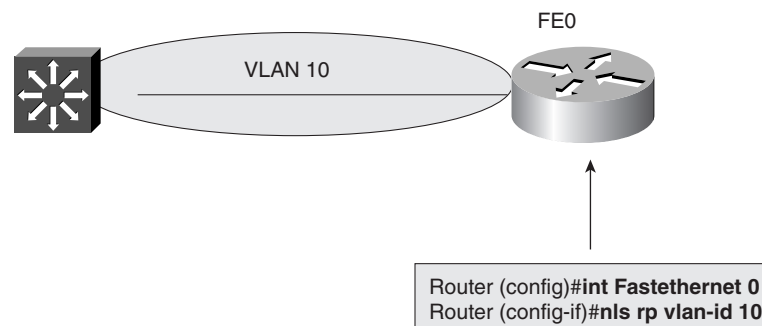
```
Router (config)#interface interface number
Router (config-if)#mls rp vlan-id vlan-id-num
```

where *vlan-id-num* represents the VLAN assigned to this interface.

To remove an interface from a VLAN, enter the **no mls rp vlan-id** *vlan-id-num* command.

Removing the VLAN ID from an interface disables MLS for that interface. Figure 8-4 demonstrates how to use these commands to assign a VLAN ID to a route processor interface.

Figure 8-4 Assigning a VLAN ID



After you determine which route processor interfaces will be MLS interfaces, you must add the interfaces to the same VTP domain as the switch. Both the switch and the MLS interfaces must be in the same domain. If the switch is not assigned to a VTP domain, you do not need to perform this task.

To place an external route processor interface in the same VTP domain as the switch, enter the following commands in interface configuration mode:

```
Router(config) interface interface number
Router(config-if)# mls rp vtp-domain domain-name
```

where *domain-name* is the name of the VTP domain in which the switch resides.

For an ISL interface, enter the **mls rp vtp-domain** command only on the primary interface. All subinterfaces that are part of the primary interface inherit the VTP domain of the primary interface.

The running configuration in Example 8-3 states that the VLAN41 interface of the MLS-RP is configured to reside in the Rigel2 VTP domain.

Example 8-3 *Determining the VTP Domain of the MLS-RP VLAN Interface*

```
Router#show run
Building configuration...
(Text deleted)
mls rp ip
!
!
interface Vlan1
 ip address 172.16.1.168 255.255.255.0
!
interface Vlan41
 ip address 172.16.41.168 255.255.255.0
 mls rp vtp-domain Rigel2
```

To remove the MLS interface from a VTP domain, enter the **no mls rp vtp-domain domain-name** command.

Displaying VTP Domain Information

Sometimes seeing VTP domain information is useful. The **show mls rp vtp-domain** command allows you to see domain information for a specific VTP domain:

```
Router#show mls rp vtp-domain vtp domain name
```

The display resulting from this command (see Example 8-4) shows a subset of the **show mls rp** command display. The following information is a result of issuing the **show mls rp vtp-domain** command:

- The name of the VTP domain(s) in which the MLS-RP interfaces reside.
- Statistical information for each VTP domain.
- The number of management interfaces defined for the MLS-RP.
- The number of VLANs in this domain configured for MLS.
- The ID of each VLAN configured for this domain MAC address.
- The number of MLS-SEs of which the router or RSM has knowledge of in this domain.
- The MAC address of each switch in this domain.

Example 8-4 *Displaying VTP Domain Information*

```
router# show mls rp vtp-domain WBU

vlan domain name: WBU
  current flow mask: ip-flow
  current sequence number: 80709115
  current/maximum retry count: 0/10
  current domain state: no-change
  current/next global purge: false/false
  current/next purge count: 0/0
  domain uptime: 13:07:36
  keepalive timer expires in 8 seconds
  retry timer not running
  change timer not running
  fcp subblock count = 7

  1 management interface(s) currently defined:
    vlan 1 on Vlan1

  7 mac-vlan(s) configured for multi-layer switching:

    mac 00e0.fefc.6000
      vlan id(s)
      1   10  91  92  93  95  100

  router currently aware of following 1 switch(es):
    switch id 0010.1192.b5ff
```

Enabling MLS

MLS is enabled on a per-interface basis. Just because you put an interface into a particular VTP domain doesn't mean that you've activated MLS. MLS must be enabled on every interface that you desire to participate in Layer 3 switching.

On a router or RSM interface, enter the following command in interface configuration mode in order to enable MLS:

```
Router (config-if)#mls rp ip
```

The running configuration in Example 8-5 shows that the VLAN19 interface of the MLS-RP is enabled to participate in MLS.

To disable MLS on an interface, enter the **no mls rp ip** command.

Example 8-5 *Determining that the MLS-RP VLAN Interface is Enabled for Multilayer Switching*

```
Router#show run
Building configuration...
(Text Deleted)
mls rp ip
!
!
interface Vlan1
 ip address 172.16.1.168 255.255.255.0
!
interface Vlan19
 ip address 172.16.41.168 255.255.255.0
 mls rp vtp-domain san-fran
 mls rp ip
```

VTP Domain Issues

When a route processor resides in a VTP domain other than the domain in which the switch resides, the switch cannot multilayer switch frames for that router. There are several ways in which a route processor and switch can end up in different VTP domains as follows:

- You can purposely place both devices in separate domains.
- You can misname or mistype the VTP domain when configuring either the switch or route processor.
- You can enter the MLS interface command prior to putting the interface in a VTP domain.

Configuring an interface for MLS by assigning the interface to a VTP domain prior to assigning it to a VTP domain places that interface in the null domain. When the interface resides in a null domain, it cannot participate in MLS with the switch.

To remove the MLS interface from a null VTP domain, disable MLS on the interface.

MLS Management Interface

When a RSM or router is configured to participate in MLS, the device uses the MLSP to send Hello messages, advertise routing changes, and announce the VLANs or MAC addresses of those interfaces on the devices participating in MLS. One interface on the MLS-RP must be identified as the management interface through which MLSP packets are sent and received. The MLSP management interface can be any MLS interface connected to the switch.

Only one management interface needs to be specified. If no management interface is configured, however, MLSP messages will not be sent.

Multiple interfaces on the same route processor can be configured as a management interface; however, this action increases the management overhead per route processor. Cisco does not recommend this practice.

To identify a management interface on an RSM or router, enter the following command in interface configuration mode:

```
Router(config-if)#mls rp management-interface
```

To disable the management interface, enter the **no mls rp management-interface** command in interface configuration mode.

The running configuration in Example 8-6 states that the VLAN41 interface on the MLS-RP is configured as the management interface.

Example 8-6 *Determining if the MLS-RP VLAN Interface Is Configured as the Management Interface*

```
Router#show run
Building configuration...

(Text Deleted)
mls rp ip
!
!
interface Vlan1
 ip address 172.16.1.168 255.255.255.0
!
interface Vlan41
 ip address 172.16.41.168 255.255.255.0
 mls rp vtp-domain bcmsn
 mls rp management-interface
 mls rp ip
```

Verifying MLS-RP

To verify the MLS configuration for an MLS-RP, enter the following command in privileged EXEC mode:

```
Router#show mls rp
```

The display resulting from this command (see Example 8-7) shows the following information:

- Whether MLS is globally enabled or disabled.
- The MLS ID for this MLS-RP.
- The MLS IP address for this MLS-RP.
- The MLS flow mask.
- The name of the VTP domain(s) in which the MLS-RP interfaces reside.
- Statistical information for each VTP domain.

- The number of management interfaces defined for the MLS-RP.
- The number of VLANs configured for MLS.
- The ID of each VLAN configured for this MAC address.
- The number of MLS-SEs to which the router or RSM is connected.
- The MAC address of each switch.

Example 8-7 *Displaying MLS RP Information*

```

router# show mls rp

multilayer switching is globally enabled
mls id is 00e0.fefc.6000
mls ip address 10.20.26.64
mls flow mask is ip-flow
vlan domain name: WBU
  current flow mask: ip-flow
  current sequence number: 80709115
  current/maximum retry count: 0/10
  current domain state: no-change
  current/next global purge: false/false
  current/next purge count: 0/0
  domain uptime: 13:03:19
  keepalive timer expires in 9 seconds
  retry timer not running
  change timer not running
  fcp subblock count = 7

1 management interface(s) currently defined:
  vlan 1 on Vlan1

7 mac-vlan(s) configured for multi-layer switching:

  mac 00e0.fefc.6000
  vlan id(s)
  1   10  91  92  93  95  100

router currently aware of following 1 switch(es):
  switch id 0010.1192.b5ff

```

Each MLSP-RP is identified to the switch by both the MLS ID and MLS IP address of the route processor. The MLS ID is the MAC address of the route processor. The MLS-RP automatically selects the IP address of one of its interfaces and uses that IP address as its MLS IP address.

The MLS-SE uses the MLS ID as a determining factor for establishing entries in the MLS cache.

This MLS IP address is used in the following situations:

- By the MLS-RP and the MLS-SE when sending MLS statistics to a data collection application.
- In the included MLS route processor list on the switch.

To verify the MLS configuration for a specific interface, enter the following command in privilege EXEC mode:

```
Router#show mls rp interface interface number
```

The display resulting from this command shows the following information:

- Whether MLS is configured on the interface.
- The VTP domain in which the VLAN ID resides.
- Whether this interface is configured as the management interface for the MLS-RP.

If the interface is not configured for MLS, the **show mls rp ip** command displays the following message:

```
Router#show mls rp ip interface Vlan41  
mls not configured on Vlan41
```

Flow Masks

The MLS-SE uses flow mask modes to determine how packets are compared to MLS entries in the MLS cache. The flow mask mode is based on the access lists configured on the MLS router interfaces. The MLS-SE learns the flow mask through MLSP messages from each MLS-RP for which the MLS-SE is performing Layer 3 switching.

MLS-SE supports only one flow mask for all MLS-RPs that are serviced by the MLS-SE. If the MLS-SE detects different flow masks from different MLS-RPs for which the MLS-SE is performing Layer 3 switching, the MLS-SE changes its flow mask to the most specific flow mask detected. However, if a more specific flowmask is in effect, a less specific flow mask then is applied.

The MLS-SE supports three flow mask modes as follows:

- **Destination-IP**—The default flow mask mode, Destination-IP represents the least-specific flow mask. The MLS-SE maintains one MLS entry for each destination IP address. All flows to a given destination IP address use this MLS entry. This mode is used if no access lists are configured on any of the MLS router interfaces.
- **Source-Destination-IP**—The MLS-SE maintains one MLS entry for each source and destination IP address pair. All flows between a given source and destination use this MLS entry regardless of the IP protocol ports. This mode is used if a standard access list is on any of the MLS interfaces.

- **IP-Flow**—This mode represents the most specific flow mask. The MLS-SE creates and maintains a separate MLS cache entry for every IP flow. An IP-Flow entry includes the source IP address, destination IP address, protocol, and protocol ports. This mode is used if there is an extended access list on any MLS interface.

When the MLS-SE flow mask changes, the entire MLS cache is purged.

You can set a flow mask on the MLS-SE without applying an access list on the route processor. You use the **set mls flow** command when you want to cache entries on a specific set of criteria to export flow statistics but not to set an access list on an interface. To set the flow mask on the MLS-SE without setting an access list on a route processor interface, enter the following command in privilege mode:


```
set mls flow [destination | destination-source | full].
```

The keywords **destination** means that you are applying the IP-Destination mode, **destination-source** means that you are applying Source-Destination-IP mode, and **full** means that you are applying IP-Flow mode. These different modes were explained earlier.

Output Lists

Figure 8-5 illustrates an output access list applied to the interface. In this case, the MLS-SE learns of this change through the MLSP process and then enforces security for the flow. Enforcement of the access list would purge any entries for flows on that interface from the MLS cache.

Figure 8-5 *Output Access Lists*



Any new flows would then be created based on the restrictions imposed by the access list. The next packet in the flow becomes a candidate packet and the process of establishing a MLS cache entry is initiated.

New entries are placed in the MLS cache once the initial packet in the flow passes the test conditions in the output access control list (ACL).

Using options like **log**, **reflexive**, or **established** forces the router to examine every packet before routing. Under MLS, the router does not examine every packet; therefore, these options are not allowed.

Input Access Lists

As with output access lists, placing an input access list on an MLS-enabled interface purges the MLS cache of all existing flows for that interface.

Because the default behavior for the input access list is to examine and route all incoming packets, however, all subsequent packets in the flow between Hosts A and B are routed.

Most input access lists can be implemented as output access lists to achieve the same effect.

Routers configured with Cisco IOS Release 11.3 or later will not automatically support input access lists on an interface configured for MLS. If an interface is configured with an input access list, all packets for a flow that are destined for that interface go through the router. Even if the router allows that flow, the flow is not Layer 3 switched.

To enable MLS to cooperate with input access lists, enter the following command in global configuration mode:

```
Router(config)#mls rp ip input-acl
```

The running configuration in Example 8-8 states that input ACLs on the MLS-RP are configured to work in a MLS environment.

Example 8-8 *Determining if Input Access Lists on the MLS-RP Can Operate in an MLS Environment*

```
Router#show run
Building configuration...

Current configuration:
!
version 11.3
(Text Deleted)
mls rp nde-address 172.16.31.113
mls rp ip input-acl
mls rp ip
```

To remove support for input access lists in an MLS environment, enter the **no mls rp ip input-acl** command in global configuration mode.

Configuring the MLS-SE

This section deals with topics involved in configuration of the switching engine or MLS-SE. Topics covered include enabling MLS, MLS caching, verifying MLS, external router support, and switch inclusion lists.

MLS is enabled by default on Catalyst series switches that support Layer 3 switching—in other words, if an RSM is on the switch. There are, however, a couple of cases where configuring the switch is necessary. The first is obvious, when the MLS-RP happens to be an external router. Because an external router is not an integral part of the switch, no knowledge of Layer 3 switching exists. The other case is when the aging time of MLS cache entries is different than the default, hence, requiring some configuration to change this parameter.

In the event that a switch has been disabled for Layer 3 switching, enter the following command in privilege EXEC mode on the switch to re-enable it:

```
Switch(enable)#set mls enable
```

The running configuration in Example 8-9 shows the entry that shows the MLS-SE is configured to support MLS.

Example 8-9 *Determining if the MLS-SE Is Configured to Support MLS*

```
Switch(enable)#show config
(Text Deleted)
#mls
set mls enable
```

Enter the **set mls disable** command to disable MLS on the MLS-SE. This command stops the MLS-SE from processing the MLSP messages from the MLS-RP and purges all existing MLS cache entries in the switch.

MLS Caching

Because the MLS cache has a size limitation, MLS entries will be deleted from the cache if certain conditions are met. This deletion, or aging, process takes into effect for the following reasons:

- Candidate entries remain in the cache for five seconds with no enabled entry before timing out.
- An MLS entry is deleted from the cache if a flow for that entry has not been detected for the specified aging time. The default aging time is 256 seconds.
- Other events, such as applying access lists, routing changes, or disabling MLS on the switch, can cause MLS entries to be purged.

The amount of time an MLS entry remains in the cache is user modifiable. To alter the value of the aging time, enter the following command in privileged EXEC mode:

```
Switch(enable)#set mls agingtime agingtime
```

where *agingtime* is the amount of time an entry remains in the cache before the entry is deleted. The range of the aging time value is from 8 to 2032 seconds. The default value is 256 seconds.

The running configuration in Example 8-10 states that entries in which no packets have been detected for a period of six minutes will be deleted from the cache.

Example 8-10 *Configuring Cache Aging*

```
Switch(enable)#show config
(Text Deleted)
#mls
set mls enable
set mls agingtime 272
```

The values for *agingtime* are entered in eight-second increments. Any *agingtime* value that is not a multiple of eight seconds is adjusted to the closest one.

Some MLS flows are sporadic or short-lived. An example of a sporadic or short-lived flow would be packets that are sent to or received from a Domain Name System (DNS) or Trivial File Transfer Protocol (TFTP) server. Because the connection may be closed after one request and one reply cycle, that MLS entry in the cache is used only once. However, that MLS entry still consumes valuable cache space until the entry is aged out. Detecting and aging out these entries quickly can save MLS entry space for real data traffic.

To solve the problem of short-lived entries in the cache, a different type of aging mechanism, called fast aging, is available. This type of aging states that if the MLS-SE does not detect a specified number of packets in a certain time period, then that entry is removed from the cache.

To configure the fast aging option, enter the following command in privilege EXEC mode:

```
Switch(enable)# set mls agingtime fast fastagingtime pkt_threshold
```

where *fastagingtime* indicates the amount of time an entry remains in the cache before the entry is deleted. Allowable configuration values are 32, 64, 96, or 128 seconds. The default is 0 seconds.

The *pkt_threshold* argument indicates the number of packets that must be detected within the specified amount of time. Allowable configuration values are 0, 1, 3, 7, 15, 31 or 63 packets. The default is 0 packets.

In the configuration in Example 8-11, we have configured a *fastagingtime* of 96 and a *pkt_threshold* of 15. So for this example, any cache entries in which no more than 15 packets have been detected for a period of 96 seconds will be deleted from the cache.

Example 8-11 *Determining Entries to Be Deleted from the Cache*

```
Switch (enable)show config
(Text Deleted)
#mls
set mls enable
set mls agingtime 272
set mls agingtime fast 96 15
```

Verifying MLS Configurations

To display information about MLS on a MLS-SE, enter the following command in privileged EXEC mode:

```
Switch (enable) show mls
```

The following information is displayed as result of executing the above command (see Example 8-12):

- Status of MLS.
- Aging time, in seconds, for an MLS cache entry.
- Fast aging time, in seconds, and the packet threshold for a flow.
- Flow mask.
- Total packets switched.
- Number of active MLS entries in the cache.
- Whether Netflow data export is enabled and, if so, for which port and host.
- MLS-RP IP address, MAC address, XTAG, and supported VLANs.

Example 8-12 *Displaying Information about MLS on an MLS-SE*

```
Switch (enable) show mls

Multilayer switching enabled
Multilayer switching aging time = 110 seconds
Multilayer switching fast aging time = 64 seconds, packet threshold = 7
Full flow
Total packets switched = 87128
Active shortcuts = 1298
Netflow Data Export disabled
Netflow Data Export port/host is not configured.
Total packets exported = 0

MLS-RP IP          MLS-RP ID      XTAG      MLS-RP MAC-Vlans
-----
192.168.1.127     0010f6fe12a3  28        00-10-f6-fe-12-a3 1,21-22
```

If you want to display information about a specific MLS-RP, enter the **show mls rp** command and designate the IP address of the target MLS-RP.

where you execute the command does make a difference. You can execute this command on both the MLS-SE and the MLS-RP. In this case, we are talking about the MLS-SE

External Router Support

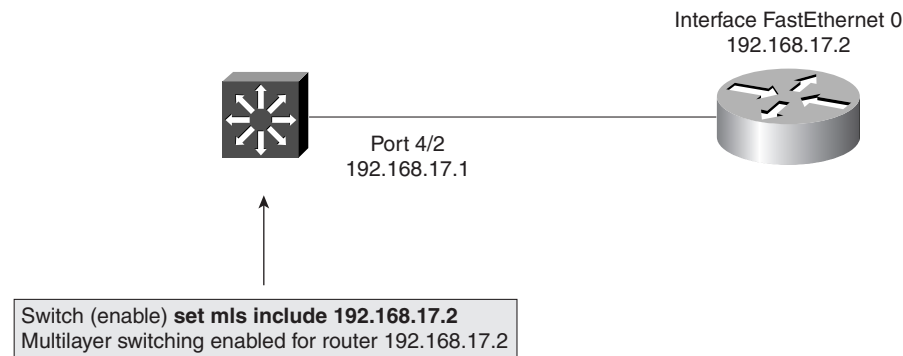
If the switch supports an externally attached MLS-RP, the switch must be manually configured to recognize that MLS-RP. To manually include an external MLS-RP, enter the following command in privilege EXEC mode on the switch:

```
Switch (enable) set mls include ip-addr
```

where *ip-addr* is the MLS IP address of the external router. To determine the IP address of the MLS-RP, enter the **show mls rp** command on the MLS-RP.

Perform this command *only* for external routers. The MLS-SE automatically includes the IP address of co-resident RSMs in the switch inclusion list. When the RSM is physically removed from the switch chassis or MLS is disabled on an RSM, the RSM IP address is removed from the inclusion list. The auto-included RSM cannot be cleared using the **clear mls include** command. Figure 8-6 demonstrates implementing the **set mls include** command to support MLS for external routers.

Figure 8-6 Including External Routers



The running configuration in Example 8-13 states that an external MLS-RP with the IP address of 172.16.41.168 has been added to the MLS include list.

To remove the MLS-RP from the switch inclusion list, enter the **clear mls include** command. A single MLS-RP can be removed by entering the IP address of a specific MLS-RP. All externally connected MLS-RPs can be removed from the switch inclusion list by entering the **clear mls include all** command.

Example 8-13 *Including External Routers in Multilayer Switching*

```
Switch (enable)#show config
(Text Deleted)
#mls
set mls enable
set mls agingtime 256
set mls agingtime fast 0 0
set mls include 172.16.41.168
```

Switch Inclusion Lists

To display the contents of the switch inclusion list to determine which MLS-RPs are participating in MLS with the MLS-SE, enter the following command in privilege EXEC mode:

```
Switch (enable) show mls include
```

The resulting display returns the IP addresses of *all* MLS-RPs that are participating in MLS with the MLS-SE.

If the IP address of an MLS-RP does not appear in the switch inclusion list, the MLS-SE will not perform Layer 3 switching for the MLS-RP. If the MLS-SE is supposed to be performing Layer 3 switching for a specific router and its IP address is not listed in the inclusion list, check the following:

- Is the router for which you manually entered the MLS IP address external?
- If the router is an RSM, is there an RSM resident and is it functional?
- Is MLS globally enabled on the MLS-RP?

Displaying MLS Cache Entries

To display the MLS cache entries, enter the following command in privilege EXEC mode:

```
Switch (enable) show mls entry.
```

This command might be used as a troubleshooting tool or just to check the status of a particular flow that you're interested in.

This command can be further defined to show MLS cache entries for the parameters defined in Table 8-2.

To remove entries from the MLS cache, enter the **clear mls entry** command in privilege EXEC mode. Table 8-3 lists how to remove MLS cache entries based on given criteria.

Table 8-2 *Displaying Specific MLS Cache Entries*

MLS Cache Entry Based On	Command to Use
Specific destination IP address	show mls entry destination <i>ip-address</i>
Specific source IP address	show mls entry source <i>ip-address</i>
Specific MLS_RP ID	show mls entry rp <i>ip-address</i>
Specific IP flow	show mls entry flow <i>protocol source-port destination-port</i>

Table 8-3 *Removing MLS Cache Entries*

Remove MLS Cache Entry Based On	Command to Use
Specific source IP address	clear mls entry source <i>ip-address</i>
Specific destination IP address	clear mls entry destination <i>ip-address</i>
Specific flow	clear mls entry flow <i>protocol src-port dst-port</i>

Refer to the “Configuring Multilayer Switching” section of the *Catalyst Series Switch Configuration Guide (4.3)*, available online at www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_4_3/config/mls.htm#41001 for details on how to format this command for each of the above instances.

Foundation Summary

The Foundation Summary is a collection of tables and figures that provide a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final prep before the exam, these tables and figures will hopefully be a convenient way to review the day before the exam.

Table 8-4 *Components of Multilayer Switching*

Component	Description
Multilayer Switching Switch Engine (MLS-SE)	The MLS-SE is a NetFlow Feature Card residing on a Supervisor Engine III card in a Catalyst switch. It can also be a Supervisor I and the PFC on the 6000 series.
Multilayer Switching Route Processor (MLS-RP)	An RSM, RSFC, MSFC or an externally connected Cisco 7500, 7200, 4500, 4700, or 8500 series router with software that supports multilayer switching.
Multilayer Switching Protocol (MLSP)	This protocol operates between the MLS-SE and MLS-RP to enable multilayer switching.

Table 8-5 *MLS Router Commands*

Command	Description
access-list <i>access-list-number</i>	Creates an access list.
ip access-group <i>access-list-number</i>	Assigns an access list to an interface.
mls rp input-acl	Supports the creation of MLS flow entries from interfaces with input ACLs.
mls rp ip	Enables multilayer switching on an MLS-RP and on a specific interface.
mls rp management-interface	Establishes a management interface through which MLSP messages are sent.
mls rp vtp-domain <i>vtp-domain-name</i>	Assigns an interface to a VTP domain.
show mls rp	Displays the MLS configuration on the MLS-RP.
show run	Displays the current configuration on the router.

Table 8-6 *MLS Switch Commands*

Command	Description
set mls agingtime <i>seconds</i>	Alters the time in which MLS entries are maintained in the MLS cache.
set mls enable	Enables multilayer switching on the MLS-SE.
show mls	Displays the MLS configuration on the MLS-SE.
show mls include	Displays the switch MLS-RP inclusion list.
show mls entry	Displays the MLS cache.
show mls rp	Displays the MLS configuration on the MLS-RP.

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however, they are designed to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess. If you get an answer wrong, review the appropriate section of this chapter to make sure you understand the reason for your mistake.

The answers to these questions can be found in Appendix A, on page 477.

- 1 What devices are the basis for Layer 3 switching as it relates in a Cisco environment?

- 2 What device is the definition of a Multilayer Switch Engine (MLS-SE)?

- 3 What devices can be used as a Multilayer Switch Route Processor (MLS-RP)?

- 4 What is the command for enabling MLS on an RP?

- 5 What two things are required to make an interface on an RP MLS-enabled?

6 What command is used to verify the MLS configuration for an MLS-RP ?

7 What are the three types of flow masks modes supported on a MLS-SE?

8 What is the command to add an input access list to a MLS flow?

9 When using an external RP to a switch, is this configured automatically or manually?

10 What is the command to enable Multilayer Switching for a Catalyst switch?

11 Assuming that MLS is running, what effect does the command **clear ip route** do on an MLS-RP?

12 What three components are required in a Cisco implementation of MLS?

13 Define a Destination-IP flow mask.

14 What is the command to display MLS entries in the cache?

Scenarios

Scenario 8-1

Refer to Figure 8-7, which depicts a simple router and switch setup for this scenario.

Figure 8-7 Scenario 8-1 Network Setup



We've decided that we need to support MLS on these two devices due to performance issues.

- 1 What commands would be necessary to implement MLS on these two devices?
- 2 Assume we are going to use the Interface VLAN12 on the RP. Also, the domain is called SJC-1. Configure accordingly.
- 3 Interface VLAN12 is also the management interface. Activate this feature.
- 4 We need to activate an input access list for VLAN12. Configure this accordingly.
- 5 On the MLS-SE, we want the MLS cache to timeout after 224 seconds. Configure this on the switch.
- 6 The RP pictured is to be included and has an IP address of 172.16.48.113. Configure accordingly.
- 7 Type the command to display included RPs.

Scenario 8-2

Refer to the output in Example 8-14 and 8-15 from **show** commands on a Catalyst switch acting as an MLS-SE, and then answer the questions that follow.

Example 8-14 Scenario 8-2 show mls Command Output

```
Switch (enable) show mls
Multilayer switching enabled
Multilayer switching aging time = 192 seconds
Multilayer switching fast aging time = 56 seconds, packet threshold = 12
Full flow
Total packets switched = 81391
Active shortcuts = 1115
Netflow Data Export disabled
Netflow Data Export port/host is not configured.
Total packets exported = 0
```

MLS-RP IP	MLS-RP ID	XTAG	MLS-RP MAC-Vlans
172.16.30.15	0010f6ad4cb2	28	00-10-f6-ad-4c-b2 1,4-5

Example 8-15 Scenario 8-2 show mls include Command Output

```
Switch (enable) show mls include
Included MLS-RP
-----
172.16.30.15
```

- 1 Use the output from Example 8-14 and 8-15 to generate a configuration of the switch as it relates to MLS.
- 2 How many VLANs are involved in MLS? What are they?
- 3 What is the XTAG for the MLS-RP?
- 4 Is the MLS-RP an RSM or an external attached router?
- 5 What type of flow is being used here?

Scenarios Answers

Scenario 8-1 Answers

- 1 To configure MLS on the RP, the command is **mls rp ip** while in global configuration mode. On the SE, in enable mode, the command is **set mls enable**.
- 2 Under the interface VLAN12, enter the command **mls rp vtp-domain sjc-1**.
- 3 Again, under the interface VLAN12, enter the command **mls rp management-domain**.
- 4 Also, under the interface VLAN12, enter the command **mls rp ip input-acl**.
- 5 On the switch, in enable mode, enter the command **set mls agingtime 224**.
- 6 On the switch, in enable mode, enter the command **set mls include 172.16.48.113**.
- 7 In order to display included RPs, enter the following command on the switch: **show mls include**.

Router Configuration for Scenario 8-1

```
Router#show run
Building configuration...

(Text Deleted)
mls rp ip
!
!
interface Vlan1
 ip address 172.16.1.1 255.255.255.0
!
interface Vlan12
 ip address 172.16.48.113 255.255.255.0
 mls rp vtp-domain sjc-1
 mls rp management-interface
 mls rp ip input-acl
 mls rp ip
```

Switch Configuration for Scenario 8-1

```
Switch (enable)show config
(Text Deleted)
#mls
set mls enable
set mls agingtime 224
set mls agingtime fast 96 15
set mls include 172.16.48.113
```

Display for show mls include Command (Question 7)

```
Switch (enable) show mls include
Included MLS-RP
-----
172.16.48.113
```

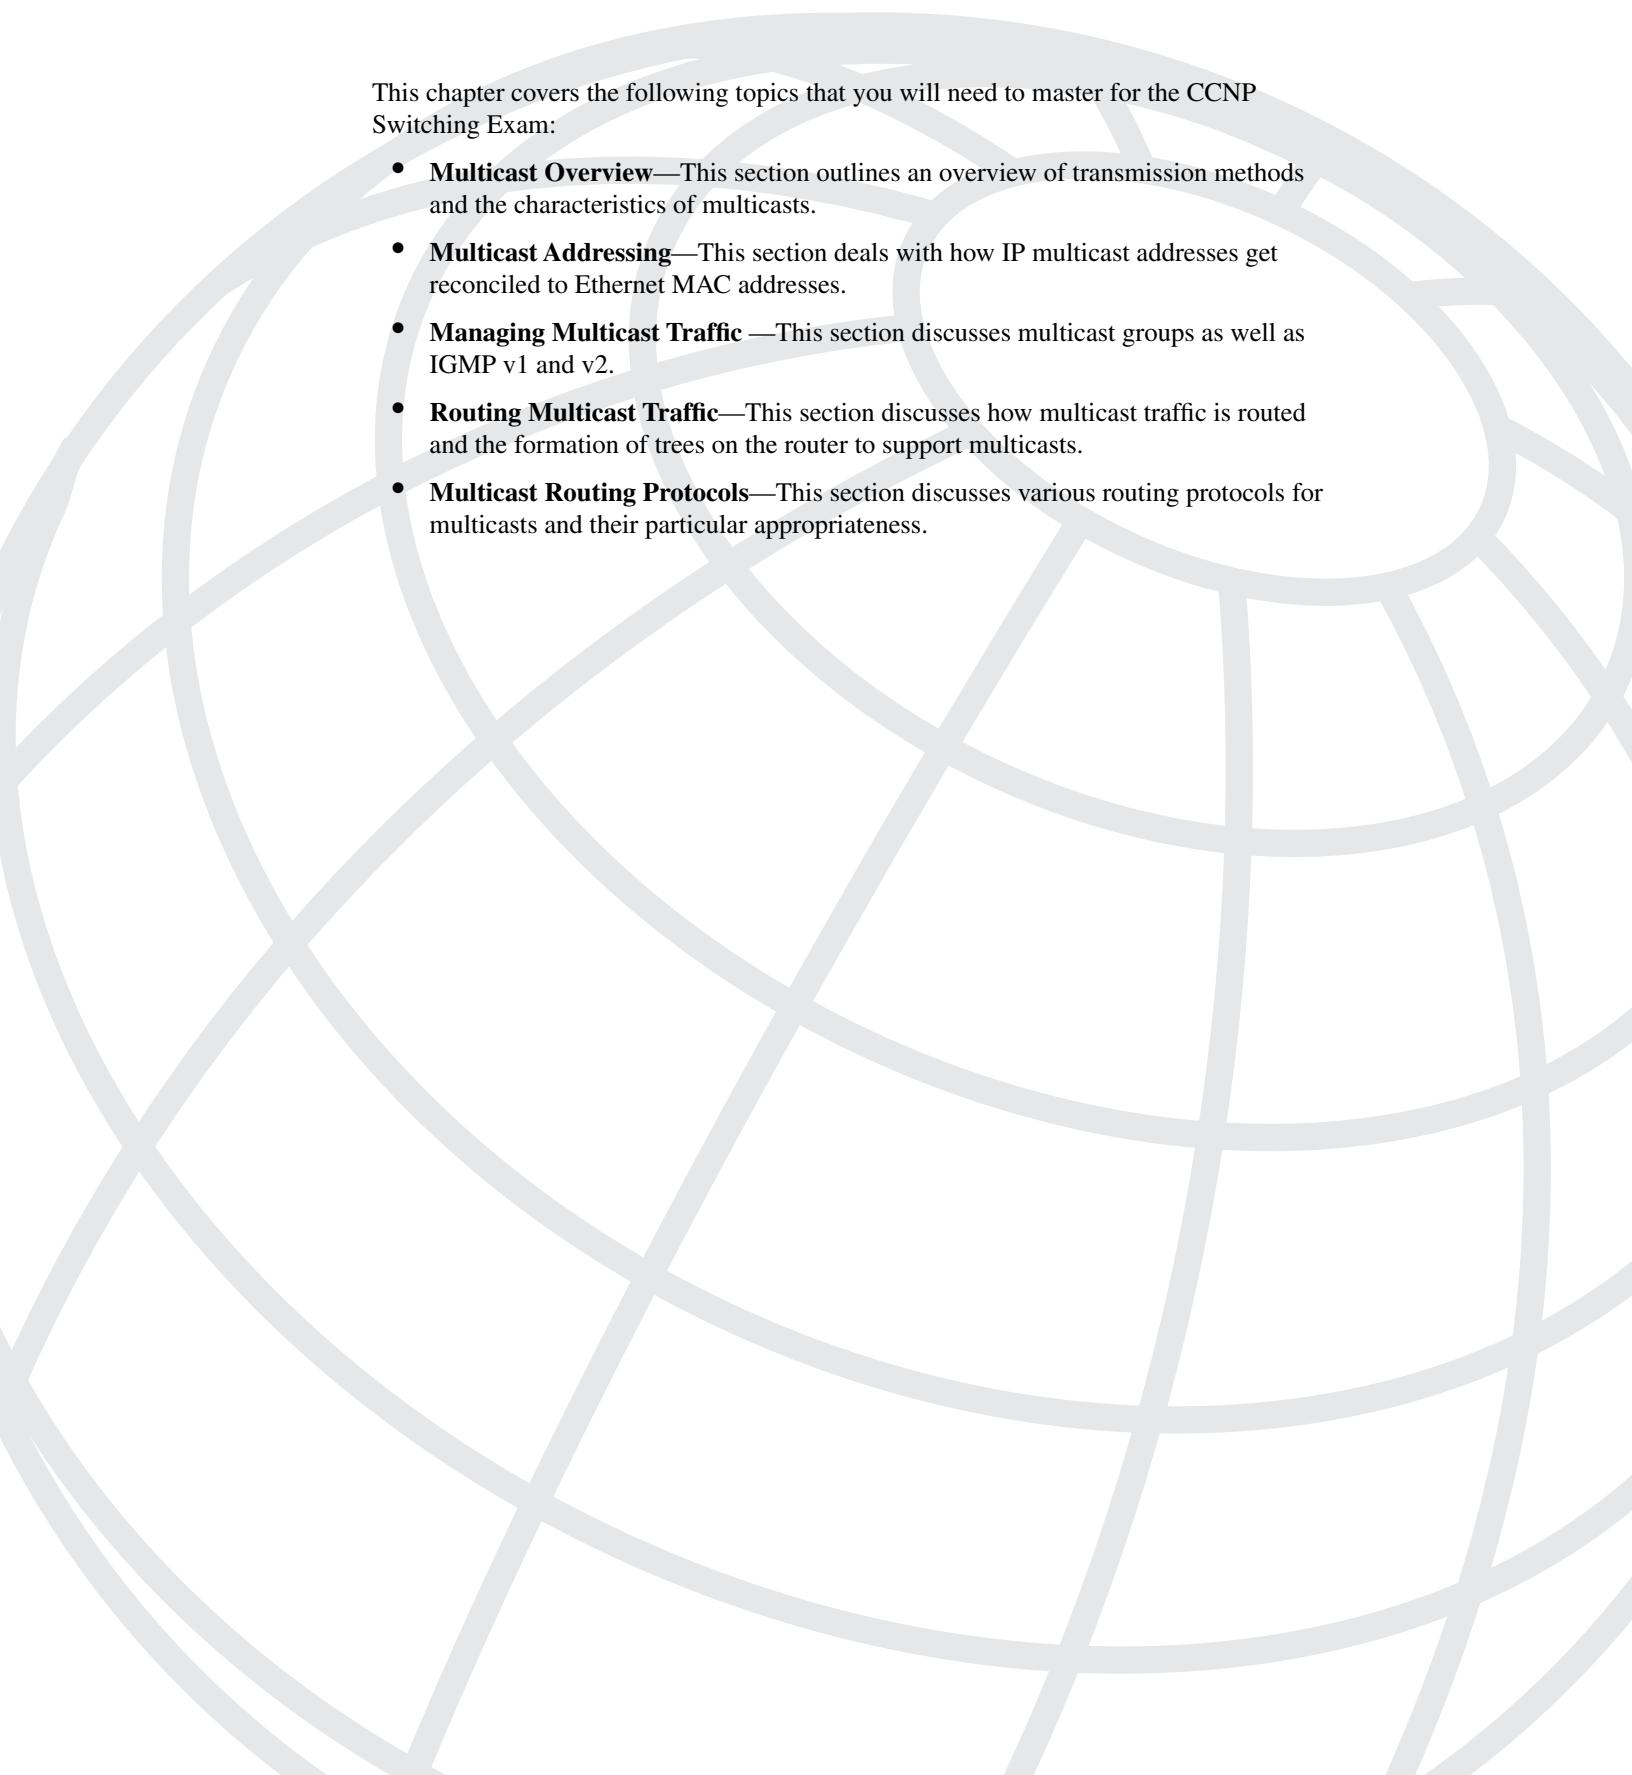

Scenario 8-2 Answers

- 1 Example 8-16 shows the correct configuration for Scenario 8-2.

Example 8-16 Scenario 8-2 Configuration

```
Switch (enable) show config
(Text Deleted)
#mls
set mls enable
set mls agingtime 192
set mls agingtime fast 56 12
set mls include 172.16.30.15
```

- 2 There are three VLANs and they are VLAN 1, VLAN 4, and VLAN5.
- 3 The XTAG for the RP is 28.
- 4 Because there is an included router, this is the sign that the RP is an external router, rather than an RSM.
- 5 This is a full flow.



This chapter covers the following topics that you will need to master for the CCNP Switching Exam:

- **HSRP Overview**—This section outlines an overview of Hot Standby Router Protocol (HSRP). Also described are the issues and nuances of using HSRP in a switched network using both traditional routers and virtual routers.
- **HSRP Operations**—This section deals with router interaction in an HSRP standby group.
- **Configuring HSRP**—This section deals with how to configure HSRP on the various devices that make up the switch block and how to ensure fault tolerant design.

Overview of Hot Standby Router Protocol

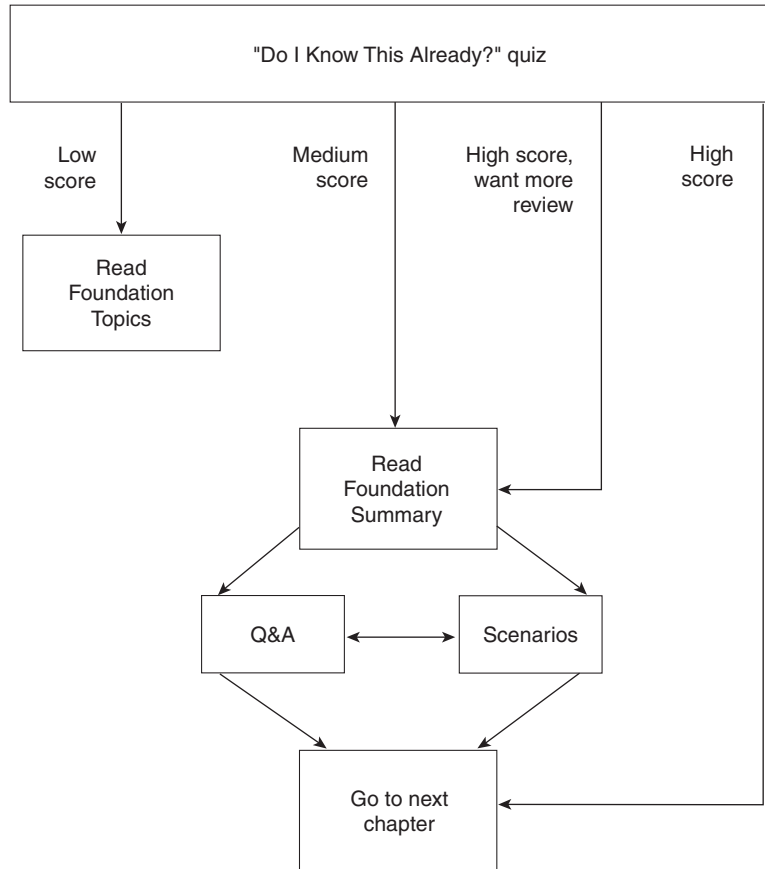
Hot Standby Router Protocol (HSRP) was conceived as a method of providing a level of fault tolerance in the network. HSRP is a Cisco proprietary protocol that is outlined in RFC 2281 (www.isi.edu/in-notes/rfc2281.txt). The HSRP protocol protects against a failure of the first-hop router. HSRP picks up where the default router left off. You can have a routing protocol capable of discovering dynamic routes, but when it comes to the default route on hosts, there isn't a means to change this in the event of a failure. So, HSRP was developed to rectify the situation. HSRP uses a monitoring function to determine the status of primary and standby router interfaces. If you configure multiple HSRP groups, you can have a backup and also do load sharing across disparate networks.

So why have a whole chapter on HSRP? As a CCNP, you'll be expected to know how to implement redundant architectures and provide load sharing and backup capabilities in today's enterprise networks.

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and answers for all your work with this book in one place for easy reference.
- Take the “Do I Know This Already?” quiz and write down your answers. Studies show retention is significantly increased through writing facts and concepts down, even if you never look at the information again.
- Use the diagram in Figure 9-1 to guide you to the next step.

Figure 9-1 *How To Use This Chapter*

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The 14-question quiz helps you make good choices of how to spend your limited study time. The quiz is sectioned into four smaller “quizlets,” which correspond to the four major headings in the “Foundation Topics” section of the chapter. Use the scoresheet in Table 9-1 to record your score.

Table 9-1 *Scoresheet for Quiz and Quizlets*

Quizlet Number	Foundation Topics Section Covering These Questions	Questions	Score
1	HSRP Overview	1–5	
2	HSRP Operations	8, 10, 11	
3	Configuring HSRP	6, 7, 9, 12–14	
All questions		1-14	

1 What is the name of the protocol that allows a set of routers that are working together to form one virtual router?

2 What is the minimum number of routers needed to perform HSRP?

3 In a properly functioning virtual router, what happens when the active router fails?

4 How many standby groups can exist on any one LAN?

5 Name the six states that an HSRP configured router can be in.

6 When configuring HSRP on a particular router interface, if the standby group is not explicitly configured, what standby group does the interface fall into by default?

7 What command is used to display the HSRP virtual router IP and MAC address?

8 Which router in an HSRP group becomes the forwarding router and how is it determined?

9 In the command **standby 35 priority 90**, what does the “35” stand for?

10 An HSRP router exchanges Hello messages with other HSRP routers. What is contained in the hello message?

11 What does the term *tracking* imply in an HSRP environment?

12 What command would allow you to debug HSRP?

13 What does the **preempt** command do in the HSRP environment?

14 What command or commands enable the preempt role in an HSRP-enabled network?

The answers to the “Do I Know This Already?” quiz are found in Appendix A on page 477. The suggested choices for your next step are as follows:

- **8 or fewer overall score**—Read the chapter. This includes the “Foundation Topics,” the “Foundation Summary,” Q&A, and scenarios at the end of the chapter.
- **2 or less on any “quizlet”**—Review the subsection(s) of the “Foundation Topics” part of this chapter based on Table 9-1. Then move into the “Foundation Summary” quiz and scenarios at the end of the chapter.
- **9–12 overall score**—Begin with the “Foundation Summary” and then follow with the Q&A and scenarios at the end of the chapter.
- **13 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary.” Then go to the Q&A and scenarios at the end of the chapter. Otherwise, move to the next chapter.

Foundation Topics

HSRP Overview

This chapter discusses the use of HSRP in the campus network. The campus network is important because that is where the switching platform is most prevalent. To have the campus network available at all times, a level of redundancy must be built in. This redundancy is where HSRP comes in. The campus model has a hierarchical look to it and HSRP fits in at the switch block layer, providing an active and a standby path to the distribution layer. HSRP not only allows for a layer of redundancy, but also has the capability of providing load sharing.

Issues with Traditional Methods

This section deals with various traditional methods that have been used, but for one reason or another are ineffective in some failure modes. These methods include default gateways, proxy ARP, RIP, and IRDP.


Default Gateways

In Figure 9-2, Router A is responsible for routing packets for Subnet A, and Router B is responsible for handling packets on Subnet B. If Router A goes down or otherwise becomes unavailable to the PC, the routing protocols used between the rest of the routers will converge at some point and connectivity will remain for at least part of the network.

Unfortunately, the PC doesn't have the capability to collect and exchange routing information. Therefore, when Router A went down, it doesn't know about any other way to route information because it is limited to a single default router. These devices typically are configured with a single default gateway IP address. If the router that is the default gateway fails, the device is limited to communicating only on the local IP network segment and is effectively disconnected from the rest of the network. Even if a redundant router exists that could serve as a default gateway, no dynamic method exists for these devices to switch to a new default gateway IP address.

Providing redundant paths and redundant equipment in key locations of the network to prevent the failure of any single device does not guarantee the availability of vital network resources to the end user.

In this example, the end-user station is configured to send packets for the Application Server to Router 192.168.20.1. Router A and Router B know about all of each others IP interfaces through the use of dynamic routing protocols. However, no way exists for the PC to dynamically switch to a new default gateway if the present default gateway fails.

Figure 9-2 *Default Gateway Illustration*

Proxy ARP

Proxy ARP is sometimes used in a network. Proxy ARP is used when the router is responding with its own MAC address as a proxy for some other host on a different subnet. So what does that buy you? Because the router knows where the destination MAC is located or via the routing table, sending packets to the router will indeed get you to the real destination.

In a proxy ARP situation, the host considers the destination to be on the same subnet as itself. The rub comes when the router being used as a proxy fails. This creates a problem in that the host station will continue to send packets destined for the far end station. The packets will unfortunately be dropped.

How is this problem rectified? A reboot of the host computer will fix the problem by clearing the ARP cache, or a force of another ARP request will pick up a failover router. Unfortunately, there is a significant delay in communications, either to reboot the host or to initiate an ARP. This delay depends entirely on the ARP timeout period.

Routing Information Protocol (RIP)

One of the first routing protocols available in networking was the Routing Information Protocol (RIP) to discover routers. In this case, the workstation holds a routing table, which lists routes and associated next hops that have a path to the destination. It is then up to the workstation to choose the best path.

RIP is a distance vector routing protocol, which means it relies on the number of hops. RIP has a limitation of 15 hops. The other limitation of using RIP—and essentially any distance vector protocols—is that network convergence is slow. In an unstable network, this can cause many problems.

ICMP Router Discovery Protocol (IRDP)

Some newer IP hosts use the ICMP Router Discovery Protocol (IRDP) to find a new router when a route becomes unavailable.

A host that runs IRDP listens for hello multicast messages from the preferred default router. As long as the end-user station detects these hello messages, the MAC address for the router generating the hello messages is used as the destination MAC address by the end-user station. As soon as an end user no longer detects the hello messages from the preferred router, the end-user station switches to an alternate router.

IRDP is preferred over RIP and default gateways because RIP takes longer to converge. Default gateways do not provide an alternative route if the default gateway becomes unavailable.

Some IP hosts use ICMP and IRDP to find a new path when the primary router becomes unavailable. IRDP is not a routing protocol like RIP or the Interior Gateway Routing Protocol (IGRP). IRDP is an extension to ICMP that provides a mechanism for routers to advertise useful default routes. IRDP offers several advantages over other methods of discovering addresses of neighboring routers. IRDP does not require hosts to recognize routing protocols, nor does it require manual configuration by an administrator.

A host that uses IRDP listens for hello multicast messages from the preferred default router. The IRDP-based advertisements are considered valid only for a predefined lifetime value. If a new advertisement is not seen during that lifetime, the router address is considered invalid and the host removes the corresponding default route. The IRDP protocol allows for varying timing values. A lifetime value is included in the header of every IRDP advertisement and applies to all addresses included in the packet. A host will use the router address only for the number of lifetime seconds after the most recent advertisement.

Advertisements are sent every seven to ten minutes; the default lifetime is 30 minutes. However, the router has complete control over the interval and lifetime values and thus can control the period of time the addresses are considered valid.

IRDP has two separate interval times: a minimum and a maximum advertisement interval. All unsolicited advertisements are sent in the window of time defined by these two values. IRDP is covered in greater detail in RFC 1256.

Hot Standby Router Protocol

So what problem are we trying to solve? We've looked at solutions that try to solve the problem of a network failure, but for various reasons don't necessarily fix the problem. Enter Hot Standby Router Protocol. HSRP addresses the problem caused by first-hop failures generally having static default gateway addresses on hosts. Previously, a failure at the default gateway address would leave the host unable to communicate outside of its own subnet. Now with HSRP, the default gateway is a virtual-router address. A failure of the active router would result in a switch to the standby router, and packets would continue to be forwarded.

Cisco routers use HSRP, which enables end stations to continue communicating throughout the network even when the default gateway becomes unavailable.

With HSRP, a set of routers works together to represent a single virtual standby router. The standby router group functions as a single router configured with a virtual IP and MAC address, distinct from the physical routers in the network.

Because the routers in the standby group route packets sent to a virtual address, packets are still routed through the network even when the router originally forwarding the packets fails.

HSRP allows one router to automatically assume the function of the second router if the second router fails. HSRP is particularly useful when the users on one subnet require continuous access to resources in the network.

If the primary or lead router of a group of HSRP routers fails, a standby router in the same group begins to forward traffic for the HSRP group.

The routers decide within the group which router forwards traffic for the virtual address. At regular intervals, the routers exchange information to determine which routers are still present and able to forward traffic.

When routers are configured to be part of an HSRP group, the routers recognize their own native MAC address, plus the HSRP group MAC address.

Routers whose Ethernet controllers only recognize a single MAC address will use the HSRP MAC address when performing as the active router and the burn-in address (BIA) when in standby mode or not speaking.

NOTE


The router also sends a gratuitous ARP when it becomes active in order to make the end stations aware of the MAC address change.

HSRP Group Members

As shown in Figure 9-3, the HSRP group consists of the following members:

- Active router
- Standby router
- Virtual router
- Other routers

Figure 9-3 *HSRP Group Members*



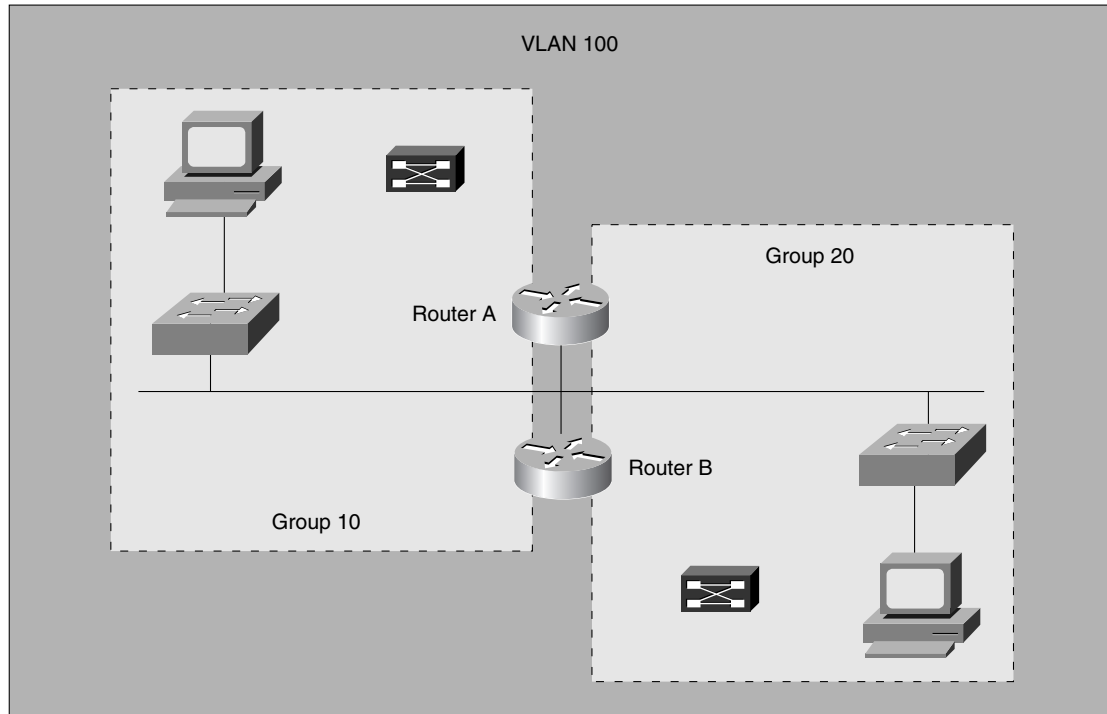
Each of these members will be discussed in detail later in this chapter.

To facilitate load sharing, a single router may be a member of multiple HSRP standby groups on a single segment. Each standby group emulates a single virtual router. There is a limit of 255 standby groups on any given LAN.

NOTE

Some platforms do not support multiple HSRPs because of the single MAC address per interface restriction. This restriction can be lifted with the use of the **standby use-bia** command.

Figure 9-4 illustrates that both Router A and Router B are members of Groups 10 and 20. However, Router A is the active forwarding router for Group 10 and the standby router for Group 20. Router B is the active forwarding router for Group 20 and the standby router for Group 10. Both Groups are members of VLAN 100.

Figure 9-4 HSRP Group Members Example

Addressing HSRP Groups Across ISL Links

As stated earlier, HSRP routers can provide for redundancy and load sharing across the same subnets, but what about different subnets?

As you are no doubt well aware, an ISL trunk provides the transit for multiple VLANs simultaneously.

NOTE HSRP over ISL was introduced in IOS 11.3 and is not possible with any earlier IOS version.

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 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.

Running HSRP over ISL allows users to configure redundancy between multiple routers that are configured as front ends for VLAN IP subnets. By configuring HSRP over ISLs, users can eliminate situations in which a single point of failure causes traffic interruptions. This feature inherently provides some improvement in overall networking resilience by providing load balancing and redundancy capabilities between subnets and VLANs.

To configure HSRP over an ISL link between VLANs, perform the following tasks:

- Step 1** Define the encapsulation type.
- Step 2** Configure the IP address.
- Step 3** Enable HSRP.

The first two steps were discussed in Chapter 7, “InterVLAN Routing.” The steps to enable HSRP are discussed later in this chapter.

CAUTION Although a route processor can theoretically support up to 32,650 subinterfaces, practical usage is dictated by the robustness of the route processor and the number of VLANs. Please monitor the CPU utilization of the route processor to determine your practical limitations.

Multiple HSRP Groups

Routers can belong to multiple groups within multiple VLANs. As members of multiple HSRP groups, routers can simultaneously provide redundant backup and perform load sharing across different IP subnets. Some of the other characteristics of multiple HSRP groups include the following:

- Although multiple routers can exist in an HSRP group, only the active router forwards the packets sent to the virtual router.
- A separate HSRP group is configured for each separate VLAN.
- Multiple standby groups may coexist on a LAN segment. Each group operates independently of other groups. Each standby group emulates a single virtual router.
- Individual routers may participate in multiple groups. The router maintains separate state and timers for each group.
- For each standby group, an IP address and a single well-known MAC address with a unique group identifier is allocated to the group.
- There can be up to 255 standby groups on any LAN.
- If multiple groups are used on a single LAN, load splitting can be achieved by distributing hosts among different standby groups.

HSRP Operations

This section deals with router interaction in an HSRP standby group and discusses the concept of the active router, which is responsible for the forwarding of packets. Also discussed is the ARP process, as well as the anatomy of an HSRP message and the different states of HSRP.

Active Router

One router in each group is elected to be the active router. The election process occurs through the sending and receiving of hello messages. The hello message contains a priority level for the sending router. The router with the highest standby priority in the group becomes the active router. The active router forwards the packets sent to the virtual router.

If the priority level is the same for each router in the group, the first router to come up and obtain the virtual router IP address becomes the active router. If the priorities are equal, the highest IP address wins.

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

Within the standby group, one router is elected to be the active router. The router with the highest standby priority in the group becomes the active router. The default priority for an HSRP router is 100. This option is user-configurable.

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.

Locating the Virtual Router MAC Address

The ARP process makes an association between Layer 3 network addresses and Layer 2 hardware addresses. An example would be an IP address and hardware Ethernet address. Each router maintains a table of resolved addresses. The router checks this ARP cache before attempting to contact a device to determine if the address has already been resolved.

The IP address and corresponding MAC address of the virtual router is maintained in the ARP table of each router in an HSRP standby group.

Example 9-1 shows the output from the **show ip arp** command (entered in privileged EXEC mode), which displays the ARP cache on a router.

Example 9-1 *show ip arp Command Output Displays ARP Cache on a Router*

Protocol	Address	Age (min)	Hardware Addr	Type	Interface
Internet	172.16.10.82	-	0010.f6b3.d000	ARPA	Vlan10
Internet	172.16.10.169	-	0010.0b79.5800	ARPA	Vlan10
Internet	172.16.10.110		0000.0c07.ac0a	ARPA	Vlan10

Note the following in Example 9-1:

- The virtual router IP address is 172.16.10.110.
- The virtual router MAC address is the well-known MAC address 0000.0c07.ac and the HSRP group identifier is 0a.
- The interface is VLAN10 on an RSM, which is the VLAN being routed.

Active and Standby Router Behavior

As mentioned earlier in this chapter, each HSRP group contains the following entities:

- An active router
- A standby router
- A virtual router

The function of the active router is to forward packets sent to the virtual router. Another router in the group is elected as the standby router. The active router assumes and maintains its active role through the transmission of hello messages.

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

The function of the virtual router is to present a consistently available router to the end user.

An HSRP standby group may contain other routers. These routers monitor the hello messages but do not respond to ARP requests. These routers do forward any packets addressed to the routers' IP addresses, but do not forward packets for the virtual router.

When the active router fails, the other HSRP routers stop receiving hello messages (they must miss three hellos for this to happen), and the standby router assumes the role of the active router.

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

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

Anatomy of an HSRP Message

All routers in a standby group send or receive HSRP messages. These messages are used to determine and maintain the router roles within the group. HSRP messages are encapsulated in the data portion of User Datagram Protocol (UDP) packets and use port number 1985. These packets are addressed to an “all router” multicast address with a Time to Live (TTL) of one (1). Figure 9-5 shows the general format for an HSRP message.

Figure 9-5 HSRP Message Format

1 Octet	1 Octet	1 Octet	1 Octet
Version	Op Code	State	Hellotime
Holdtime	Priority	Group	Reserved
Authentication Data			
Authentication Data			
Virtual IP Address			

The HSRP message contains the following information:

- The Version field indicates the version of the HSRP.
- The Op Code describes the type of message contained in this packet.
- Hello messages are sent to indicate that a router is running and capable of becoming either the active or standby router.
- Group messages are sent when a router wants to become the active router.
- Reserved messages are sent when a router no longer wants to be the active router.
- Internally, each router in the standby group implements a state machine. The State field describes the current state of the router sending the message.
- The Hellotime field is only meaningful in hello messages. This field contains the approximate period between the hello messages that the router sends. The time is given in seconds.
- The Holdtime field is only meaningful in hello messages. This field contains the amount of time that the current hello message should be valid. The time is given in seconds.

- The Priority field is used to elect the active and standby routers. When comparing priorities of two different routers, the router with the numerically higher priority wins. In the case of routers with equal priority, the router with the higher IP address wins.
- The Group field identifies the standby group. Values in the range of 0 and 255 are valid.
- The Authentication Data field contains a clear-text, eight-character reused password.
- The Virtual IP Address field contains the IP address of the virtual router used by this group.

Only the active and the standby routers send periodic HSRP messages after the protocol has completed the election process.

HSRP States

HSRP defines six states in which an HSRP configured router may exist. When a router exists in one of these states, the router performs the necessary actions required in that state. The HSRP states are as follows:

- **Initial state**—All routers begin in the initial state. This starting state indicates that HSRP is not running. This state is entered via a configuration change or when an interface first comes up.
- **Learn state**—In the learn state, the router is still waiting to hear from the active router. The router has not yet seen a hello message from the active router, nor learned the IP address of the virtual router.
- **Listen state**—In the listen state, the router knows the virtual IP address, but is neither the active router nor the standby router. The router listens for hello messages from those routers.
- **Speak state**—In the speak state, the router sends periodic hello messages and is actively participating in the election of the active and/or standby router. A router cannot enter the speak state unless the router has the IP address of the virtual router.
- **Standby state**—In the standby state, the router is a candidate to become the next active router and sends periodic hello messages. There must be at least one standby router in the HSRP group.
- **Active state**—In the active state, the router is currently forwarding packets that are sent to the virtual MAC address of the group. The active router sends periodic hello messages. At least one active router must be in the HSRP group.

Not all HSRP routers will transition through all states. For example, a router that is not in the standby or active router will not enter the standby or active states.

Configuring HSRP

This section discusses the configuration of HSRP. The active components of the configuration are HSRP standby interface, standby preempt, and hello message timers. This section also covers tracking the HSRP interface and troubleshooting HSRP.

Configuring an HSRP Standby Interface

To configure a router as a member of an HSRP standby group, enter the following command in interface configuration mode:

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

where the optional *group-number* argument 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.

The *ip-address* argument indicates the IP address of the virtual HSRP router

While running HSRP, it is important that the end-user stations do not discover the actual MAC addresses of the routers in the standby group. Any protocol that informs a host of the router's actual address must be disabled. To ensure that the actual addresses of the participating HSRP routers are not discovered, enabling HSRP on a Cisco router interface automatically disables ICMP redirects on that interface.

NOTE

Support for HSRP and ICMP redirects was introduced with IOS 12.1(3)T.

After the **standby ip** command is issued, the interface changes to the appropriate state. The following is an example of the state message generated. This message is automatically generated upon successful execution of the **standby ip** command:

```
3w1d   : %STANDBY-6-STATECHANGE: Standby: 50: Vlan100 state Speak    -> Standby
3w1d   : %STANDBY-6-STATECHANGE: Standby: 50: Vlan100 state Standby  -> Active
```

Example 9-2 illustrates that interface VLAN100 is a member of the HSRP standby group 50, that the virtual router IP address for that group is 192.168.100.50, and that ICMP redirects are disabled.

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

Example 9-2 *show run* Command Output Displays VLAN Membership, HSRP Group Address, and the Status of ICMP Redirects

```
Router#show run
Building configuration...

Current configuration:
!
(text deleted)
interface Vlan100
 ip address 192.168.100.1 255.255.255.0
 no ip redirects
 standby 50 ip 192.168.100.50
!
```

Configuring HSRP Standby Priority

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 the following command in interface configuration mode:

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

where *group-number* indicates the HSRP standby group. This number can be in the range of 0 to 255. The *priority-value* argument indicates the number that prioritizes a potential HSRP router. The range is 0 to 255; the default is 100.

The router in an HSRP group with the highest priority becomes the forwarding router.

Example 9-3 illustrates that interface VLAN100 has a priority value of 160 in HSRP standby group 50. If this priority value is the highest number in that HSRP standby group, the RSM on which this interface resides is the active router for that group.

Example 9-3 *show run* Command Output Displays the HSRP Group Priority Value

```
Router#show run
Building configuration...

Current configuration:
!
(text deleted)
interface Vlan100
 ip address 192.168.100.1 255.255.255.0
 no ip redirects
 standby 50 priority 160
 standby 50 ip 192.168.100.50
```

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

Configuring 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 on the active router:

```
Router(config-if)#standby group-number preempt
```

After the **standby preempt** command is issued, the interface changes to the appropriate state. The following is an example of the state message generated.

```
3w1d : %STANDBY-6-STATECHANGE: Standby: 50: Vlan100 state Standby -> Active
```

This message is automatically generated as soon as the router becomes active in the network.

Example 9-4 states that interface VLAN100 is configured to resume its role as the active router in HSRP group 50, assuming interface VLAN100 on this router has the highest priority in that standby group.

Example 9-4 show run Command Output Displays the Role of the Active Router

```
Router#show run
Building configuration...

Current configuration:
!
(text deleted)
interface Vlan100
 ip address 192.168.100.1 255.255.255.0
 no ip redirects
 standby 50 priority 160
 standby 50 preempt
 standby 50 ip 192.168.100.50
```

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

Configuring the 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 standby router. The Hello message contains the priority of the router, as well as a *hellotime* and *holdtime* value. The *hellotime* value indicates the interval between the hello messages that the router sends. The *holdtime* value contains 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.

TIP In general, the value of the holdtime should be at least three times the value of the hellotime and is required to be greater than the hellotime.

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

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

where the optional *group-number* argument indicates the group number on the interface to which the timers apply. The default is 0. The *hellotime* argument indicates the hello interval in seconds. This interval is an integer from 1 through 255. The default is 3 seconds. The *holdtime* argument indicates the 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.

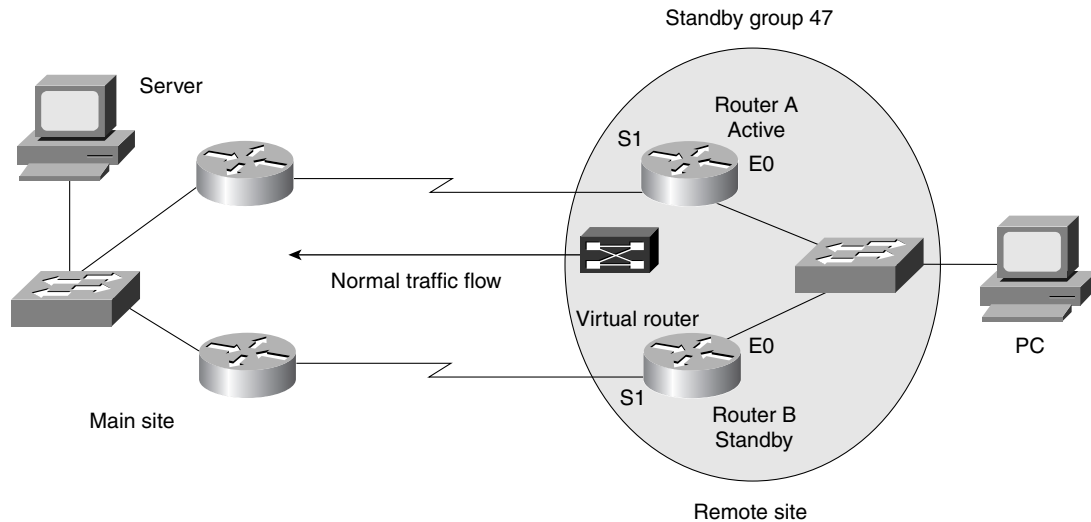
To reinstate the default standby timer values, enter the **no standby group timers** command.

Understanding 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 the campus LAN example in Figure 9-6, Router A and Router B reside in a branch office. These two routers each support a serial 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. Router A and Router B are exchanging hello messages through their E0 interfaces.


However, the serial link between the active forwarding router for the standby group and headquarters has an outage. Without HSRP enabled, Router A detects the failed link and sends an ICMP redirect to Router B. When HSRP is enabled, however, ICMP redirects are disabled. Therefore, neither Router A nor the virtual router sends an ICMP redirect and, 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 the headquarters cannot be routed.

Figure 9-6 Standby Group Normal Operation

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 the campus LAN example in Figure 9-7, the E0 interface on Router A tracks the S1 interface. If the link between the S1 interface and headquarters fails, the router automatically decrements its priority on that interface and stops transmitting hello messages out interface E0. Router B assumes the active router role when no hello messages are detected for the specific holdtime period.

Figure 9-7 Standby Group Experience Failure



Configuring HSRP Tracking

To configure HSRP tracking, enter the following command in interface configuration mode:

```
Router(config-if)#standby group-number track type-number interface-priority
```

The command arguments for this command are defined as follows:

- *group-number*—This optional argument indicates the group number on the interface to which the tracking applies. The default number is 0.
- *type*—This argument indicates the interface type (combined with the interface number) to be tracked.
- *number*—This argument indicates the interface number (combined with the interface type) to be tracked.
- *interface-priority*—When the interface becomes disabled, this optional argument indicates the amount by how much the HSRP priority for the router is decremented. 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.

HSRP Status

To display the status of the HSRP router, enter the following command in privileged EXEC mode.

```
Router#show standby type-number group brief
```

The command options for this command are defined as follows:

- *type-number*—This optional argument indicates the target interface type and number for which output is displayed.
- *group*—This optional argument indicates a specific HSRP group on the interface for which output is displayed.
- **brief**—This option displays a single line of output summarizing each standby group.

If these optional interface parameters are not indicated, the **show standby** command displays HSRP information for all interfaces. Example 9-5 presents the output resulting when the *type-number* and *group* parameters are specified.

Example 9-5 show standby Command Output with type-number and group Parameters Specified

```
Router#show standby Vlan100 50
Vlan100 - Group 50
  Local state is Active, priority 150, may preempt
  Hello time 3 holdtime 10
  Next hello sent in 00:00:02.944
  Hot standby IP address is 192.168.100.50 configured
  Active router is local
  Standby router is 192.168.100.1 expires in 00:00:08
  Standby virtual mac address is 0000.0d05.ab10
  Tracking interface states for 1 interface, 1 up:
    Up Vlan51 Priority decrement: 40
```

Example 9-6 demonstrates output resulting when the **brief** parameter is specified.

Example 9-6 show standby brief Command Output

```
Router#show standby brief
Interface  Grp  Prio P State   Active addr   Standby addr   Group addr
Vl100     50  160 P Active   local         192.168.100.1  192.168.100.50
Vl112     12  100 Standby  192.168.102.1 local          192.168.12.10
```

Troubleshooting HSRP

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

```
Router#debug standby
```

CAUTION Activating the **debug** command can cause your system to be unusable due to the high priority assigned to this process in the CPU. Use with caution!

Example 9-7 displays the **debug standby** command output as the router with the IP address 192.168.100.1 initializes and negotiates for the role of the active router.

Example 9-7 *debug standby Command Output*

```
3w1d : %STANDBY-6-STATECHANGE: Standby: 0: Vlan100 state Init -> Listen
3w1d : %STANDBY-6-STATECHANGE: Standby: 0: Vlan100 state Listen -> Speak
3w1d : SB0:Vlan100 Hello out 192.168.100.1 Speak pri 160 hel 3 hol 10 ip
192.168.100.50
3w1d : SB0:Vlan100 Hello out 192.168.100.1 Speak pri 160 hel 3 hol 10 ip
192.168.100.50
3w1d : SB0:Vlan100 Hello out 192.168.100.1 Speak pri 160 hel 3 hol 10 ip
192.168.100.50
3w1d : SB0:Vlan100 Hello out 192.168.100.1 Speak pri 160 hel 3 hol 10 ip
192.168.100.50
3w1d : %STANDBY-6-STATECHANGE: Standby: 0: Vlan100 state Speak -> Standby
3w1d : %STANDBY-6-STATECHANGE: Standby: 0: Vlan100 state Standby -> Active
3w1d : SB: Vlan100 Adding 0000.0d05.ab10 to address filter
```

The output of this **debug** command shows the various states of HSRP. The states are initial, listen, learn, speak, standby, and active. In the end, VLAN100 is in an active state.

To disable the debugging feature, enter either the **no debug standby** command or **the no debug all** command. You may also use **undebg standby** or **undebg all**.

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however, they are designed to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to these questions can be found in Appendix A, on page 477.

- 1 What is the name of the protocol that allows a set of routers that are working together to form one virtual router?

- 2 What problem makes HSRP necessary?

- 3 What is the minimum number of routers needed to perform HSRP?

- 4 What is the RFC that pertains to HSRP?

- 5 In a properly functioning virtual router, what happens when one of the routers fails?

6 How many standby groups can exist on any one LAN?

7 What constitutes an HSRP group?

8 What is the role of the active router?

9 Name the six states that an HSRP configured router can be in.

10 When configuring HSRP on a particular router interface, if the standby group is not explicitly configured, what standby group does the interface fall into by default?

11 Assume you are using five VLANs within your network and want to implement HSRP. How many HSRP groups would you need to create?

12 What command is used to display the HSRP virtual router IP and MAC address?

13 Which router in an HSRP group becomes the forwarding router and how is it determined?

14 In the command **standby 35 priority 90**, what does the “35” stand for?

15 An HSRP router exchanges hello messages with other HSRP routers. What is contained in the hello message?

16 What does the term *tracking* imply in an HSRP environment?

17 What command would allow you to debug HSRP?

18 What does the **preempt** command do in the HSRP environment?

19 What command or commands enable the preempt role in an HSRP-enabled network?

20 What command shows the status of an HSRP router?


Scenarios

The following scenarios and questions are designed to draw together the content of the chapter and exercise your understanding of the concepts. There is not necessarily a right answer. The thought process and practice in manipulating the concepts is the goal of this section.

Scenario 9-1

Examine the network in Figure 9-8 and answer the questions that follow.

Figure 9-8 Scenario 9-1 Network Topology



In Figure 9-8, Host X is using an application that requires uninterruptible access to the Internet. We must enable HSRP on the two routers so that one will take over in the event of a failure of either of the T1 access links to the Internet.

- 1 Assuming we are using a standby group of 99 and the virtual router IP address is 192.168.1.2, configure basic HSRP on Router X. Assume the interface is VLAN 100.
- 2 Configure Router B to be the active router.
- 3 Assume the T1 connected to Router B fails. Configure Router B to return to the active router state when the T1 circuit is recovered.
- 4 Configure the hellotime of 10 and then the holdtime, using established parameters.

Scenario Answers

The answers provided in this section are not necessarily the only possible correct answers. They merely represent one possibility for each scenario. The intention is to test your base knowledge and understanding of the concepts discussed in this chapter.

Should your answers be different (as they likely will be) consider the differences. Are your answers in line with the concepts of the answers provided and explained here? If not, go back and read the chapter again, focusing on the sections related to the problem scenario.

The key here is for you to gain an understanding of the topics.

Scenario 9-1 Answers

- 1 To activate HSRP on an interface, you must use the **standby** command as follows:

```
interface Vlan10
  ip address 192.168.1.1 255.255.255.0
  no ip redirects
  standby 99 ip 192.168.1.2
```

- 2 Configuring Router B to be the active router involves the use of the **priority** command as shown in the following sample configuration:

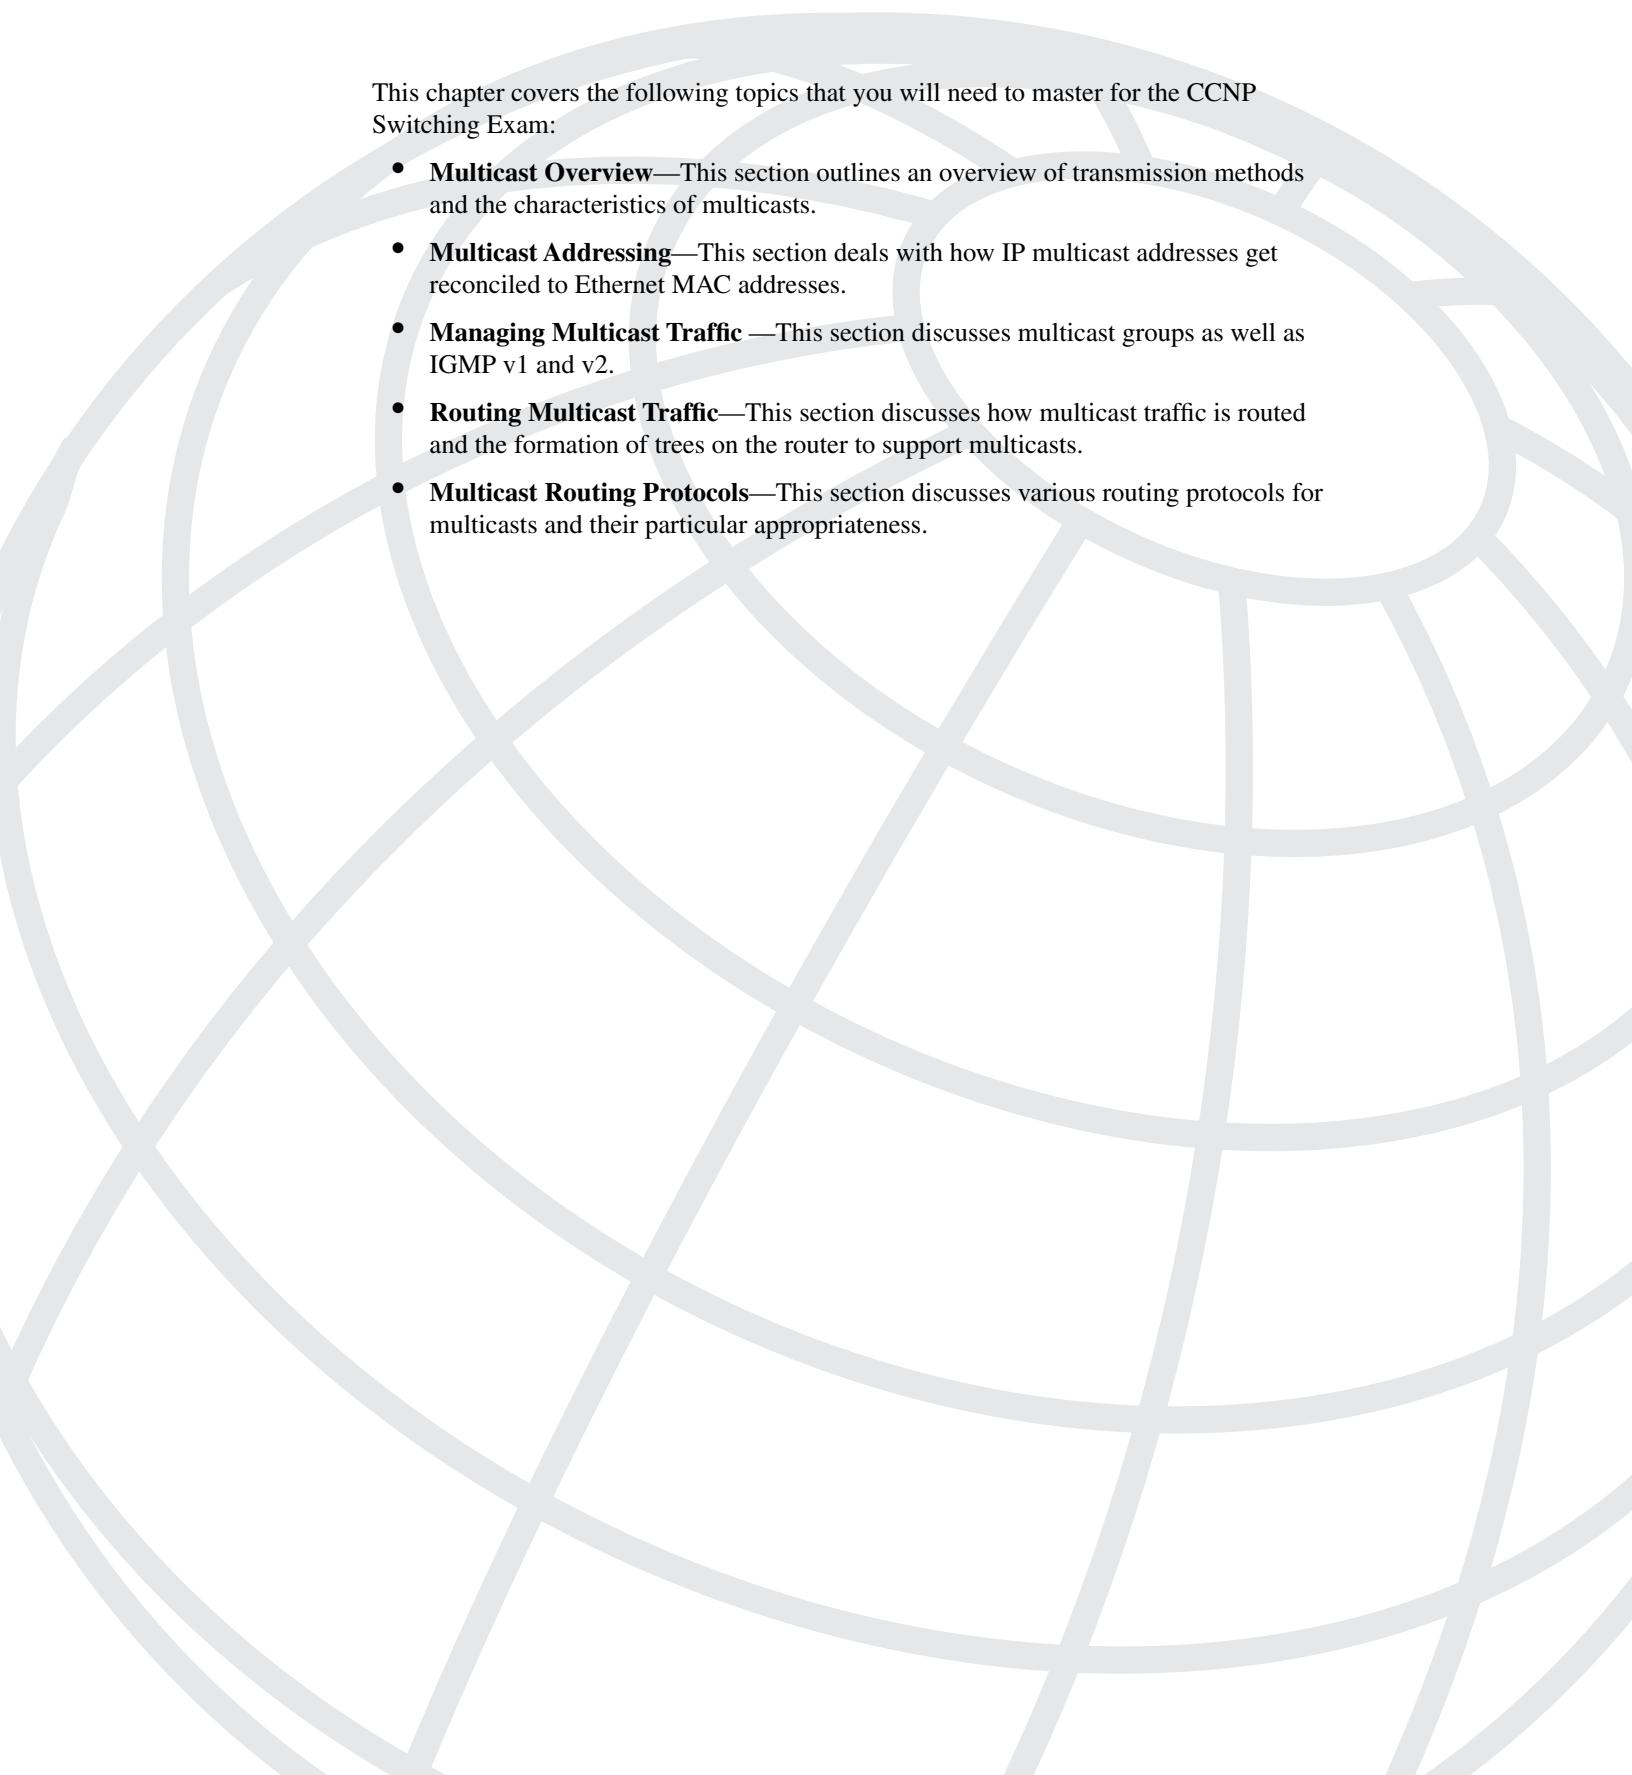

```
interface Vlan10
  ip address 192.168.1.3 255.255.255.0
  no ip redirects
  standby 99 ip 192.168.1.2
  standby 99 priority 150
```

- 3 Building on the configuration as we go forward, we add the **preempt** command to allow Router B to recover as the active router:

```
interface Vlan10
  ip address 192.168.1.3 255.255.255.0
  no ip redirects
  standby 99 ip 192.168.1.2
  standby 99 priority 150
  standby 99 preempt
```

- 4 Continue to build on the configuration here by adding the **timers** command. Use the standard practice of having holdtime at least three times the hellotime.

```
interface Vlan10
  ip address 192.168.1.3 255.255.255.0
  no ip redirects
  standby 99 ip 192.168.1.2
  standby 99 priority 150
  standby 99 preempt
  standby 99 timers 10 30
```

This chapter covers the following topics that you will need to master for the CCNP Switching Exam:

- **Multicast Overview**—This section outlines an overview of transmission methods and the characteristics of multicasts.
- **Multicast Addressing**—This section deals with how IP multicast addresses get reconciled to Ethernet MAC addresses.
- **Managing Multicast Traffic** —This section discusses multicast groups as well as IGMP v1 and v2.
- **Routing Multicast Traffic**—This section discusses how multicast traffic is routed and the formation of trees on the router to support multicasts.
- **Multicast Routing Protocols**—This section discusses various routing protocols for multicasts and their particular appropriateness.

Multicasts

Today's campus networks support intranet applications that operate between one sender and one receiver. The world is changing rapidly and this model will be superseded by a new paradigm. In the emerging campus network, a demand is there for intranet and multimedia applications where one sender will transmit to a group of receivers simultaneously. Applications like these include transmitting all-hands messages to employees, video and audio broadcasting, interactive video distance learning, transmitting data from a centralized data warehouse to multiple departments, communication of stock quotes to brokers, and collaborative computing. For example, an internal technical support facility might use multicast file transfer software to send software updates to multiple users in the campus simultaneously in one stream of data.

With the advent of MP3s, streaming video, MPEGs, voice, and other cutting edge multimedia applications, the toll on the network infrastructure has never been greater. Reaping the benefits of these applications and integrating them into the campus network can be a challenge. High bandwidth use is the order of the day. The capability of today's networks to combine the text, graphics, audio, and video components and effectively deliver them is increasingly difficult.

To that end, multimedia traffic types can traverse the network in one of the following forms:

- Unicast
- Broadcast
- Multicast

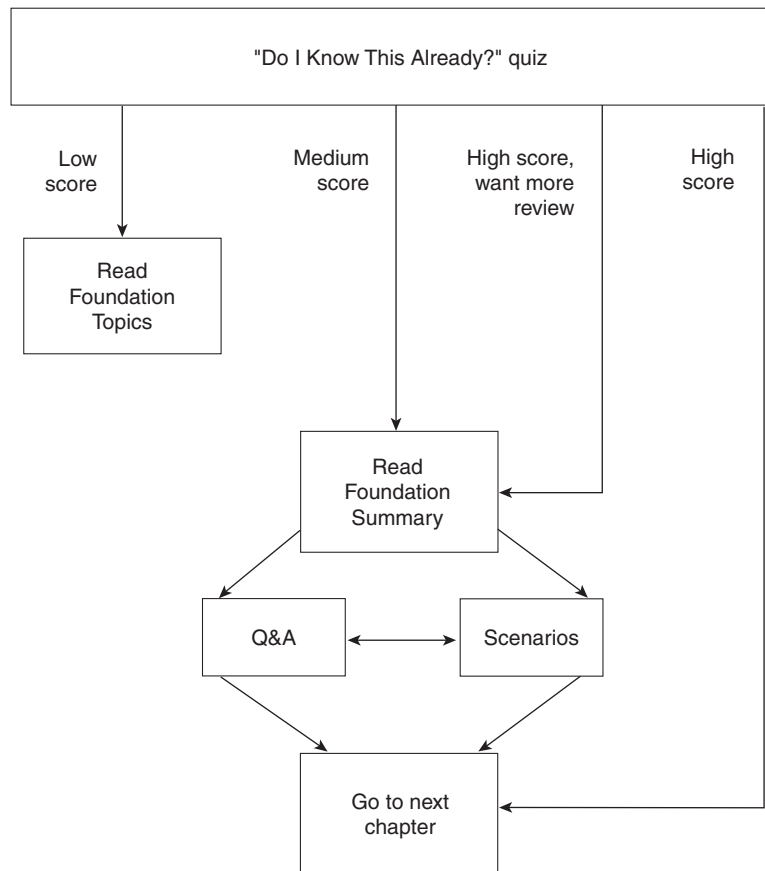
Each one of these methods of transmission has a different effect on network bandwidth, as you will see in the "Foundation Topics" section of the chapter.

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and answers for all your work with this book in one place for easy reference.
- Take the “Do I Know This Already?” quiz and write down your answers. Studies show retention is significantly increased through writing facts and concepts down, even if you never look at the information again.
- Use the diagram in Figure 10-1 to guide you to the next step.

Figure 10-1 *How to Use This Chapter*



“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The quiz helps you make good choices of how to spend your limited study time. The quiz has five sections, or “quizlets,” which correspond to the five major headings in the “Foundation Topics” section of the chapter. Use the scoresheet in Table 10-1 to record your score.

Table 10-1 *Scoresheet for Quiz and Quizlets*

Quizlet Number	Foundation Topics Section Covering These Questions	Questions	Score
1	Multicast Overview	1–2	
2	Multicast Addressing	3–4	
3	Managing Multicast Traffic	5–7	
4	Routing Multicast Traffic	8–10	
5	Multicast Routing Protocols	11–13	
All questions		1–13	

1 Name the three types of traffic available in today’s multimedia environment.

2 What Layer 4 protocol is used to carry multicast traffic?

3 What Class of IP address is used in a multicast environment?

4 Describe the makeup of the Class D multicast address by octet or bits.

5 What is the name of the protocol used to report their multicast group membership with neighboring multicast routers?

6 What is the special name assigned to the one multicast router that performs host membership queries to determine which groups have members?

7 What does a host send to the multicast group address to join a group?

8 Which type of routing involves transmitting packets from one source to one source?

9 Define a distribution tree.

10 Name the two types of distribution trees.

11 Name the three types of dense mode routing protocols.

12 Name the two types of sparse mode routing protocols.

13 Which multicast routing protocol is widely used on the MBONE?

The answers to the “Do I Know This Already?” quiz are found in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections,” on page 477. The suggested choices for your next step are as follows:

- **8 or fewer overall score**—Read the chapter. This includes the “Foundation Topics,” the “Foundation Summary,” Q&A, and scenarios at the end of the chapter.
- **9–12 overall score**—Begin with the “Foundation Summary” and then follow with the Q&A and scenarios at the end of the chapter.
- **13 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” and then go to the Q&A and scenarios at the end of the chapter. Otherwise, move to the next chapter.

Foundation Topics

Multicast Overview

This section discusses the comparison of unicast, broadcast, and multicast transmission and why having multicast capability is important. Also covered are the characteristics of multicasts.

Unicast Traffic

In a unicast architecture, a given application will send a single copy of a packet to every client unicast address that is called out. Unicast is basically a one-to-one relationship, potentially carried out many times. There are scaling issues related to doing unicasts. In the event that the unicast group is large and diverse, the potential to carry the same traffic multiple times is great.

Using the effects of Moore's law, which states that every 18 months the capacity of transistors on a semiconductor chip doubles, we see that the technology has matured enough to make it possible to afford to outfit every user with a unicast connection to the Internet. While possible and probable, contrast this with a video application like IP TV. The bandwidth required is, simply put, huge.

The concerns of network managers when it comes to unicast traffic consist of the number of user connections and the amount of replicated unicast transmissions.

Let's take the case of an IP TV server in a unicast example.

NOTE

IP TV is a streaming video server and application capable of doing both unicast and multicasts.

The server must send a separate TV stream for each client requesting access to the application. For example, an IP TV server sends a single channel of broadcast content to each client in the network. Let's assume for the moment that we need approximately 2 Mbps of bandwidth to support the application. The formula for required bandwidth is fairly simple:


$$2.0 \times n \text{ Mbps of link bandwidth}$$

where n is equal to the number of client viewers.

Assuming a 10-Mbps Ethernet interface on the server, five server-to-client streams would completely saturate the network interface. If we bump that up to a higher bandwidth, such as Fast Ethernet, a theoretical limitation of the interface is around 50 2.0-Mbps video streams. In practice, even that wouldn't be achievable.

Replicated unicast transmissions consume bandwidth within the network. The path between server and client must take into account the number of router and switch hops that occur between the two points. As routers are added to the path, the data is replicated across the link, as demonstrated by Figure 10-2.

Figure 10-2 *Replicated Unicast Traffic*



If 100 clients are separated from the server by two router hops and two switch hops, a single multi-unicast channel will consume 300 Mbps of router bandwidth and 300 Mbps of switch bandwidth.


Even if the video stream bandwidth is scaled back to 100 kbps, which provides acceptable quality in smaller windows on the end-station screen, the multi-unicast will consume 20 Mbps of both router and switch bandwidth.

Because other choices are available for sending multimedia traffic, unicast multimedia is used on a limited basis. Replicated unicast cannot scale up to efficiently deliver traffic to large numbers of end stations, but may be suitable for small numbers of destinations.

Broadcast Traffic

In a broadcast design, an application sends only one copy of each packet using a broadcast address. If this technique is used, however, broadcasts either must be stopped at the broadcast domain boundary with a Layer 3 device or transmitted to all devices in the campus network. Broadcasting a packet to all devices can be inefficient if only a small group in the network actually needs to see the packet as demonstrated in Figure 10-3.

Figure 10-3 *Broadcast Traffic*




Broadcast multimedia is dispersed throughout the network just like normal broadcast traffic. As with normal broadcasts, every client has to process the broadcast multimedia data frame. However, unlike standard broadcast frames which are generally small, multimedia broadcasts can reach as high as 7 Mbps or more of data. Even if an end station is not using a multimedia application, the device still processes the broadcast traffic. This requirement can use most, if not all, of the allocated bandwidth for each device. For this reason, the broadcast multimedia method is rarely implemented.

Multicast Traffic

The most efficient solution for transmitting multimedia is one in which a multimedia server sends one copy of each packet, addressing each packet to a special multicast address. Unlike the unicast environment, a multicast server sends out a single data stream to multiple clients. Unlike the broadcast environment, the client device decides whether to listen to the multicast address. Multicasting saves bandwidth and controls network traffic by forcing the network to replicate packets only when necessary. By eliminating traffic redundancy, multicasting reduces network and host processing.

In the example shown in Figure 10-4, the IP TV server transmits a single TV stream for each multicast group. *Multicast group* indicates which hosts have joined a particular group for the purposes of receiving multicast traffic. In this case two multicast groups are defined, A and B. In this example, two hosts are defined as part of Multicast Group A and just one in Multicast Group B.

Figure 10-4 *Multicast Traffic*



The example in Figure 10-4 illustrates that only clients subscribed to a particular multicast address and part of a multicast group can receive the IP TV broadcast.

If we assume that this IP TV broadcast is utilizing 2.0 Mbps of the bandwidth, we can see that the remaining bandwidth is free to be used for other applications. In addition, this bandwidth is only being used if the particular host is subscribed.

Characteristics of Multicast Traffic

The concept of IP multicast is defined as sending IP packets to a group of hosts on the network. One of the obvious benefits of this technology is the preservation of bandwidth by sending a single data stream to a group of clients instead of sending to all clients or having multiple streams at once. IP multicast is first described in RFC 1112, *Host Extensions for IP Multicasting*. A more current RFC, RFC 2236, describes IGMP, Version 2.

IP multicasting has the following characteristics:

- Facilitates transmission of an IP datagram to a multicast group comprised of zero or more hosts identified by a single IP destination address
- Delivers a multicast datagram to all members of the multicast group with the same “best-effort” reliability as regular unicast IP datagrams
- Supports dynamic membership of a multicast group
- Supports all multicast groups regardless of the location or number of members
- Supports the membership of a single host in one or more multicast groups
- Upholds multiple data streams at the application level for a single group address
- Supports a single group address for multiple applications on a host

Another benefit of multicasting is that it is limited in network delay. This limitation is due to the one-to-many nature of multicasting, which limits the path. In contrast, unicasting transmits multiple copies of the same stream to potentially large numbers of hosts.

Multicasting carries with it the benefit of being anonymous. This anonymity is accomplished because any given server transmits to a single multicast group address, representing an entire group of recipients. The server never knows the unicast network address of any given recipient.

Multicast traffic is handled at the transport layer using the User Datagram Protocol (UDP). Unlike the Transmission Control Protocol (TCP), UDP has no reliability functionality, which means no error correction or flow control. Because of the simplicity of UDP, data packet headers contain fewer bytes and consume less network overhead than TCP.

Multicast Addressing


This section covers multicast Class D addressing structure and range, as well as how to map Ethernet MAC addresses to Class D IP addresses.

Multicast Address Structure

IP multicasting is the transmission of an IP data frame to a multicast group, identified by a single IP address. Because the multicast group is identified by a single IP address rule, the IP multicast datagram contains a specific combination of the destination MAC address and a destination IP address.

The range of IP addresses is divided into classes based on the high order bits of a 32-bit IP address. IP multicast uses Class D addresses. A Class D address consists of 1110 as the higher order bits in the first octet, followed by a 28-bit group address. Unlike Class A, B, and C IP addresses, the last 28 bits of a Class D address are unstructured, as illustrated by Figure 10-5.

Figure 10-5 IP Multicast Uses Class D Addresses



These remaining 28 bits of the IP address identify the multicast group ID. This multicast group ID is a single address typically written as decimal numbers in the range 224.0.0.0 through 239.255.255.255. The high-order bits in the first octet identify this 224-base address.

Multicast addresses may be dynamically or statically allocated. Dynamic multicast addressing provides applications with a group address on demand. Because dynamic multicast addresses have a specific lifetime, applications must request this type of address only for as long as it is needed.

Statically allocated addresses are reserved for specific protocols that require well-known addresses. The Internet Assigned Numbers Authority (IANA) assigns these well-known addresses. These addresses are called *permanent host groups* and are similar in concept to the well-known TCP and UDP port numbers. Table 10-2 lists some of the “well-known” Class D addresses.

Address 224.0.0.1 identifies the all-hosts group. Every multicast-capable host must join this group at the start. If a **ping** command is issued using this address, all multicast-capable hosts on the network must answer the **ping** request.

Address 224.0.0.2 identifies the all-routers group. Multicast routers must join that group on all multicast-capable interfaces.

Table 10-2 *Well-Known Class D Addresses*

Well-Known Class D Address	Purpose
224.0.0.1	All hosts on a subnet
224.0.0.2	All routers on a subnet
224.0.0.4	All Distance Vector Multicast Routing Protocol (DVMRP) routers
224.0.0.5	All Open Shortest Path First (OSPF) routers
224.0.0.6	All OSPF designated routers
224.0.0.9	All Routing Information Protocol, version 2 (RIP-2) routers
224.0.0.13	All Protocol Independent Multicast (PIM) routers

Addresses ranging from 224.0.0.0 through 224.0.0.255 are reserved for local purposes, such as administrative and maintenance tasks. Multicast routers do not forward datagrams destined to this range of addresses.

Similarly, the address range 239.0.0.0 to 239.255.255.255 is reserved for administrative scoping. An administrative scope zone is defined by a set of routers surrounding a region within the network. These routers are configured to deny multicast traffic in a particular address range from entering or leaving the zone. This technique is useful in containing high-bandwidth traffic to a specific region in the campus network. Administrative scoping is outlined in RFC 2365, *Administratively Scoped IP Multicast*.

You can retrieve a document containing the current list of multicast assigned address at <ftp://ftp.isi.edu/in-notes/iana/assignments/>.

Mapping IP Multicast Addresses to Ethernet

Ethernet frames have a 48-bit destination address field. To avoid invoking the Address Resolution Protocol (ARP) to map multicast IP addresses to Ethernet addresses, the IANA designated a range of Ethernet addresses for multicast. The lower 23 bits of the Class D address are mapped into a block of Ethernet addresses that have been reserved for multicast. This block includes addresses in the range 00:00:5e:00:00:00 through 00:00:5e:ff:ff:ff. The IANA allocates half of this block for multicast addresses. Given that the first byte of any Ethernet address must be 01 to specify a multicast address, the Ethernet addresses corresponding to IP multicasting are in the range 01:00:5e:00:00:00 through 01:00:5e:7f:ff:ff.

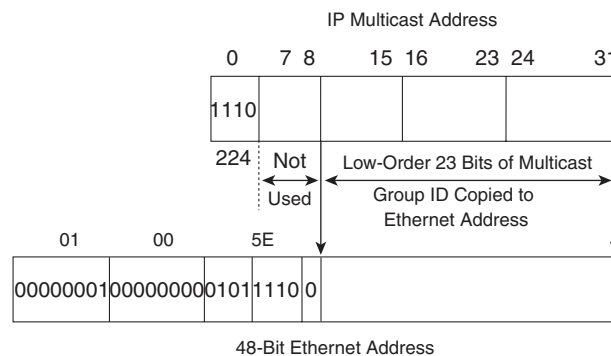
NOTE Multicast MAC and IP addresses only belong in destination addresses.

The prefix 01-00-5e identifies the frame as multicast; the next bit is always 0, leaving only 23 bits for the multicast address. Because IP multicast groups are 28-bits long, the mapping cannot be one-to-one. Only the 23 least-significant bits of the IP multicast group are placed in the frame. The remaining five high-order bits are ignored, resulting in 32 different multicast groups being mapped to the same Ethernet address. For example, if the IP multicast group address is 224.10.8.5, the destination MAC address becomes 01-00-5E-0A-8-5.

When mapping multicast addresses to MAC addresses, all 28 bits of the IP multicast cannot be mapped to the available 23 bits of MAC address space, meaning that five bits of address are lost. This also means that 32, or 2^5 , addresses can be ambiguous when mapped. Therefore, each IP multicast MAC address is capable of representing 32 IP multicast addresses.

Figure 10-6 illustrates a multicast group address being mapped into an IEEE-802 multicast address. In multicast addressing, the high-order nine bits of the IP address are not mapped into the MAC-layer multicast address. In the example in Figure 10-6, the mapping places the low-order 23 bits of the IP multicast group ID into the low order 23 bits of the IEEE-802 multicast address.

Figure 10-6 Mapping Multicast Address to Ethernet



Managing Multicast Traffic

In the world of multicasting, the concept of sending a multicast within a single broadcast domain is a somewhat trivial process. However, when we extend this into multiple segments within a campus environment over multiple switches and routers, we complicate matters significantly.

The sending process specifies a destination address defined as a multicast address. The device driver in the sending server converts this address to the corresponding Ethernet address and sends the packet out on the network. The receiving devices, or clients, must indicate that they want to receive datagrams destined for a given multicast address. Complications arise when multicasting is extended beyond a single physical network and multicast packets pass through routers.

Sending and receiving multimedia require coordination from all devices participating in the multicast. These devices include the server, the host, the router, and the switch. Some of the issues in facilitating multimedia traffic in the campus network are

- Coordinating the multicast operations of the different devices in the network
- Establishing a path between source and destination devices
- Forwarding multicast traffic through the network

IP multicast traffic for a particular source/destination group pair is transmitted from the source to the multicast group via a distribution tree. This distribution tree connects all the hosts in the group. Different IP multicast routing protocols use different techniques to construct these multicast Spanning Trees; after the tree is constructed, however, all multicast traffic for a specific group is distributed over this tree.

Before multicast traffic can traverse the network, routers need to know which hosts, if any, on a specific physical network belong to a given multicast group. Because the emerging campus network model is comprised of both routers and switches, switches also need to know how to direct multicast traffic. Cisco switches do this through the use of CGMP, which is discussed later in this chapter.

Subscribing and Maintaining Groups

The Internet Group Management Protocol (IGMP) provides a means to report their multicast group membership with neighboring multicast routers. The IGMP protocol, version 1, is defined in RFC 1112, *Host Extensions for IP Multicasting*. RFC 2236, *Internet Group Management Protocol, Version 2*, defines IGMP, version 2.

IGMP manages multicast traffic throughout networks through the use of special multicast queriers and hosts. A *querier* is a network device, such as a router, that sends IGMP queries. A set of queriers and hosts that receive multicast data streams from the same source is called a *multicast group*. Queriers and hosts use IGMP messages to join and leave multicast groups.

IGMP supports two specific message structures, as follows:

- Query messages are used to discover which network devices are members of a given multicast group.
- Report messages are sent by hosts in response to query messages to inform the querier of a host membership.

NOTE

Within the IP multicast model there is no notion of membership between a source and the receivers. A source does not have to be a member of a group to send traffic to that group. Conversely, receivers do inform routers of what groups the receivers want to have membership in so that the router can forward the appropriate traffic flows.

Several versions of IGMP are available; IGMP versions 1 and 2 are now in production, with version 3 in development. Each version has its own set of behavior characteristics.

IGMP Version 1

IGMP uses IP datagrams to transmit information about multicast groups. The datagram consists of a 20-byte IP header and an 8-byte IGMP message.

According to the IGMPv1 specification, one multicast router per LAN must periodically transmit Host Membership Query messages to determine which host groups have members on the querier's directly attached networks. IGMP query messages are addressed to the all-host group (224.0.0.1) and have an IP Time-To-Live (TTL) equal to one. This TTL ensures that the Query messages sourced from a router are transmitted onto the directly attached network but are not forwarded by any other multicast routers.

When the end station receives an IGMP query message, the end station responds with a host membership report for each group into which the end station belongs.

IGMP messages are specified in the IP datagram with a protocol value of 2. Table 10-3 describes the fields of the IGMP message.

Table 10-3 *IGMPv1 Message Format Fields*

Field Name	Value
Type	Two types of IGMP messages are of concern to hosts: 1 = Host Membership Query 2 = Host Membership Report
Unused	Unused field, zeroed when sent, ignored when received.
Checksum	The checksum is the 16-bit one's complement of the one's complement sum of the 8-octet IGMP message. For computing the checksum, the checksum field is zeroed.
Group Address	In a Host Membership Query message, the group address field is zeroed when sent, ignored when received. In a Host Membership Report message, the group address field holds the IP host group address of the group being reported.

Joining a Group Using IGMP Version 1

Hosts joining a group do not have to wait for a query to join. When a host wants to join a multicast group, the host sends a Host Membership Report to the group address. This unsolicited request reduces join latency for the end system when no other members of that group are present on that network segment.

General Queries Using IGMP Version 1

Multicast routers send Host Membership Query messages to discover which host groups have members on their attached local networks. General queries go to the all-hosts (224.0.0.1) multicast address and carry a TTL of one (1). One member from each group on the segment will respond with a report. General queries are sent out periodically based on the setting of the **ip igmp query-interval** command. (The default setting is 60 seconds.)

No formal IGMP query router election process exists within IGMPv1 itself. Instead, the election process is left up to the multicast routing protocol, and different protocols use different mechanisms. This process often results in multiple queriers on a single network segment supporting multiple multicast-enabled routers.

Membership Queries Using IGMP Version 1

To ensure the viability of group membership on a given network segment, the router multicasts periodic IGMPv1 membership queries to the all-hosts (224.0.0.1) group address. Only one member per group responds with a report to a query. This action saves bandwidth on the network segment and processing by the hosts. This process is called *report suppression*. The report suppression mechanism is accomplished as follows:

When a host receives the query, it starts a countdown timer for each multicast group for which the host is a member. The countdown timers are each initialized to a random count within a given time range. In IGMPv1, the time range is a fixed range of 10 seconds. Therefore, the countdown timers are randomly set to some value between 0 and 10 seconds.

When a countdown timer reaches zero, the host sends a membership report for the group associated with the countdown timer to notify the router that the group is still active. However, if a host receives a membership report before the associated countdown timer reaches zero, the host cancels the countdown timer associated with the multicast group thereby suppressing the host's own report.

Leaving a Group Using IGMP Version 1

No special leave mechanism was defined in IGMPv1. Instead, IGMPv1 hosts leave a group passively or quietly at any time without any notification to the router.

Multicast routers periodically transmit IGMP queries to refresh their knowledge of the group members present on each network interface. This process updates the local group database of the router. Eventually, the router should be able to detect that no members of a group are present on an interface any longer and, if possible, remove itself from the multicast delivery tree for this group. If the router does not receive a report from any members of a group after a number of queries, the router assumes that no group members are present on a particular interface.

When a router is just starting up or if multicast routing has just been enabled, a router may send several IGMP queries in rapid succession in order to quickly learn which groups have local members.

IGMP Version 2

As a result of some of the limitations discovered in IGMPv1, work began on IGMPv2 in an attempt to remove these limitations. Most of the changes between IGMPv1 and IGMPv2 are primarily to address the issues of Leave and Join latencies, as well as address ambiguities in the original protocol specification

Version 2 of IGMP made some enhancements to the previous version, including the definition of a Group-Specific Query. This type of message allows the router to transmit a Specific Query to one particular group. IGMPv2 also defines a Leave Group Message for the hosts, which results in lower leave latency.

Four types of IGMP messages are of concern to the host-router interaction:

- Membership query
- Version 2 membership report
- Leave report
- Version 1 membership report

The Version 1 membership report is used for backward-compatibility with IGMPv1. New message types may be used by newer versions of IGMP or by multicast routing protocols. Any unrecognized or other message types are silently ignored.

Table 10-4 provides a description of the IGMPv2 message fields.

Table 10-4 *IGMPv2 Message Fields*

Field Name	Value
Type	0 x 11 = Membership Query 0 x 12 = Version 1 Membership Report 0 x 16 = Version 2 Membership Report 0 x 17 = Leave Report 0 x 12 = Version 1 Membership Report
Maximum Response Time	10 seconds = Default value. Meaningful only in a membership query. Specifies the maximum allowed time before sending a responding report in units of 1/10 second. 0 = All other messages.

continues

Table 10-4 *IGMPv2 Message Fields (Continued)*

Field Name	Value
Checksum	Calculated the same as for the ICMP checksum.
Group Address	0 in a general query. Group address queried in a Group Specific Query. Multicast group address in a report.

Joining a Group Using IGMP v2

The process of joining a multicast group is the same in IGMPv2 as it is in IGMPv1. Like IGMPv1, IGMPv2 hosts joining a group do not have to wait for a query to join. When a host wants to join a multicast group, the host sends a host membership report to the multicast group address.

A host sends an IGMP join message when the host wants to join a multicast group. If the host and server reside in different subnets, the join message must go to a router. When the router intercepts the message, the router looks at its IGMP table. If the network number is not in the table, the router adds the information contained in the IGMP message.

Using queries and reports, a multicast router builds a table detailing which of the router interfaces have one or more hosts in a multicast group. When the router receives a multicast datagram, the router forwards the datagram to only those interfaces that have hosts with processes belonging to that group.

After a host has joined a multicast group, the host appears in the router's group database.

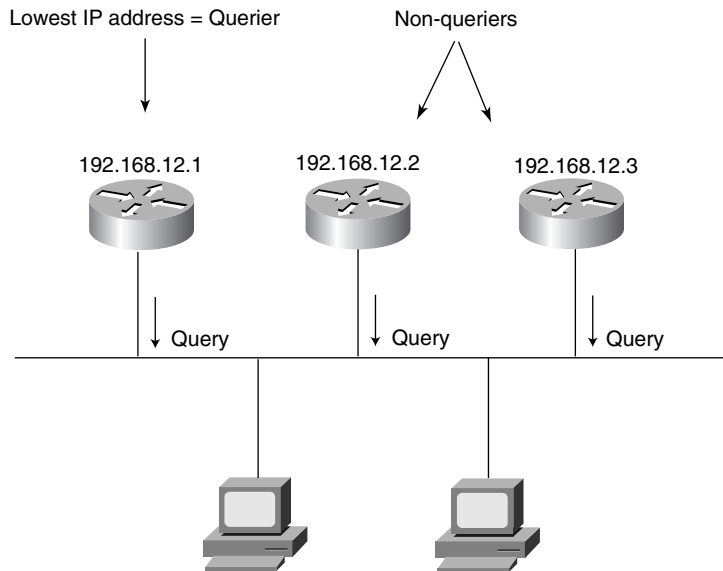
Querier Election Using IGMPv2

IGMPv2 defines a procedure for election of the multicast querier for each network segment. In IGMPv2, the multicast router with the lowest IP address on the LAN segment is elected the multicast querier.

Initially, every router on the segment believes itself to be the querier for every one of the router's interfaces that are multicast-enabled. When a router is first multicast-enabled, the router begins transmitting query messages. If the router subsequently detects a queried message that is sourced from a numerically lower IP address, the router ceases to act as a querier on that interface as demonstrated in Figure 10-7.

The Query-Interval Response time has been added to IGMPv2 to control the burstiness of reports. This value is indicated in queries to convey to the membership the time interval in which members have to respond to a query with a report message.

Figure 10-7 IGMPv2 Multicast Querier



A Group-Specific Query also was added in IGMPv2 to allow the router to query membership in only a single group instead of all groups. This addition is an optimized way to quickly find out if any members are left in a group without asking all groups for a report.

The difference between the Group-Specific query and the General Query is that a General Query is multicast to the all-hosts (224.0.0.1) address while a Group-Specific Query for group G is multicast to the group G multicast address.

To locate and verify the elected querier, enter the **show ip igmp** command in user or privileged EXEC mode as demonstrated in Example 10-1.

Example 10-1 Locating the Elected Querier with **show ip igmp**

```
RTR100>show ip igmp interface type number
Ethernet0 is up, line protocol is up
 Internet address is 198.92.37.6, subnet mask is 255.255.255.0
 IGMP is enabled on interface
 IGMP query interval is 60 seconds
 Inbound IGMP access group is not set
 Multicast routing is enabled on interface
 Multicast TTL threshold is 0
 Multicast designated router (DR) is 198.92.37.33
 No multicast groups joined
```

continues

Example 10-1 *Locating the Elected Querier with show ip igmp (Continued)*

```
Ethernet1 is up, line protocol is up
  Internet address is 198.92.36.129, subnet mask is 255.255.255.0
  IGMP is enabled on interface
  IGMP query interval is 60 seconds
  Inbound IGMP access group is not set
  Multicast routing is enabled on interface
  Multicast TTL threshold is 0
  Multicast designated router (DR) is 198.92.36.131
  Multicast groups joined: 225.2.2.2 226.2.2.2
Tunnel0 is up, line protocol is up
  Internet address is 10.1.37.2, subnet mask is 255.255.0.0
  IGMP is enabled on interface
  IGMP query interval is 60 seconds
  Inbound IGMP access group is not set
  Multicast routing is enabled on interface
  Multicast TTL threshold is 0
  No multicast groups joined
```

The designated router is a different function and is listed separately in the display above. The concept of the designated router is discussed in more detail in Chapter 11, “Configuring Multicast Networks.”

Maintaining a Group Using IGMPv2

Similar to IGMPv1, the IGMPv2 router multicasts periodic Membership Queries to the all-hosts (224.0.0.1) group address. Only one member per group responds with a report to a query. All other group members suppress their Membership Reports.

When a host receives a General Query, the host sets delay timers for each group, excluding the all-systems group (of which the host is a member on the interface from which the query was received). Each timer is set to a different random value, using the highest clock granularity available on the host.

When a host receives a Group-Specific Query, the host sets a delay timer to a random value for the group being queried if the host is a member of that group. If a timer for the group is already running, the host timer is reset to the random value only if the requested max response time is less than the remaining value of the running timer. When a group timer expires, the host multicasts a Version 2 Membership Report to the group, with an IP TTL of 1.

Leaving a Group Using IGMPv2

A Leave Group message was also added in IGMPv2. Whenever any end station wants to leave a group, the host transmits a Leave Group message to the all-routers group (224.0.0.2) with the group field indicating the group being left. This transmission allows end systems to tell the router the hosts are leaving the group. Thus, this action reduces the leave latency for the group on the segment when the member leaving is the last member of the group.

Switching Multicast Traffic Using CGMP

In the multilayer campus model, IP multicast traffic traverses a Layer 2 switch, especially at the access layer. Because IP multicast traffic maps to a corresponding Layer 2 multicast address, multicast traffic is delivered to all ports of a Layer 2 switch.

For example, a video client wants to watch a 1.5-Mbps IP multicast-based video feed sent from a corporate video server. The video client sends an IGMP join message to the video server. The next-hop router for the client logs the IGMP join message. IP multicast traffic is transmitted downstream to the video client. The switch detects the incoming traffic and examines the destination MAC address to determine where the traffic should be forwarded. Because the destination MAC address is a multicast address and there are no entries in the switching table directing the traffic, the 1.5-Mbps video feed is simply sent to all ports.

Switches must have an architecture that allows multicast traffic to be forwarded to a large number of attached group members without unduly loading the switch fabric. This function allows the switch to provide support for the growing number of new multicast applications without impacting other traffic. Layer 2 switches also need some degree of multicast awareness to avoid flooding multicasts to all switch ports.

Multicast control in Layer 2 switches can be accomplished in several ways:

- Virtual LANs (VLANs) can be defined to correspond to the boundaries of the multicast group. This approach is simple; however, it does not support dynamic changes to group membership and adds to the administrative burden of unicast VLANs.
- Layer 2 switches can snoop IGMP queries and reports to learn the port mappings of multicast group members. This action allows the switch to dynamically track group membership. However, snooping every multicast data and control packet consumes a lot of switch processing capacity and therefore can degrade forwarding performance and increase latency. However, most Cisco switches that implement IGMP snooping use specialized hardware (Application-Specific Integrated Circuits, or ASICs) and avoid the performance penalty.
- The traditional role of the router as a control point in the network can be maintained by defining a multicast router-to-switch protocol. The Cisco Group Management Protocol (CGMP) allows the router to work with the switch to configure the multicast forwarding table to correspond with the current group membership.

CGMP is a proprietary protocol developed by Cisco to enable Cisco Catalyst switches to learn about the existence of multicast clients from Cisco routers and Layer 3 switches.

CGMP is based on a client-server model. The router is considered a CGMP server, with the switch taking on the client role. The basis of CGMP is that the IP multicast router sees all IGMP packets and therefore can inform the switch when specific hosts join or leave multicast groups. The switch then uses this information to construct a forwarding table.

When the router sees an IGMP control packet, the router creates a CGMP packet. This CGMP packet contains the request type (either a join or a leave), the multicast group address, and the actual MAC address of the client. The packet is sent to a well-known address to which all switches listen. Each switch then interprets the packet and creates the proper entries in a forwarding table.

Building on the previous video example, the client starts by sending an IGMP join message to the neighboring multicast router. Now when the next-hop router receives the IGMP join message, however, the router records the source MAC address of the IGMP message and issues a CGMP join message downstream to the Catalyst switch. The Catalyst switch uses the CGMP message to dynamically build an entry in the switching table that maps the multicast traffic to the switch port of the client. In this example, the server delivers the 1.5-Mbps video feed only to those switch ports that are in the switching table. The ports on the switch that do not support any hosts in the multicast group do not propagate the traffic.

Routing Multicast Traffic

Campus networks typically have a large number of subnetworks, each being their own broadcast domain. As you know, routers must connect these subnetworks so that they can be routed from one broadcast domain to the next. This is, by definition, the function of the IP protocol.

Each host on the Internet has an address that identifies the physical location of the host. Part of the address identifies the subnet on which the host resides and part identifies the individual host on that subnet. Routers periodically send routing update messages to adjacent routers, conveying the state of the network as perceived by that particular router. This data is recorded in routing tables that are then used to determine optimal transmission paths for forwarding messages across the network.

Unicast transmission involves transmission from a single source to a single destination. The transmission is directed toward a single physical location that is specified by the host address. This routing procedure is relatively straightforward because of the binding of a single address to a single host.

Routing multicast traffic is a more complex problem. A multicast address identifies a particular transmission session rather than a specific physical destination. An individual host is able to join an ongoing multicast session by using IGMP to communicate this desire to the subnet router.

Because the number of receivers for a multicast session can potentially be quite large, the source does not need to know all the relevant addresses. Instead, the network routers must somehow be able to translate multicast addresses into host addresses. The basic principal involved in multicast routing is that routers interact with each other to exchange information about neighboring routers.

Multicast routing is based upon the construction of “trees,” connecting the members of the various multicast groups. The following sections discuss the types of trees and how they are constructed.

Distribution Trees

For efficient transmission of multicast traffic, designated routers construct a tree that connects all members of an IP multicast group. A *distribution tree* specifies a unique forwarding path between the subnet of the source and each subnet containing members of the multicast group.

A distribution tree has just enough connectivity so that there is only one loop-free path between every pair of routers. Because each router knows which of its lines belong to the tree, the router can copy an incoming multicast datagram onto all the outgoing branches. This action generates the minimum needed number of datagram copies. Because messages are replicated only when the tree branches, the number of copies of the messages transmitted through the network is minimized.

Because multicast groups are dynamic with members joining or leaving a group at any time, the distribution tree must be updated. Branches that contain new members must be added. Branches in which no listeners exist must be discarded, or pruned.

There are two basic tree construction techniques: source-specific trees and shared, or center-specific, trees.


Source-Specific Distribution Trees

Source-specific distribution trees require finding a shortest path from the sender to each receiver, resulting in multiple minimal delay trees for a group.

The source-specific method builds a Spanning Tree for each potential source, or subnetwork. These Spanning Trees result in source-based delivery trees emanating from the subnetworks directly connected to the source stations. Because many potential sources for a group exist, a different delivery tree is constructed, rooted at each active source. Figure 10-8 illustrates a source-specific distribution tree.

Source-based trees are constructed using a technique called *Reverse Path Forwarding (RPF)*. If a packet arrives on a link that the local router believes to be on the shortest path back toward the source of the packet, the router forwards the packet on all interfaces except the incoming interface. If the packet does not arrive on the interface that is on the shortest path back toward the source, the packet is discarded.

NOTE RPF is used for source trees and RP-rooted shared trees.

Figure 10-8 *Source-Specific Distribution Tree*

The interface over which the router expects to receive multicast packets from a particular source is referred to as the *parent* link. The outbound links over which the router forwards the multicast packet are called the *child* links for this source.

The RPF algorithm also reduces unnecessary packet duplication. If the local router making the forwarding decision can determine that a neighboring router on a child link is downstream, the packet is not forwarded to the upstream neighboring router. A downstream neighbor is a neighboring router that considers the local router to be on the shortest path back toward a given source.


If the link between the local router and the neighboring router is not the shortest path, the packet is not forwarded on that child link.

Shared Distribution Trees

The shared tree makes use of distribution centers and constructs a single multicast tree, resulting in a low overhead method but sacrificing minimal end-to-end delay.

Unlike source, or shortest path, tree algorithms that build a source-based tree for each source or each (source, group) pair, shared-tree algorithms construct a single delivery tree shared by all members of a group. The shared-tree approach is quite similar to the Spanning-Tree Algorithm except that the shared tree allows the definition of a different shared tree for each group. Devices wanting to receive traffic for a multicast group must explicitly join the shared delivery tree. Multicast traffic for each group is sent and received over the same delivery tree regardless of the source. Figure 10-9 illustrates a shared distribution tree.

Figure 10-9 *Shared Distribution Tree*



A shared tree may involve a single router or set of routers, which comprises the “core” of a multicast delivery tree. Shared-tree algorithms make efficient use of router resources because this technique requires only a router to maintain state information in each group, not each (source, group) pair.

Multicast routing protocols build distribution trees by examining a unicast reachability protocol routing table.

Scope of Delivery

As in unicast routing, the multicast TTL field controls the live time of the packet. The function of TTL is to prevent packets from being looped forever due to routing errors. However, the TTL field in multicasting also carries the concept of a “threshold.”

Multicast-enabled routers have a TTL threshold assigned to each interface. A multicast router will forward a multicast packet across an interface only if the TTL field in the IP header is

greater than the TTL threshold assigned to the interface. If the TTL field in the IP header of the packet is equal to or less than the TTL threshold assigned to the interface, the packet is discarded. If the interface has no assigned TTL threshold, the packet is forwarded. The router then decrements the packet TTL upon sending the packet out the interface

Table 10-5 provides a list of TTL thresholds and their associated scope.

Table 10-5 *TTL Thresholds*

Value	Action
0	Restricted to the same host; not output by any interface
1	Restricted to the same subnet; not forwarded by a router
15	Restricted to the same site, organization, or department
63	Restricted to the same region
127	Worldwide
191	Worldwide; limited bandwidth
255	Unrestricted in scope; global

For example, a multicast packet with a TTL of less than 16 is restricted to the same department, or site, and will not be forwarded across an interface to other sites in the same region. Defining the scope of a site or region is the responsibility of the network administrator.

Multicast Routing Protocols

A multicast routing protocol is responsible for the construction of multicast delivery trees and is necessary to permit the forwarding of multicast packets. Different IP multicast routing protocols use different techniques to construct multicast Spanning Trees and forward packets.

In general, IP multicast routing protocols follow one of two basic methods, depending largely on the number of multicast group members in the network.

Dense Mode Routing Protocols

The first method for multicast routing is based on the assumption that the multicast group members are densely distributed throughout the network and bandwidth is plentiful, meaning that almost all hosts on the network belong to the group. These dense mode multicast routing protocols rely on periodic flooding of the network with multicast traffic to set up and maintain the distribution tree.

Dense mode routing protocols include the following:

- Distance Vector Multicast Routing Protocol (DVMRP)
- Multicast Open Shortest Path First (MOSPF)
- Protocol Independent Multicast Dense Mode (PIMDM)

Dense mode routing protocol operations assume that almost all routers in the network will need to distribute multicast traffic for each multicast group. The dense mode protocols are most appropriate in environments with densely clustered receivers and the available bandwidth to tolerate flooding.

DVMRP

DVMRP is described in RFC 1075. DVMRP is widely used on the Internet multicast backbone (MBONE).

DVMRP uses a process called *reverse path flooding*. This process is similar to the split horizon process on a Cisco router.

When a router receives a packet, it floods the packet out all paths except the one that leads back to the packet source. This technique allows a data stream to reach all LANs. If a router is attached to a set of LANs that do not want to receive a particular multicast group, the router can send a prune message back up the distribution tree to stop subsequent packets from traveling where there are no members.

DVMRP periodically floods packets in order to reach any new hosts that want to receive a particular group. A direct relationship exists between the time it takes for a new receiver to get the data stream and the frequency of flooding.

DVMRP implements its own unicast routing protocol in order to determine which interface leads back to the source of the data stream. This unicast routing protocol is similar to RIP and is based purely on hop counts. As a result, the path that the multicast traffic follows may not be the same as the path that the unicast traffic follows.

NOTE

PIM is supported on Cisco routers and these routers are intelligent enough to interact with external DVMRP neighbors, but do not implement DVMRP natively.

MOSPF

MOSPF is a link-state multicast routing protocol described in RFC 1584, *Multicast Extensions to OSPF*. MOSPF is a protocol that should be used in a single domain or organization. MOSPF is dependent on the use of OSPF as the accompanying unicast routing protocol. In an OSPF/MOSPF network, each router maintains an up-to-date image of the topology of the entire

network. MOSPF, as an extension to the OSPF protocol, includes multicast information within a standard OSPF link-state advertisement (LSA). By using this, an MOSPF router finds out about multicast groups that are active and what router they are active on.

The link-state information is utilized to build multicast distribution trees. The distribution tree has each (source, group) pair listed within and then creates a tree to send to active sources. As with OSPF, any link-state change requires a re-computation of the tree.

NOTE MOSPF is not supported on Cisco routers.

PIMDM

PIMDM is similar to DVMRP. This protocol is best suited when there are “dense” members of multicast groups. PIM uses flooding as a mechanism to reach all routers in the network and then prunes those routers that do not support members of that particular multicast group.

The conditions under which PIMDM is most useful are

- 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.

Two Internet standards-track drafts describe PIM, a multicast protocol that can be used in conjunction with all unicast IP routing protocols. These documents are draft-ietf-idmr-pim-arch-05.txt, *Protocol-Independent Multicast (PIM): Motivation and Architecture* and *Protocol-Independent Multicast (PIM): Protocol Specification*. PIM is also discussed in RFC 2362.

NOTE PIMDM is supported on Cisco routers.

Sparse Mode Routing Protocols

The other method of multicast routing is based on a sparse distribution of multicast group members. Because the multicast group members are located sparsely throughout the network, taking the previous approach of flooding would be a waste of bandwidth. Therefore, employing a more efficient method to accomplish multicast routing becomes necessary.

Sparse mode multicast routing protocols use the assumption that explicit requests are used to join a multicast distribution.

Sparse mode routing protocols include the following:

- Core-Based Trees (CBTs)
- Protocol Independent Multicast Sparse Mode (PIMSM)

Sparse mode protocols are widely used in WAN environments, as opposed to the campus, largely because of the notion that only a few routers are involved.

CBT

CBT was originally described in RFC 2201. Since then, the CBT protocol has been updated in RFC 2189. Unfortunately, CBTv2 isn't backward compatible. In general, neither CBTv1 nor CBTv2 has been widely implemented. The CBT protocol constructs a single tree that is shared by all members of the group. Multicast traffic for the entire group is sent and received over the same tree, regardless of the source. The use of a shared tree can lighten the load on individual routers relative to the amount of multicast routing information stored.

CBT has a core router that is used to construct the tree. When routers are ready to join the tree, they send a join message to the core router. When the core router sends a reply, it travels the reverse path, thereby forming a branch of the tree. Because the CBT join request has a TTL of 1 set, CBT routers in the network forward the message hop by hop until the core is reached or until a CBT router that is already on the shared tree is reached.

PIMSM

PIMSM is used in those environments where the number of receivers are "sparse," hence the name. PIMSM can also be used when multicast traffic is sporadic.

Because of this, a different means of determining the status of multicast members is used. Because the number is relatively small, or "sparse," it makes more sense to institute a proxy of sorts, commonly called the *rendezvous point* (RP). So instead of widespread flooding, the host receiver (or sender) must register with the RP. In short, to make any multicast traffic flow under PIMSM, the host must register with the RP.

PIM is a flexible protocol in that some multicast groups can be dense mode and can coexist together with other groups that might be sparse mode.

Foundation Summary

The Foundation Summary is a collection of tables and figures that provides a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final preparation before the exam, these tables and figures will hopefully be a convenient way to review the day before the exam.

Table 10-6 *Multimedia Traffic Types*

Traffic Type	Characteristic
Unicast	One-to-one, carried out multiple times.
Broadcast	One-to-many, you get it whether you want it or not.
Multicast	One-to-subscribed group, most efficient.

Table 10-7 *Well-Known Class D Addresses*

Well-Known Class D Address	Purpose
224.0.0.1	All hosts on a subnet
224.0.0.2	All routers on a subnet
224.0.0.4	All Distance Vector Multicast Routing Protocol (DVMRP) routers
224.0.0.5	All Open Shortest Path First (OSPF) routers
224.0.0.6	All OSPF designated routers
224.0.0.9	All Routing Information Protocol, version 2 (RIP-2) routers
224.0.0.13	All Protocol Independent Multicast (PIM) routers

Table 10-8 *Multicast Routing: Types of Distribution Trees*

Type of Distribution Tree	Characteristic
Source-specific	Most efficient—most direct path from the sender to each receiver
Shared	Single multicast tree—one path for all

Table 10-9 *Multicast Routing Protocols*

Protocol	Characteristic
DVMRP	Reverse Path Flooding.
MOSPF	Multicast OSPF. Link-state routing protocol.
PIM	Protocol Independent Multicast. Two types: Sparse Mode and Dense Mode.
CBT	Core router used to construct a tree.

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; they are designed, however, to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to the questions can be found in Appendix A, on page 477.

- 1 Name the three types of traffic available in today’s multimedia environment.

- 2 What Layer 4 protocol is used to carry multicast traffic?

- 3 What Class of IP address is used in a multicast environment?

- 4 Describe the makeup of the Class D multicast address by octet or bits.

- 5 What is the name of the protocol used to report their multicast group membership with neighboring multicast routers?

6 What is the special name assigned to the one multicast router that performs host membership queries to determine which groups have members?

7 What does a host send to the all-router group address of 224.1.1.1 to join a group?

8 Which type of routing involves transmitting packets from one source to one source?

9 Define a distribution tree.

10 Name the two types of distribution trees.

11 Name the three types of dense mode routing protocols.

12 Name the two types of sparse mode routing protocols.

13 Which multicast routing protocol is widely used on the MBONE?

14 Name three characteristics of IP multicasting.

15 Certain traditional routing protocols use multicasts to carry routing information. Name one routing protocol and the multicast address it uses.

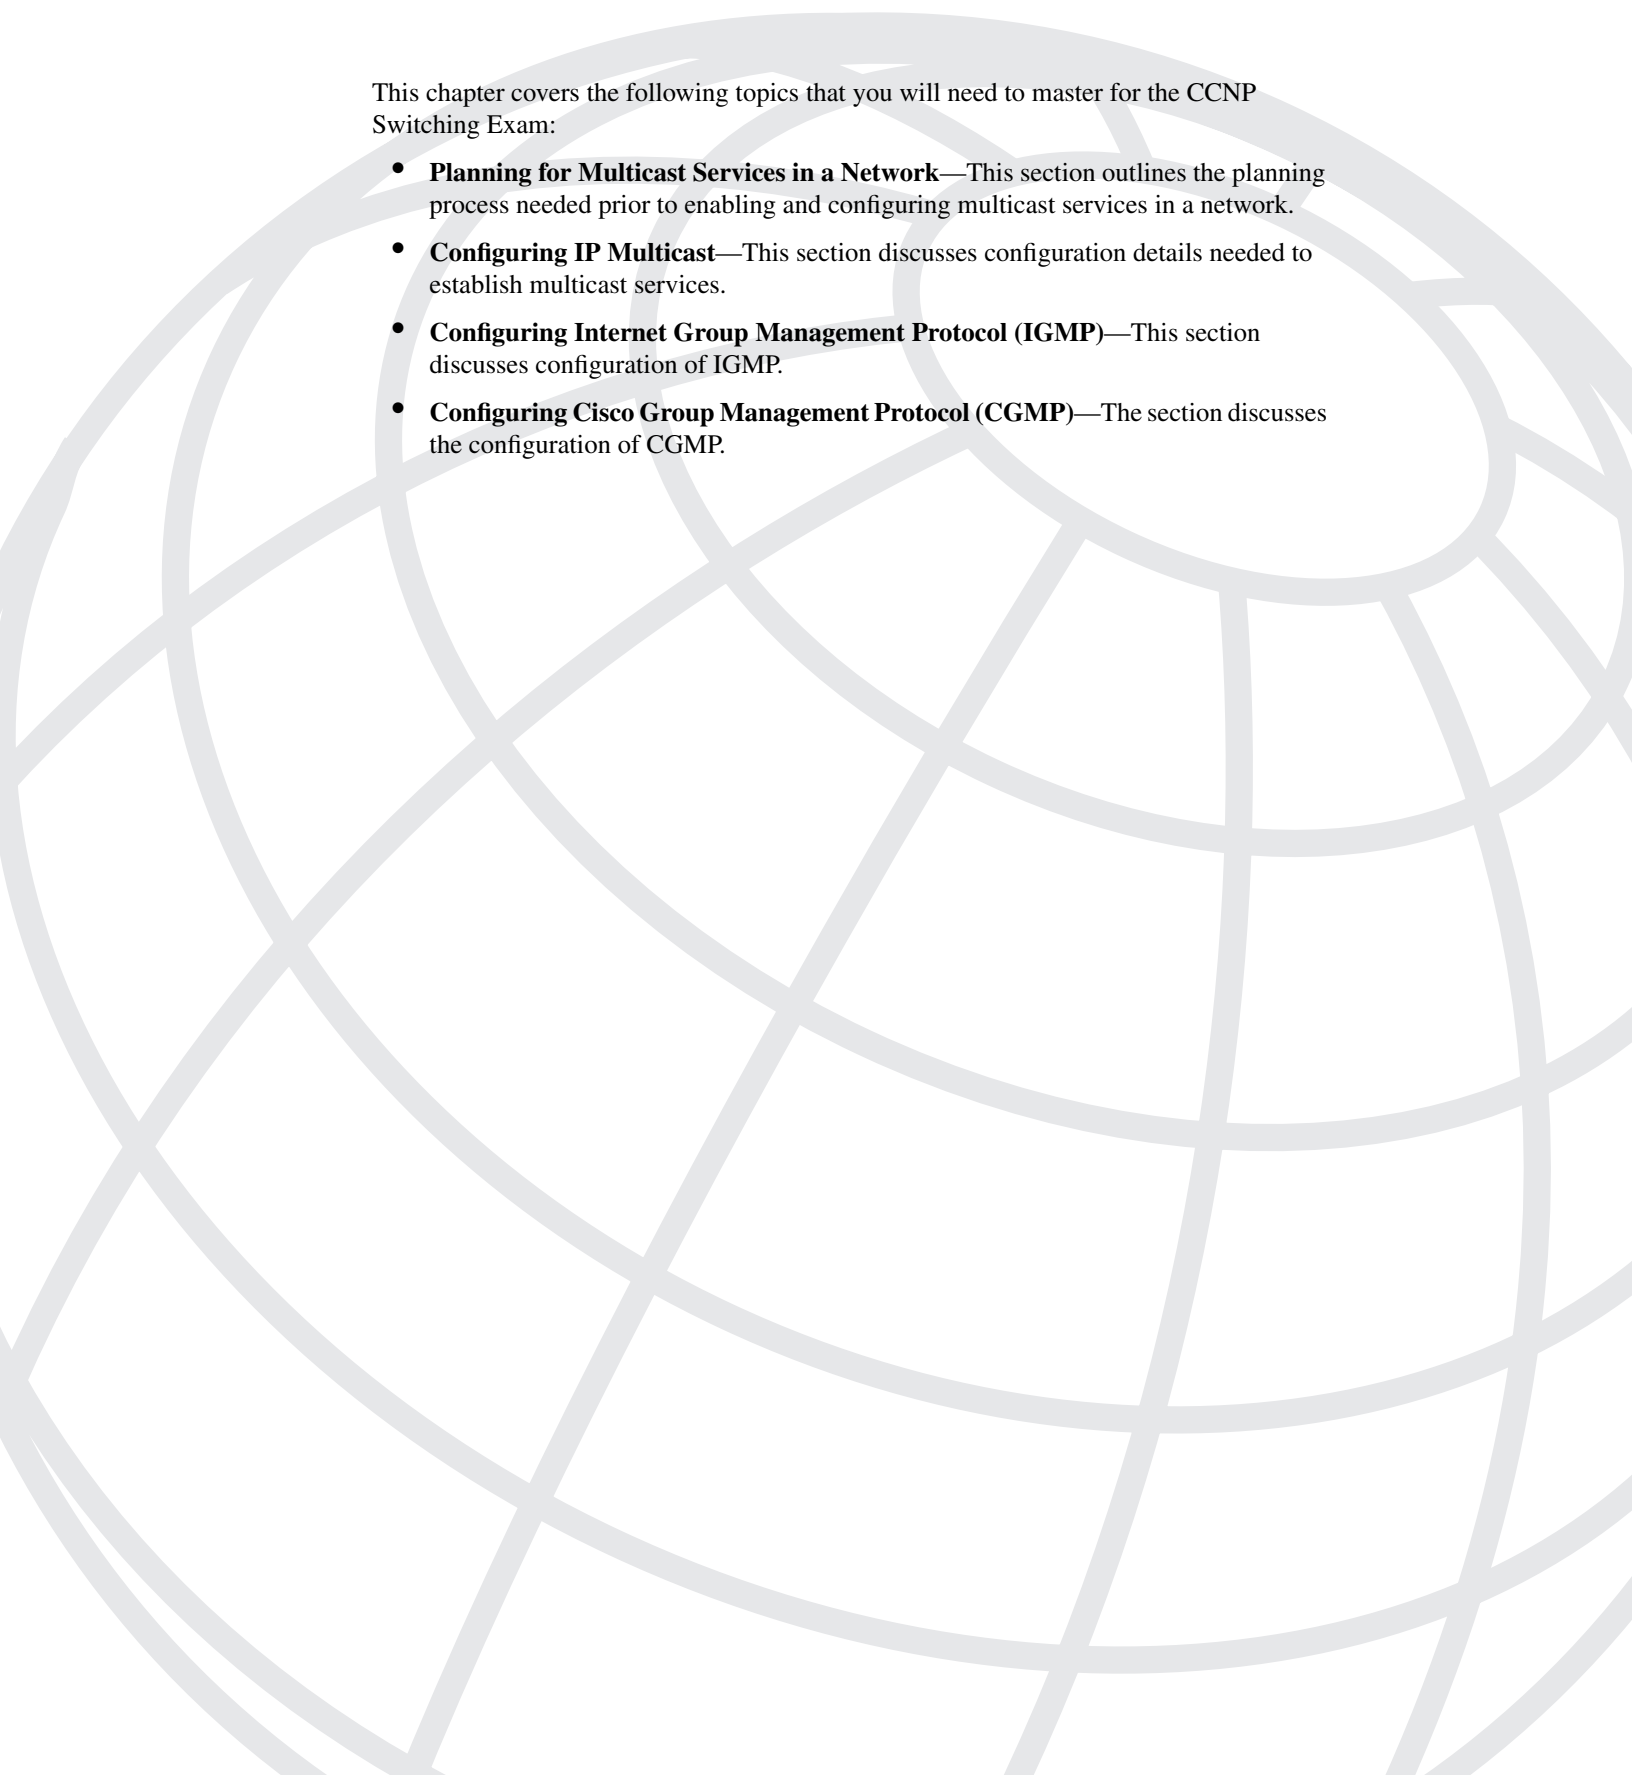

16 What is the name of the Cisco-specific protocol that is used with routers and switches to configure the multicast forwarding table to represent group membership?

17 What is the algorithm used in a source-specific distribution tree?

18 What is used to manage the scope of multicast delivery?

19 What two characteristics describe when PIMSM is most useful?

20 MOSPF is best suited to which type of environment?



This chapter covers the following topics that you will need to master for the CCNP Switching Exam:

- **Planning for Multicast Services in a Network**—This section outlines the planning process needed prior to enabling and configuring multicast services in a network.
- **Configuring IP Multicast**—This section discusses configuration details needed to establish multicast services.
- **Configuring Internet Group Management Protocol (IGMP)**—This section discusses configuration of IGMP.
- **Configuring Cisco Group Management Protocol (CGMP)**—The section discusses the configuration of CGMP.

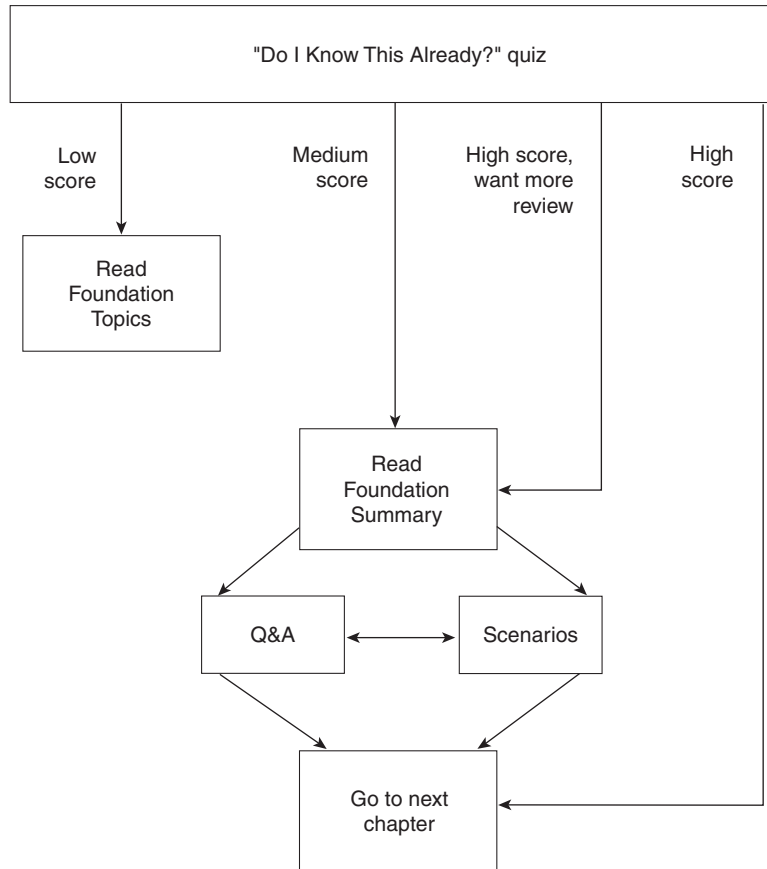
Configuring Multicast Networks

This chapter describes how to configure basic multicast networks. You can find a more complete description of the IP multicast routing commands used in this chapter on Cisco Connection Online (CCO) at www.cisco.com. The information in this chapter builds on that covered in Chapter 10, “Multicasts.”

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and answers for all your work with this book in one place for easy reference.
- Take the “Do I Know This Already?” quiz and write down your answers. Studies show retention is significantly increased through writing facts and concepts down, even if you never look at the information again.
- Use the diagram in Figure 11-1 to guide you to the next step.

Figure 11-1 *How to Use This Chapter*

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The quiz helps you make good choices of how to spend your limited study time. The quiz is divided into four sections that correspond to the four major headings in the “Foundation Topics” section of this chapter. Use the scoresheet in Table 11-1 to record your score.

Table 11-1 *Scoresheet for Quiz and Quizlets*

Quizlet Number	Foundation Topics Section Covering These Questions	Questions	Score
1	Multicast Planning	1–2	
2	Configuring Multicasts	3–4	
3	Configuring IGMP	5–6	
4	Configuring CGMP	7–9	
All questions		1–9	

1 Which Internet Request for Comments (RFC) deals with multicasts?

2 What is the name of the industry standard protocol that deals with multicast groups? The Cisco proprietary protocol?

3 What command enables multicast routing on a Cisco router?

4 What command is issued to enable PIM in sparse mode?

5 What is the default type of IGMP used in a Cisco router?

6 What command would you use to display all multicast packets received and transmitted on a router?

7 What is the status of CGMP in default mode on a Cisco router?

8 How is CGMP enabled on a Cisco router? On a Catalyst switch?

9 What is the purpose of CGMP leave?

You can find the answers to the “Do I Know This Already?” quiz in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Section,” on page 477. The suggested choices for your next step are as follows:

- **4 or fewer overall score**—Read the chapter, including the “Foundation Topics,” the “Foundation Summary,” Q&A, and scenarios at the end of the chapter.
- **6–7 overall score**—Begin with the “Foundation Summary” and then go to the Q&A and scenarios at the end of the chapter.
- **8 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” and then go to the Q&A and scenarios at the end of the chapter. Otherwise, you are ready to move to the next chapter.

Foundation Topics


Planning for Multicast Services in a Network

So what's needed to deploy IP multicast? A lot of patience and planning is needed. OK, so maybe patience and planning aren't the only things needed, but they help. All devices in the network should be able to interpret multicasts. What devices are important? How about routers, switches, and servers, for starters. All your clients receiving multicasts might be helpful. The point here is that the network interface cards (NICs) should support multicasts.

The use of switches is generally a requirement for today's high-speed, high-bandwidth LAN environment. However, not all switches are up to the task of performing effective multicasting. The switch backplane or switching fabric must allow a high percentage of multicast traffic to be passed. So, when shopping for a switch architecture, be aware of the need for a robust switching fabric. The Catalyst switch architecture is designed for efficient multicasting.

In enterprise campus networks, the use of Layer 2 switches is not the end solution. A Layer 3 routing device, such as a router switch module or a full blown external router, is also needed. Figure 11-2 illustrates the network devices typically needed to support a network with multicast services.

Figure 11-2 *IP Multicast Components*



Cisco routing devices provide support for multicasts through the use of multicast routing protocols such as Distance Vector Multicast Routing Protocol (DVMRP), Protocol Independent Multicast (PIM), and Core-Based Trees (CBT), as discussed in detail in Chapter 10.

Configuring IP Multicast

IP multicast and the task configuring it are somewhat advanced topics. Fortunately, the material on the CCNP Switching exam covers only the basics of configuring IP multicast. We will do the same here, but will list as optional a few of the advanced tasks.

The two basic tasks in enabling multicast are

- Enabling IP multicast routing
- Enabling PIM on an interface

Advanced tasks are optional and include the following:

- Configuring a rendezvous point
- Configuring the Time To Live (TTL) threshold
- Debugging IP Multicast
- Configuring Internet Group Management Protocol (IGMP)
- Enabling Cisco Group Management Protocol (CGMP)

Enabling IP Multicast Routing

Enabling IP multicast routing allows the Cisco IOS software to forward multicast packets. Much like enabling other routing protocols, you must make an entry in global configuration mode to turn this on for the entire router. Then, using interface commands, you can turn on various modes of multicast routing using only specific interfaces. To enable IP multicast routing on the router, enter the following command in global configuration mode.

```
Router(config)#ip multicast-routing
```

To disable IP multicast routing, enter the **no ip multicast-routing** command. By default, multicast routing is disabled on an interface.

Enabling PIM on an Interface

When you enable multicast routing on a route processor or router, it is processed on an individual interface basis. Enabling each individual interface used with a specific multicast routing protocol is necessary.

The command to enable PIM on an interface is

```
DallasR1>(config-if)#ip pim {dense-mode | sparse-mode | sparse-dense-mode}
```

The options for this command are defined as follows.

- **dense-mode**—Enables dense mode of operation. Dense mode is used when all routers in the network will need to distribute multicast traffic for each multicast group.
- **sparse-mode**—Enables sparse mode of operation. Sparse mode is used when relatively few routers in the network will be involved in each multicast.

- **sparse-dense-mode**—The interface is treated in the mode in which the group operates. If the group is operating in sparse mode, the interface does as well. If the group is operating in dense mode, then the interface does.

PIM can be implemented in three different modes: sparse mode, dense mode, or a combination called sparse-dense mode, as specified by the **ip pim** options previously listed. The mode used depends on the “density” of the hosts in an area.

Enabling PIM in Dense Mode

In dense mode, the source of the multicast and its receivers are all in close proximity, such as a campus environment. Dense mode should be used when bandwidth is plentiful, due to periodic flooding. Pruning is used in this environment to avoid unnecessary multicast packets being flooded to a router with no directly connected neighbors.

Outgoing interface lists (oilst) are used to display those interfaces that meet one of the following criteria:

- An interface is heard from a PIM neighbor.
- A host has joined a group that uses that interface.
- A particular interface has been manually configured to join a group.

Enabling PIM in Sparse Mode

Sparse mode should be configured when there are few multicast hosts, hence the term sparse mode. Because they are sparse, a rendezvous point is needed in this scheme.

Sparse mode protocols are more appropriate for large internetworks where dense mode protocols would waste bandwidth by flooding packets to all parts in the network and then prune back all unwanted connections. Sparse mode protocols use explicit join messages to set up distribution trees. Tree state is then set up only on routers on the distribution tree, and data packets are only forwarded to LANs that have hosts who join the group. Dense mode protocols build only source-distribution trees—they determine the location of receivers by flooding data throughout the network and then pruning off the receiverless branches. PIM sparse mode can be used for any combination of sources and receivers, whether densely or sparsely populated.

Rendezvous points (RPs) are configured and act as a sort of proxy for multicast hosts. Multicast senders use the RP to announce the hosts, and multicast receivers use RP to learn about other hosts.

Enabling PIM in Sparse-Dense Mode

To configure PIM sparse-dense mode, use the following commands on all PIM routers inside the PIM domain, beginning in global configuration mode:

```
Router#ip multicast-routing
Router#interface type number
Router#ip pim sparse-dense-mode
```

Verifying PIM Configuration

You can verify the PIM configuration that you just completed by typing the following command:

```
DallasR1>#show ip pim interface [type number] [count]
```

where *type* is the interface type and the *number* is the number of the interface. The **count** parameter represents the number of packets received and sent out the interface.

Selecting a Designated Router

In a normal functioning multicast network, PIM queries are sent periodically to discover other routers in the network running PIM. For multi-access networks such as Ethernet, PIM queries are sent to the well-known multicast address of 224.0.0.2, otherwise known as “all routers.” In a multi-access network, a designated router is elected. If this process sounds familiar to you, you might remember it as a process also used in OSPF networks.

The election process is also very similar to OSPF, which uses the highest IP address received in PIM query messages from a network device’s neighbors. If no PIM queries are received after a given time period, the election process for Designated Router runs again.

When running PIM in sparse mode, the designated router is responsible for sending multicast join messages to the RP on behalf of host computers on the network. No designated router exists when running PIM in dense mode.

Displaying PIM Neighbors

As is typical on a Cisco router, a **show** command will display PIM neighbors:

```
DallasR1>#show ip pim neighbor type number
```

where *type* is the type of interface and *number* is the number of the interface. On a Route Switch Module, the *type* would be a particular VLAN and the *number* a particular VLAN number (VLAN 10 is an example). Example 11-1 shows some typical output from the **show ip pim neighbor** command.

Example 11-1 show ip pim neighbor Command Output Displays Information About PIM Neighbors

```
DallasR1>show ip pim neighbor
PIM Neighbor Table
Neighbor Address Interface    Uptime    Expires    Mode
192.168.100.1 FastEthernet0 1w2d      00:01:33  Dense
192.168.101.1 FastEthernet0 1w4d      00:01:21  Dense (DR)
192.168.21.1  FastEthernet0 1w1d      00:01:16  Dense
192.168.57.1  FastEthernet0 2w5d      00:01:33  Dense
192.168.99.1  Serial0.4      22:00:01  00:01:08  Dense
192.168.78.1  Serial0.1      22:00:15  00:01:19  Dense
```

Most of the information in Example 11-1 is self-explanatory. The **Expires** parameter is defined as the length of time that the PIM neighbor is considered active. Note that the designated router is shown out to the right as (DR).

Configuring a Rendezvous Point

One of the features that you have to configure if you use PIM in sparse mode is a Rendezvous Point (RP).

The routers learn that they are RPs automatically. RPs are used by multicast senders in a sparse mode environment to announce their existence. Through the destination, receivers learn about new senders.

Multi-RP environments can be configured for any given multicast group. One term used in the description of RPs is *leaf routers*. Leaf routers are either directly connected to a multicast group member or to a sender.

The RP address is used by first-hop routers to send PIM register messages on behalf of a host sending a packet to the group. The RP address is also used by last-hop routers to send PIM join/prune messages to the RP to inform about group membership. The RP address is configured only on the first-hop and last-hop routers (leaf routers).

A PIM router can be configured as an RP for more than one group. A group can also have more than one RP configured. An access list is used to determine the groups for which the router is an RP.

Although a group can have more than one RP, only one RP address is used per group at any given time. You can configure multiple redundant RPs, but only one is used.

To configure the address of the RP, use the following command in global configuration mode:

```
DallasR1>#ip pim rp-address ip-address [group-access-list-number][override]
```

As usual, you can disable the RP address by using the **no** form of the command.

Auto-RP

Auto-RP is a feature that automates the distribution of group-to-RP mappings in a PIM network. This feature has the following benefits:

- Easy-to-use multiple RPs within a network to serve different group ranges.
- Allows load splitting among different RPs and arrangement of RPs according to the location of group participants.
- Avoids inconsistent, manual RP configurations that can cause connectivity problems.

Multiple RPs can be used to serve different group ranges or serve as hot backups of each other. To make Auto-RP work, you must designate a router as an RP Mapping Agent that receives the RP-announcement messages from the RPs and arbitrates conflicts. The RP mapping agent then sends the consistent group-to-RP mappings to all other routers. Thus, all routers automatically discover which RP to use for the groups they support.

One way to start is to place (preserve) the default RP for all global groups at or near the border router of your routing domain, while placing another RP in a more centrally located router for all local groups using the administratively scoped addresses (239.x.x.x).

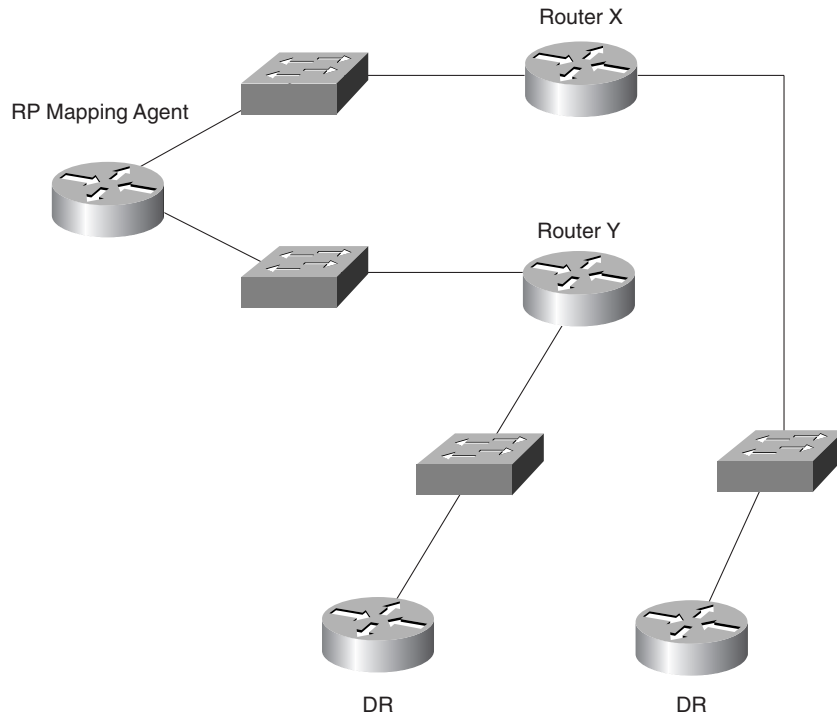
As illustrated by Figure 11-3, Router X announces to the Mapping Agent that it would like to be the RP for Group X and Router Y announces to the Mapping Agent that it would like to be the RP for Group Y. The Mapping Agent then passes this information to the Designated Routers (DR).

NOTE

If you configure PIM in sparse mode or sparse-dense mode and do not configure Auto-RP, you must statically configure an RP as described in the section “Configuring a Rendezvous Point” later in this chapter.

If you are setting up Auto-RP in a new internetwork, you do not need a default RP. For the initial deployment of Auto-RP into an existing sparse mode cloud, use the following guidelines to allow minimal disruption to the existing multicast infrastructure:

- Step 1** Choose a default RP.
- Step 2** Announce the RP and the group range it serves.
- Step 3** Assign the RP mapping agent.
- Step 4** Verify the group-to-RP mapping.

Figure 11-3 *Rendezvous Point Announcement*

Follow the same steps to set up Auto-RP in a new internetwork, but omit Step 1 (default RPs are not required in a new internetwork). The sections that follow describe each of these steps in further detail.

Step 1: Choose a Default RP

Sparse mode environments need a default RP; sparse-dense mode environments do not. If you have sparse-dense mode configured everywhere, you do not need to choose a default RP.

Adding Auto-RP to a sparse mode cloud requires a default RP. In an existing PIM sparse mode region, you must define at least one RP across the network that has good connectivity and availability. By “existing PIM sparse mode region,” we mean that the **ip pim rp-address** command is already configured on all routers in this network.

Use that RP for the global groups (for example, 224.x.x.x and other global groups). You do not need to reconfigure the group address range that RP serves. RPs discovered dynamically through Auto-RP take precedence over statically configured RPs. Using a second RP for the local groups is desirable.

Step 2: Announce the RP and the Group Range It Serves

Find another router to serve as the RP for the local groups. The RP mapping agent can double as an RP itself. Assign the whole range of 239.x.x.x to that RP or assign a subrange of that (for example, 239.2.x.x).

To designate that a router is the RP, perform the following task in global configuration mode:

```
DallasR1>#ip pim send-rp-announce type number scope ttl group-list access-list-number
```

To change the group ranges this RP optimally serves in the future, change the announcement setting on the RP. If the change is valid, all other routers automatically adopt the new group-to-RP mapping.

The following example advertises the IP address of Ethernet 0 as the RP for the administratively scoped groups.

Example 11-2 Advertising the RP Address for Administratively Scoped Groups

```
DallasR1>#ip pim send-rp-announce ethernet0 scope 16 group-list 1
DallasR1>#access-list 1 permit 226.0.0.0 0.255.255.255
```

Step 3: Assign the RP Mapping Agent

The RP mapping agent is the router that sends the authoritative Discovery packets telling other routers which group-to-RP mapping to use. Such a role is necessary in the event of conflicts (such as overlapping group-to-RP ranges).

Find a router whose connectivity is not likely to be interrupted and assign it the role of RP mapping agent. All routers within TTL number of hops from the source router receive the Auto-RP Discovery messages. To assign the role of RP mapping agent in that router, enter the following command in global configuration mode:

```
DallasR1>#ip pim send-rp-discovery scope ttl
```

Step 4: Verify the Group-to-RP Mapping

To see if the group-to-RP mapping has arrived, issue one of the following commands in EXEC mode on the designated routers:

```
DallasR1>#show ip pim rp mapping
DallasR1>#show ip pim rp [group-name | group-address] [mapping]
```

The first command is a generic show all mapping, but the second one allows you to specify either the individual group or group address.

Configuring Time-To-Live

Time-To-Live (TTL) works in this situation just like it does in other routing environments. Simply put, any packet that comes in with a higher TTL than the one configured will be forwarded and the TTL value decreased by one. TTL is expressed as a number that signifies the number of router hops. The default value of TTL is 0. A TTL of zero means that every packet is forwarded. Configuring the TTL limit is done on a per-interface basis. To configure a value other than the default, type the following in interface mode:

```
DallasR1>(config-if)#ip multicast ttl-threshold ttl-value
```

Example 11-3 demonstrates how to configure the TTL limit using the **ip multicast ttl-threshold** command.

Example 11-3 ip multicast ttl-threshold *Configures the TTL Limit on a Router Interface*

```
DallasR1>#show running-config
hostname DallasR1>
!
!
ip multicast routing
!
interface VLAN100
ip pim sparse-dense-mode
ip multicast ttl-threshold 16
```

Debugging Multicast

Many potential commands can be used when debugging multicast. We will go over a couple of the most important here. You can find documentation on other debugging commands at CCO (www.cisco.com).

The first command of significance, **show ip pim neighbor**, displays the PIM neighbor table as demonstrated earlier in Example 11-1.

Another relevant command, **show ip mroute**, shows the entries in the multicast routing table. The general syntax for this command is as follows:

```
DallasR1> show ip mroute [group-name | group-address] [source] [summary] [count]
[active kbps]
```

Example 11-4 demonstrates output generated by the **show ip mroute** command.

Routing entries are categorized as (S,G). S is the source and G is the destination multicast group. The other kind of entry is a (*,G), which is an entry made by the Designated Router on behalf of a host that wants to join a group. The G is the group multicast address.

Another useful debugging tool is to capture all multicast packets to the console screen through the **debug** facility. As is the case when using any **debug** command, you must exercise caution because this is a CPU-intensive task.

Example 11-4 `show ip mroute` Displays the Entries in the Multicast Routing Table

```
DallasR1>#RSM114#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, J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, Next-Hop, State/Mode

(*, 224.1.150.1), 00:30:17/00:02:59, RP 0.0.0.0, flags: D
(172.16.50.1/32, 224.1.150.1), 00:30:17/00:02:58, flags: CT

(*, 172.16.53.1), 2d9h/01:00:00, RP 0.0.0.0 flags: DCT
```

To log all IP multicast packets received and transmitted by a router, enter the following command in EXEC mode:

```
DallasR1>#debug ip mpacket [detail] [access-list] [group]
```

Example 11-5 shows sample output generated by issuing this command.

Example 11-5 `debug ip mpacket` Command Output Displays IP Multicast Packets Received and Transmitted by a Router

```
Router#debug ip mpacket
IP multicast packets debugging is on

2d03h : IP: s=192.168.24.1 (Vlan21) d=224.4.204.15 len 60, mforward
2d03h : IP: s=192.168.24.1 (Vlan22) d=224.2.204.15 len 60, mforward
2d03h : IP: s=192.168.24.1 (Vlan21) d=224.0.1.18 len 65, mforward
```

Configuring Internet Group Management Protocol (IGMP)

Internet Group Management Protocol (IGMP) is an important part of IP that must be supported by all multicast hosts on a network. Multicast routers use IGMP to keep track of multicast hosts on a network.

Although two versions of IGMP are available, version 1 and version 2, IGMP version 2 is the default in all Cisco routers running IOS Release 11.3(2)T and later.

To configure the multicast router to join a particular multicast group, enter the following command in the relevant interface configuration mode:

```
DallasR1>(config-if)#ip igmp join-group group-address
```

where *group-address* is the multicast address of the group.

One way to verify the multicast group is to issue a **ping** to the specified group address. All configured routers should respond to the **ping**.

Configuring Cisco Group Management Protocol (CGMP)

Cisco Group Management Protocol (CGMP) runs on Catalyst switches and Cisco routers. CGMP is used in conjunction with IGMP running on Cisco routers to determine forwarding information. CGMP messages are sent to the well-known multicast MAC address of 01-00-0c-dd-dd-dd. Catalyst switches discover CGMP routers via a hello mechanism.

CGMP is capable of operating correctly only when it is working in conjunction with a router. The router is needed to detect IGMP packets and communicate with the CGMP-enabled switch. The switch receives CGMP packets created by the router.

NOTE CGMP is disabled by default on the XDI interface of Cisco Catalyst switches.

This section covers the configuration of CGMP on a Cisco router and also on a Catalyst switch. To configure CGMP on a Cisco router, enter the following command for a particular interface:

```
DallasR1>(config-if)#ip cgmp
```

Upon hitting the **Enter** key, a CGMP Join message is sent.

Example 11-6 demonstrates a router configuration with CGMP for a router (RouterLA1) connected to a Catalyst switch (CatalystLA1).

Example 11-6 Router Configuration with CGMP

```
RouterLA1>#show running-config
hostname RouterLA1>
!
ip multicast routing
interface FastEthernet 3/1.0
ip address 192.168.3.15 255.255.255.0
ip pim dense-mode
ip cgmp
!
```

The command to disable CGMP is the same as any within the Cisco IOS; that is, use **no** in front of the previous command:

```
DallasR1>(config-if)#no ip cgmp
```

Disabling CGMP triggers a CGMP Leave message similar to the CGMP Join message triggered when enabling CGMP.

You also need to configure the switch in a multicast environment. Enabling CGMP on a Catalyst switch requires the following command in enable mode:

```
DallasS1>(enable)set cgmp enable
```

Running **show config** on the Catalyst switch reveals that CGMP is enabled as demonstrated in Example 11-7.

Example 11-7 *Verifying CGMP Configuration on a Catalyst Switch*

```
CatalystLA1>(enable)show config
set prompt CatalystLA1>
set interface sc0 192.168.1.1 255.255.255.0
set cgmp enable
```

As is typical syntax on a Catalyst switch, entering the **set cgmp disable** command turns CGMP off.

Configuring CGMP Leave

In some cases, you may want multicast group to be removed from the forwarding tables, freeing up bandwidth. The command to accomplish this, called CGMP leave, is as follows:

```
Dallas_SW(enable)set cgmp leave
```

A multicast router sends out group queries periodically. In a normal participating mode, the multicast hosts would send a reply to these queries. If, after a given number of queries no response is given by any members of a group, that group is then eligible to be pruned from the forwarding tables of the switch.

Foundation Summary

The Foundation Summary is a collection of tables and figures that provides a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final prep before the exam, these tables and figures will hopefully be a convenient way to review the day before the exam.

Table 11-2 *Multicast Switch Commands*

Command	Description
set cgmp enable	Enables CGMP processing on the switch
show cgmp	Displays current CGMP settings
set cgmp leave	Allows pruning of multicast groups
show config	Displays the configuration of the switch

Table 11-3 *Multicast Router Commands*

Command	Description
ip cgmp	Enables support of CGMP
ip multicast-routing	Enables IP multicast routing
ip pim sparse-dense-mode	Enables the PIM protocol on an interface
ip pim rp-address	Specifies an RP address to enable PIM sparse mode
ip igmp join-group group-address	Enables router to join IGMP group
show ip pim neighbor	Displays PIM neighbors
debug ip mpacket [detail] [access-list] [group]	Enables debugging of multicast packets
show ip mroute	Displays the contents of the multicast routing table
show ip mroute host-ip group-ip	Displays the contents of the multicast routing table for a specified IP host address and specified IPmulticast address
show run	Displays the current configuration
ip multicast ttl-threshold ttl-value	Configures a TTL threshold

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however, they are designed to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to these questions can be found in Appendix A, on page 477.

- 1 Which Internet Request for Comments (RFC) deals with multicasts?

- 2 What is the name of the industry standard protocol that deals with multicast groups? The Cisco proprietary protocol?

- 3 What command enables multicast routing on a Cisco router?

- 4 What command is issued to enable PIM in Sparse Mode?

- 5 What is the default type of IGMP used in a Cisco router?

6 What command would you use to display all multicast packets received and transmitted on a router?

7 What is the status of CGMP in default mode on a Cisco router?

8 How is CGMP enabled on a Cisco router? On a Catalyst switch?

9 What is the purpose of CGMP leave?

10 Name one multicast routing protocol that is used by Cisco routers for router-to-router communication?

11 What is the MBONE? What multicast routing protocol is used throughout the MBONE?

12 What are the two basic tasks associated with configuring IP multicasts on a Cisco router?

13 What process is used within PIM router-to-router communication that is also used in an OSPF network over a multi-access network?

14 Name one of the three reasons why an interface would be placed in the oolist for a multicast group.

15 Define Rendezvous Point and determine under what circumstances Rendezvous Point would be used.

16 What troubleshooting command can be used to determine which routers belong to an IGMP group?

17 Assume that you have a router named Router1 connected to other networks from a wide area perspective and a Catalyst switch connected to Router1 on port Ethernet 0. What commands would be required to enable basic CGMP in this network?


18 What happens when a CGMP-enabled router receives an IGMP control packet?

Scenarios

Scenario 11-1

In this scenario, depicted in Figure 11-4, we have an IP streaming video application being downloaded from a satellite dish. The video server is taking the video stream, buffering it, and converting it to a multicast application. From there the video stream is being sent to two different multicast groups. Given this information, complete the tasks that follow to make this application work.

Figure 11-4 Scenario 11-1 Network Setup



- 1 Configure IP Multicast support in the router.
- 2 Configure IP PIM in sparse mode on the router.
- 3 Display PIM information about interface VLAN10.
- 4 Display the PIM neighbor tables.
- 5 Configure an RP with address 172.16.1.2.
- 6 Configure CGMP on both the router and the switch.

Scenarios Answers

Scenario 11-1 Answers

- 1 To configure IP multicast on the router, enter the following command at the global configuration prompt:

```
Router(config)#ip multicast-routing
```
- 2 To configure PIM in sparse mode, enter the following command at the interface prompt:

```
Router (config-if)# ip pim sparse-mode
```
- 3 To display PIM information about interface VLAN10, enter the following command at the prompt:

```
Router#show ip pim interface
```
- 4 To display the PIM neighbor tables, enter the following command at the prompt:


```
Router#show ip pim neighbor
```
- 5 To configure an RP, enter the following command while in global configuration mode:

```
Router>#ip pim rp-address 172.16.1.2
```
- 6 To configure CGMP on the router, enter the following command:

```
Router>#ip cgmp
```

CGMP is normally disabled on a Cisco router. The configuration of a switch is also very simple. The configuration is as follows:

```
Switch>(enable)set cgmp enable
```

This chapter covers the following topics that you will need to master for the CCNP Switching Exam:

- **Access Policies**—This section discusses the need for an access policy. Access policies define who has access to what.
- **Managing Network Devices**—This section discusses the different methods for controlling and managing network devices.
- **Access Layer Policy**—This section discusses the type of policy needed at the access layer of the switch block.
- **Distribution Layer Policy**—This section discusses the type of policy needed at the distribution layer of the switch block.
- **Core Layer Policy**—This section discusses the type of policy needed at the core layer of the switch block.

Controlling Access in the Campus Environment

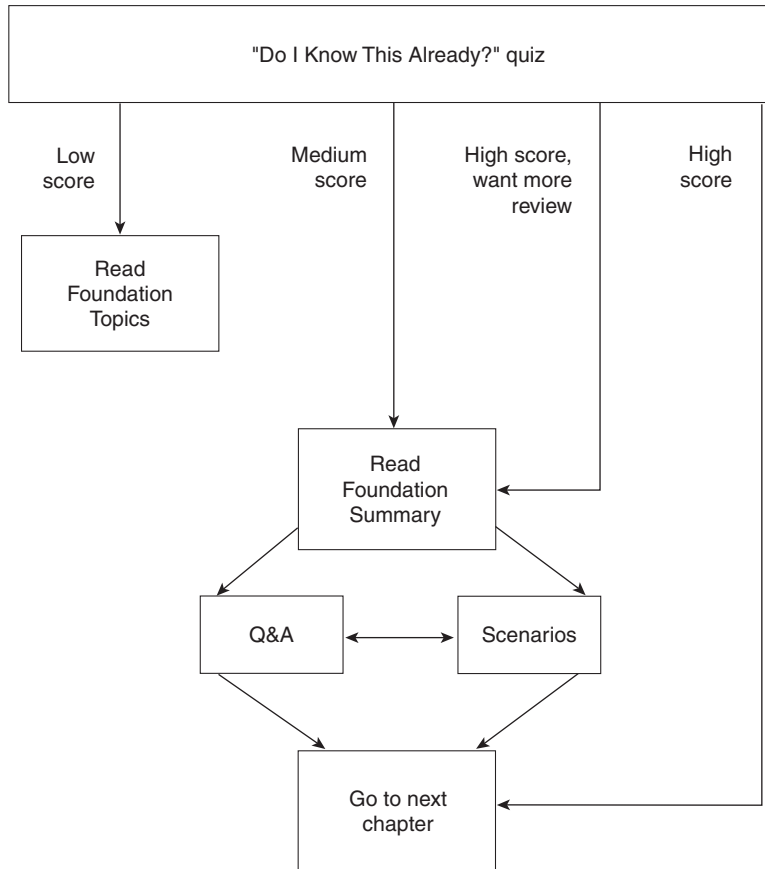
In this age of increased activity of interlopers on the Internet and other networks, the need for access control is greater than ever. This chapter covers some of the preventative measures that can be used in a Cisco campus environment. The first preventative measure involves creating an access policy. The components of an access policy are discussed, followed by the policies of each layer within the campus block. Certain configurations of security measures on Cisco devices are also discussed.

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and answers for all your work with this book in one place for easy reference.
- Take the “Do I Know This Already?” quiz and write down your answers. Studies show retention is significantly increased through writing facts and concepts down, even if you never look at the information again.
- Use the diagram in Figure 12-1 to guide you to the next step.

Figure 12-1 *How to Use This Chapter*



“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The 13-question quiz helps you make good choices of how to spend your limited study time. Use the scoresheet in Table 12-1 to record your score.

Table 12-1 *Scoresheet for Quiz*

Quizlet Number	Foundation Topics Section Covering These Questions	Questions	Score
1	Access Policies	1–2	
2	Managing Network Devices	3–7	
3	Access Layer Policy	8–9	
4	Distribution Layer Policy	10–12	
5	Core Layer Policy	13	
All questions		1–13	

1 Define an access policy.

2 What is the access layer defined as?

3 Is HTTP access normally enabled on a Cisco router? What is the main purpose of using HTTP?

4 Name at least two components relating to controlling access to network devices.

5 What way of accessing a network device requires a password?

6 What feature of the Cisco IOS protects a console connection left unattended?

7 What does the **access-class** command do when applied to a virtual terminal configuration?

8 What VLAN is the default VLAN for a Catalyst switch and why is it a good idea to change this?

9 What does port security do on a Catalyst series switch?

10 What is the range of numerical representation of a standard IP access list? An extended access list?

11 Should a standard or an extended access list be used when filtering a particular host?

12 When implementing route filtering, what type of access list is used—a standard or an extended access list?

13 In general, what type of policies should be implemented in the core layer?

The answers to the “Do I Know This Already?” quiz are found in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A Sections,” on page 477. The suggested choices for your next step are as follows:

- **8 or fewer overall score**—Read the chapter. This reading includes the “Foundation Topics,” the “Foundation Summary,” Q&A, and scenarios at the end of the chapter.
- **9–12 overall score**— Begin with the “Foundation Summary,” and then follow with the Q&A and scenarios at the end of the chapter.
- **13 or more overall score**—If you want more review on these topics, skip to the “Foundation Summary” and then go to the Q&A and scenarios at the end of the chapter. Otherwise, move to the next chapter.

Foundation Topics

Access Policies

Access policies are the defining guidelines that are necessary to create a level of access control. An access policy is a firm's documented standard of network access for the firm's users. Access policies may vary widely just like the business itself. Different size businesses may require a different type of access policy. In general, a corporate network security policy, whether or not it covers access, is designed to protect to the level of the data it is securing.

An access policy may define the following:

- Management and configuration of network devices, including physical security, logical security, and access control.
- The means of controlling users' access to the network through the use of mechanisms such as switch port security and VLAN management.
- Controlling access to distributed and enterprise services.
- Determining the traffic allowed out of a distribution switch and into the core network, as well as how traffic is managed.
- Route filtering to determine the routes that should be seen by the core network—the distribution and access blocks.

In terms of the campus environment, an access policy is designed to police that traffic going to and from the campus. The policy should allow only what is necessary in order to do business. An access policy should also provide a measure of protection to those network devices in the campus.

In Figure 12-2, each layer can—and probably should—have a different access policy, mostly because each layer has a different task associated with it. Some access policies could, however, apply to all devices in the network. Others could be defined individually at each layer.


Figure 12-2 Access Policies for Hierarchical Layers of a Network

Table 12-2 follows up on Figure 12-2 by summarizing the different characteristics and access policies for each hierarchical layer of a given network.

Table 12-2 Access Policy Guidelines

Hierarchical Network Layer	Characteristics/Access Policies
Access Layer	The access layer is the entry point for the users to the campus network. The use of Port Security and passwords are used here to protect the network.
Distribution Layer	The distribution layer carries the bulk of all policy decisions. This layer defines what traffic enters to or from either side, being the core or access layers. This could determine whose traffic is going where or specific paths. Or, it could be advertising routes that traverse from the access layer to the core. Many of the network device access policies could be the same as the access layer.
Core Layer	The core layer is a high bandwidth backbone, which is capable of handling the aggregate traffic of all the other devices in the network. There really shouldn't be any policy at this layer because the job here is to pass traffic at a high speed. Any policy implemented would slow down the flow.

Managing Network Devices

The policy to control access to network devices should be one of the first components of the access policy. All devices at every layer in the campus network should have a plan to provide for the following:

- Physical security
- Passwords
- Privilege levels to allow limited access to a network device
- Limiting virtual terminal or Telnet access

Physical Access

Virtually all devices provide a way of gaining control of a given device, assuming that you have physical access to it. That is why defining a physical access policy is so important. If the physical device isn't secured, chances are your network isn't secure either. Therefore, every network device should be secured in some manner.

You can physically secure your network by doing the following:

- Establish a configuration, control, and change management policy for all devices at each of the respective layers.
- Establish a security plan for all physical locations. Include details on physical and link security.
- Provide the proper physical environment. The physical environment should have provisions for locking the room, proper ventilation and temperature controls, and backup power.
- Control direct access to the device. Lock racks when possible and apply passwords to console and auxiliary ports. Disable ports not being used, such as the auxiliary port.
- Secure access to network links. Provide the same type of security for the wiring closet that you would for the physical equipment.

Passwords

There are several different ways to access every Cisco device. Every method of accessing the device should have a password applied to prevent unauthorized access.

Out-of-band management options include the console port and the auxiliary port.

In-band management options include Trivial File Transfer Protocol (TFTP) servers and Simple Network Management Protocol (SNMP)-based network management systems, such as CiscoWorks 2000.

Virtual terminal ports that are used for terminal access and are referred to as *vt*y ports. There are five vty ports by default on each Cisco device. You can create more vty ports if you need to have more than five users accessing a device simultaneously. Example 12-1 demonstrates how you would configure passwords for the console port and the vty ports on a Cisco device.

Example 12-1 *Modifying Console Port Passwords on a Cisco Device*

```
R1(config)#line console 0
R1(config-line)#login
R1(config-line)#password lisbon

R1(config)#enable password bilbao
R1(config)#login local
R1(config)#username student password cisco
```

The **login** option that appears in Example 12-1 indicates where to find the login information. If the login is specified without a keyword, as in the case of the console port, the system will use the line as the login. The user will be prompted for the password of the line itself (in this case, *lisbon*). The other options indicate that the specific user must log in. The keyword after **login** indicates where to find the user information. The **login local** statement indicates that the information will be found locally in the **username student password cisco** statement. Other options include **login authentication** or **login tacacs**. These options indicate that the login information is contained on a centralized authentication server. Centralizing usernames, passwords, and profile information makes maintaining a large number of users or devices easier.

It is recommended that users log in to the system with a username and password rather than having everyone use the password of the line. Having users log in to the device makes it easier to track who has access and when.

By default, passwords are stored in clear text format in the router's configuration. The only exception to this is the enable secret password, which is automatically encrypted. Password encryption can be compromised so it should be used in combination with other methods of security.

NOTE

More information on Terminal Access Controller Access Control System Plus (TACACS+) and other authentication services are covered in the Cisco IOS Security Configuration Guide.

Assigning passwords prevents users from initiating a session with the network device. If the console is left unattended in privileged mode, any user can modify the network device's configuration. A timeout for an unattended session provides additional security. Example 12-2 demonstrates configuring a session timeout for console and vty ports.

Example 12-2 *Configuring Session Timeouts for Console and vty Ports on a Cisco Device*

```
R1(config)#line console 0
R1(config-line)#exec-timeout 5 10
R1(config)#line vty 0 4
R1(config-line)#exec-timeout 5 2
```

NOTE

In Example 12-2, the two numbers (5 and 10) following **exec-timeout** indicate minutes and seconds. These figures should be adequately long enough to do configuration work but short enough to not leave this open for extended periods.

Routers and high-end switches calculate timeouts in minutes. An option is also available to calculate seconds in addition to minutes on routers. The Cisco IOS command-based switches calculate timeouts in seconds.

Privilege Levels

The two default levels of access are user and privileged. The user level allows the user to perform certain commands but does not give them the ability to modify the configuration or perform a debug. At the other end of the spectrum, the privileged level allows the user to issue all commands, including configuration and debug commands.

Cisco IOS provides different levels of privileges for users with the use of the **privilege level** command. This command allows network administrators to provide a more granular set of rights to Cisco network devices.

There are 16 different levels of privilege that can be set, ranging from 0 to 15. Level 1 is the default user EXEC privilege. The highest level, 15, allows the user to have all rights to the device. Level 0 can be used to specify a more limited subset of commands for specific users or lines. For example, you can allow user “guest” to use only the **show users** and **exit** commands.

At other privilege levels, you must specify the commands that the privilege level should be able to complete. Example 12-3 demonstrates the capability to set privilege levels above that of EXEC user but below that of full enable level.

Example 12-3 *Setting Privilege Levels on a Cisco Device*

```
R1(config)#privilege configure level 3 username
R1(config)#privilege configure level 3 copy run start
R1(config)#privilege configure level 3 ping
R1(config)#privilege configure level 3 show run
R1(config)#privilege configure level 3 show
R1(config)#enable secret level 3 cisco
```

Use the **privilege** command to define the commands that can be entered at that privilege level:

```
Router (config)#privilege mode level level command
```

Where mode equals one of the following:

- **configuration**—Global configuration
- **controller**—Controller configuration
- **exec**—EXEC
- **hub**—Hub configuration
- **interface**—Interface configuration
- **ipx-router**—IPX router configuration
- **line**—Line configuration
- **map-class**—Map class configuration
- **map-list**—Map list configuration
- *route-map*—Route map configuration
- *router*—Router configuration

Use the **enable secret level level password** command to set the password for the privilege level.

Example 12-4 shows a user named **student** logging in with a privilege level of 3. The privilege level 3 has been assigned a password of **dallas**. The user will inherit all the commands that have been listed under the **privilege level 3** command as shown previously in Example 12-3.

Example 12-4 *Setting User Privilege Level*

```
Router(config)#enable secret level 3 dallas
Router(config)#enable secret san-fran
Router(config)#username student password cisco

Trying x.x.x.x ... Open

Username: student
Password: cisco
Router>enable 3                Restricted ENABLE privileges
Password: dallas
Router#show privilege          Displays current privilege level
Current privilege level is 3
```

Upon entry to the network device, a banner or message should greet the user. This banner is referred to as the message of the day, having evolved from the UNIX world.

The banner should be a warning and indicate how serious security breaches are to your firm. Computer security practitioners advise not to use the word “welcome” in the message or in any way indicate that you are advocating any entry to the system. Hackers or other intruders have been found not guilty in court due to the simple fact that the word “welcome” was part of the message of the day. Clearly state your security policy and what will happen to violators, if you have room.

The **banner** command uses a delimiter to indicate the end of the message. Any character is valid in the message except the delimiter. The delimiter can also be any character as long as it is not used anywhere else in the message. Example 12-5 demonstrates configuration of the banner message as well as the message displayed upon a user Telnetting to the router.

Example 12-5 *Banner Message Configuration and Display*

```
R1(config)#banner motd 'Unauthorized access will be prosecuted!'
#telnet 192.168.2.5

Unauthorized access will be prosecuted!
Login:
```

Virtual Terminal Access


By default, there are five vtys (otherwise known as Telnet sessions) on each Cisco device. You can create as many as you need. The vtys that you received is based on the number of vtys that are currently in use. Because you will never know exactly which vty line you are using, you should set identical restrictions on all lines.

The **line vty-number vty-range** command takes you into the selected configuration mode of the vtys. The most common use of this command is **line vty 0 4**. This command indicates that you are modifying the first five vtys.

The **access-class** command applies the access list to the interface. The access list is a standard access list that indicates the source addresses that are either permitted or denied. The **in | out** condition at the end of the **access-class** statement indicates whether the source address should be allowed to establish a Telnet session with this device or allowed to Telnet out of this device.

Use caution with the **access-class** command. If you do not match any of the test conditions in the access list, you will be denied Telnet access into the device. The “implicit deny any” at the end of every access list means that when you get to the end, you will deny all other traffic!

Figure 12-3 shows a user with IP address 192.168.2.5 attempting to Telnet to the router.


Figure 12-3 *Configuring vty Access*

Example 12-6 shows the *vtty* access configuration for this user.

Example 12-6 *Configuring vty Access*

```
R1(config)#access-list 1 permit 192.168.2.5
R1(config)# line vty 0 4
R1(config)# access-class 1 in
```

Starting in release 11.0 (6) and later, Cisco allows web browser access to configure your Cisco network device. This access is provided via HTTP and, while easier, it does create some potential security issues. If you turn on HTTP server, no security is default for this command. In other words, anyone can access the router via a web browser. For that reason, applying an access list (covered further in this section) is imperative. The default setting for HTTP access is off. Figure 12-4 illustrates a user with IP address 192.68.10.7 attempting to establish HTTP access.

Figure 12-4 *HTTP Access*

Example 12-7 demonstrates how to enable and configure HTTP access, given the setup in Figure 12-4.

Example 12-7 *Configuring HTTP Access*

```
Router3(config)#access-list 1 permit 192.168.10.7
Router3(config)#ip http server
Router3(config)#ip http access-class 1
Router3(config)#ip http authentication local
Router3(config)#username student password cisco
```

To enable HTTP access, enter the following command:

```
Switch(config)#ip http server
```

You would be wise to apply an access list that has only the required access and nothing more. In Example 12-7, the access list explicitly permits the station 192.168.10.7 and implicitly denies everyone else. By applying the access list with the **ip http access-class 1** statement, all stations other than 192.168.10.7 are denied access to the HTTP software.

Password security for web access can be applied similar to console and virtual terminal access. The following command is used to specify what kind of authentication is being used:

```
Switch(config)#ip http authentication [aaa | enable | local | tacacs]
```

where

- **aaa** indicates that authentication, authorization, and accounting (AAA) should be used for authentication.
- **enable** indicates that the enable password should be used. This is the default method.
- **local** indicates that the local user database is used for authentication information.
- **tacacs** indicates that a TACACS server should be used for authentication.

Access Layer Policy

The access layer is the entry point for users to access the network. Cable connections are generally pulled from an access layer switch to offices and cubicles in a company. For this reason, the network devices of the access layer are the most physically vulnerable. Anyone can plug a station into an access layer switch.

You should take a couple of precautions at the access layer, including

- **Port security**—Limit the Media Access Control (MAC) addresses allowed to use the switch to prevent unauthorized users from gaining access to the network at all.
- **VLAN management**—The default VLAN of all ports is VLAN1. VLAN1 is traditionally the management VLAN. This means that users entering the network on ports that were not configured would be in the management VLAN of the switch block. Cisco recommends that the management VLAN be moved to another VLAN to prevent users from entering the network on VLAN1 on an unconfigured port.

Access Layer Port Security

Port security is a feature of the Cisco Catalyst switches that allows the switch to block input from a port when the MAC address of a station attempting to access the port is different from the configured MAC address. This situation is referred to as a *MAC address lockdown*.

When a port receives a frame, the port compares the source address of the frame to the secure source address that was originally learned by the port. If the addresses do not match, the port is disabled and the LED for the port turns orange.

Port security cannot be applied to trunk ports where addresses may change frequently. Not all hardware supports port security. Check with your documentation or Cisco Connection Online (CCO) to see if your hardware supports this feature.

Configuring Port Security at the Access Layer

By default, the switch allows all MAC addresses to access the network. For network security purposes, the switch relies on mechanisms such as file server operating systems and applications. Port security allows a network administrator to configure a set of allowed devices or MAC addresses to provide additional security. If port security is enabled, only the MAC addresses that are explicitly allowed can use the port. A MAC address can be allowed as follows:

- **Static assignment of the MAC address**—The network administrator can code the MAC address when port security is assigned. This method is the more secure of the two options; however, it is difficult to manage.
- **Dynamic learning of the MAC address**—If the MAC address is not specified, the port turns on learning for security. The first MAC address seen on the port becomes the secure MAC address.

Enabling and Verifying Port Security Using the **set** CLI on **set** Command-Based Switches

Use the following commands to enable and verify port security on a **set** command-based switch:

```
Switch (enable) set port security mod_num/port_num...enable mac address
Switch (enable) show port mod_num/port_num
```

For example, consider the setup in Figure 12-5.

Figure 12-5 *Enabling and Verifying Port Security*

Example 12-8 demonstrates how to enable and then verify port security for the **set** command-based switch in Figure 12-5.

Example 12-8 *Enabling/Verifying Port Security on a set Command-Based Switch*

```
Switch (enable) set port security enable 4/1 02-60-8c-12-34-56
```

```
show port 4/1
```

Port	Security	Secure Src-address	Last Src-address	Shutdown	Trap	IF-index
4/1	enabled	02-60-8c-12-34-56	02-60-8c-12-34-56	no		270

Enabling and Verifying Port Security on Cisco IOS Command-Based Switches

Use the following commands to enable and verify port security on Cisco IOS command-based switches:

```
Switch(config-if)#port secure [max-mac-count maximum-MAC-count]
Switch#show mac-address-table security [type module/port]
```

The **port secure max-mac-count** command allows the network administrator to define the maximum number of MAC addresses that can be supported by this port. The maximum number can range from 1 to 132. The default value is 132.

Distribution Layer Policy

Most of the access control policy will be implemented at the distribution layer. This layer is also responsible for ensuring that data stays in the switch block unless that data is specifically permitted outside of the switch block. This layer is also responsible for sending the correct routing and service information to the core.

A good policy at the distribution layer ensures that the core block or the WAN blocks are not burdened with traffic that has not been explicitly permitted. A distribution layer policy also protects the core and the other switch blocks from receiving incorrect information, such as incorrect routes, that may harm the rest of the network.

Access control at the distribution layer falls into three different categories:

- Defining which user traffic makes it between VLANs and thus ultimately to the core. This control can be done in the form of an access list applied to an interface to permit only certain data to pass through.
- Defining which routes are seen by the core block and the switch block. This control can be done through the use of distribution lists to prevent routes from being advertised to the core.
- Defining which services the switch block will advertise out to the rest of the network. Service control could also be used to define how the network finds the server-aggregation block in order to get services like Dynamic Host Control Protocol (DHCP) and Domain Name System (DNS).

Filtering Traffic at the Distribution Layer

Many of the access control methods used at the distribution layer rely on the creation of an access control list. Two types of IP access lists are available—standard and extended.

Each type of access list is a series of permits and denies based on a set of test criteria. The standard access list allows for a test criteria of only the source address. The extended access list allows for greater degree of control by checking the source and destination addresses as well as the protocol type and the port number or application type of the packet. A standard access list is easier for the router to process; an extended access list, however, provides a greater degree of control.

Access lists are created for a variety of applications. Access lists can be used for controlling access in the campus network by applying them in different capacities. These include the following:

- Applying the access list to the interface for traffic management purposes through the use of the *protocol* **access-group** command, where *protocol* is the Layer 3 protocol that is being managed.
- Applying the access list to a line for security purposes through the use of the **access-class** command. This list determines the users of a specific line. This chapter focuses on the vtys.
- Managing routing update information through the use of the **distribution-list** command. This access list determines which routes are learned by the router and which routes are advertised out of the router.
- Managing services update information through the use of commands such as **ipx output-sap-filter** in order to determine which services are advertised.

IP Standard Access List Overview

IP standard access lists include the following characteristics:

- Test condition is based on the source address only.
- Numbered standard access lists are 1 to 99.
- Access list is processed from the top down. As soon as a match is found, the access list stops processing.
- There is an implicit deny of everything at the end of every access list. If no match is found in the access list, it will ultimately match the implicit deny at the end of the list.
- The creation of the access list does nothing until the access list is applied.
- Access lists can be applied either inbound or outbound. An inbound access list checks the packet as it enters the interface before it has been routed. An outbound access list checks as the packet goes out an interface after the packet has been routed.

Use the **access-list** command to create an entry in a standard traffic filter list:


```
Router(config)#access-list access-list-number {permit | deny} source-address  
[source-wildcard]
```

where

- *access-list-number* identifies the list to which the entry belongs. For an IP standard access list, use a number from 1 to 99.
- **permit** | **deny** indicates what the result will be if the test condition is matched. A permit will allow the test condition either in or out of the interface. A deny will drop the packet and send an ICMP message back to the source.
- *source-address* identifies the source IP address to match.
- *source-wildcard* indicates how much of the address to match. A 0 indicates that it must match the corresponding bit in the source address; a 1 indicates that the corresponding bit can be any value.

The access control list can now be applied to the interface for traffic management purposes. To apply the access list to the interface, use the **ip access-group** *access-list-number* **in** | **out** command.

By default, the **access-group** command is set for outbound processing. This action means that the packet will be checked after it has been routed and just before the packet exits the interface. You can modify this access list for inbound checking by applying the **in** keyword at the end of the **access-group** command. For example, consider the setup in Figure 12-6.

Figure 12-6 Restricting Access with Access Lists

Example 12-9 demonstrates how you would configure an access list for the router in Figure 12-6.

Example 12-9 Configuring a Standard IP Access List

```
access-list 1 permit 192.168.2.5
interface vlan 10
ip address 192.168.4.1 255.255.255.0
access-group 1 out
ip access-group 1 out
```

In Example 12-9, the **access-list 1** is configured, which permits only a specific network to be passed. The **access-group** command is then used on an interface basis (interface VLAN10) and is used on an outbound basis.

IP Extended Access List Overview

An extended access list follows many of the same principals of a standard access list. However, an extended list provides for a higher degree of control by enabling filtering based on the source address as well as the destination address, the protocol type, and the application or port number.

Extended access lists have the following characteristics:

- Top-down processing of the access list. As soon as a match is made in the access list, it stops processing and either permits or denies the packet based on the statement in the access list.

Numbered access lists use a range of 100 to 199. In IOS 12.0, however, this is updated to include 1300-1999.

- Test conditions include protocol type, source address, destination address, application port, and session layer information.
- There is an implicit deny of everything at the end of the access list.
- The creation of the access list does nothing until the access list is applied using the appropriate command.

After you have defined your policy for traffic management, apply it to the distribution layer.

Consider the access list in Example 12-10.

Example 12-10 *Configuring an Extended IP Access List*

```
access-list 101 permit tcp any 192.168.7.0 0.0.0.255
access-list 101 permit tcp any host 192.168.2.5 eq smtp
access-list 101 permit icmp any any echo
access-list 101 permit icmp any any echo-reply
!
interface VLAN10
ip access-group 101 out
```

The access list in Example 12-10 does all of the following:

- Allows all TCP traffic coming from any host going to the subnetwork of 192.168.7.0 0.0.0.255.
- Allows any device to reach the host of 192.168.2.5 if the application is the Simple Mail Transfer Protocol (SMTP) (mail).
- Allows Internet Control Message Protocol (ICMP) echo and echo reply (ping).
- Denies all other traffic (implicit).

Use the **access-list** command to create an entry in an extended traffic filter list:

```
Router(config)#access-list access-list-number { permit | deny {protocol |
protocol-keyword}}{source source-wildcard | any}
{destination destination-wildcard | any}[protocol-specific options] [log]
```

where

- *access-list-number* identifies the list to which the entry belongs. For an IP extended access list, use a number from 100 to 199.
- **permit** | **deny** indicates what the result will be if the test condition is matched. A **permit** will allow the test condition either in or out of the interface. A **deny** will drop the packet and send an ICMP message back to the source.
- *protocol-keyword* indicates the protocol type to match. Options include **ip**, **tcp**, **udp**, **icmp**, **igrp**, **eigrp**, **ospf**, **nos**, or any number in the range of 0 to 255. To match any protocol, use the keyword **ip**.
- *source* and *destination* indicate the IP addresses of both the source and the destination.

- *source-wildcard* and *destination-wildcard* indicate the wildcard mask to indicate the number of address bits to match. A 0 indicates to match the bit exactly; a 1 indicates that the bit can be anything.
- **log** causes informational logging messages about the packet that matches the entry. Use this command with caution, because it consumes CPU cycles.

Controlling Routing Update Traffic

Controlling the routing table of the core block has several advantages:

- Reduces the size of the routing table at the core block allowing it to process packets faster.
- Prevents users from getting to networks that have not been advertised unless they have a static or default route to get there.
- Prevents incorrect information from propagating through the core block.

Two methods are available for controlling the routing information that is sent to the core block, as follows:

- **Route summarization**—Depending on the routing protocol used, a summarized entry of all the available routes of the switch block can be sent from the distribution layer to the core.
- **Distribution lists**—A distribution list can be used to indicate what routes the distribution layer can advertise to the core, or conversely, what the core can accept from the switch block.

NOTE Route summarization is another way to limit the size of the routing table at the core block, but this method is not covered here.

Configuring Route Filtering

The basic method for configuring route filtering is by using the **distribute-list** command. This method is used frequently in large routed networks but can be used by Route Switch modules (RSMs) in a large switched network as well.

The basic command syntax for configuring route filtering for inbound routing updates is

```
R1(config-router)# distribute-list access-list-number | name in [type number]
```

Similarly, the command syntax for configuring route filtering for outbound routing updates is

```
R1(config-router)# distribute-list access-list-number | name out [interface-name]  
routing-process | autonomous-system-number
```

The command arguments for the **distribute-list** command are described as follows:

- *access-list-number*—Number of the previously created standard access list.
- **in** | **out**—Define the filtering on either incoming routing updates (**in**) or outgoing routing updates (**out**).
- *interface-name*—Name of the interface. Indicates that the networks in the access list will be filtered if they came from or are going to a specific interface.
- *routing-process autonomous-system-number*—Name of the routing process including the keywords of **static** and **connected**. This option applies only to outbound distribution filters.

You can filter routing update traffic for any protocol by defining an access list and applying it to a specific routing protocol.


To configure a filter, perform the following steps:

- Step 1** Identify the network addresses that you want to filter and create a standard access list.
- Step 2** Determine whether the routing protocol should be filtered incoming or outgoing on the interface.
- Step 3** Assign the access filter to routing updates.

IP Route Filtering

Consider the network device setup in Figure 12-7.

Figure 12-7 IP Route Filtering



The command syntax in Example 12-11 indicates that the routing process of Enhanced Interior Gateway Routing Protocol (EIGRP) will send the network of 140.20.0.0 255.255.0.0 in its routing updates out E0/0 (Ethernet) but will filter all other networks. If the core is connected to VLAN10, it will receive only 140.20.15.0 and only 140.20.15.0.0 will be allowed to traverse the core.

Example 12-11 *Configuring IP Route Filtering*

```
router eigrp 100
 network 140.20.0.0

 distribute-list 7 out int VLAN10
 !
 access-list 7 permit 140.20.15.0 0.0.0.255
```

The options for the networks of 140.20.x.0, except 140.20.16.0, include the following:

- All other networks will be able to send and receive data in the switch block but will not be allowed to get to any other switch block or to the core block. For this setup to work, a static or default route will have to be configured.
- All other networks will not be seen by the core block and other switch blocks. A default or static route will allow them to send and receive data to other switch blocks, including the core.

Core Layer Policy

The core block is responsible for moving data quickly. All the devices that are designed to be core block solutions are optimized to move data as quickly as possible. For this reason, the core block should have little to no policy.

The only policies that should be applied at the core block are those that relate to quality of service (QoS) commands for congestion management and congestion avoidance.

QoS implementations vary, depending on hardware used and versions of IOS. Please see your IOS-specific documentation for details.

Foundation Summary

The Foundation Summary is a collection of tables and figures that provides a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final preparation before the exam, these tables and figures will hopefully be a convenient way to review the day before the exam.

Table 12-3 *Route Switch Module or Router and Switch Commands*

Command	Description
access-list <i>access-list</i>	Creates an access list
distribute-list <i>access-list</i> [in out]	Applies an access list to a routing protocol
line <i>line-type line-number</i>	Selects a line to configure
login [local tacacs]	Indicates where the login should look for information
privilege mode level <i>level command</i>	Enters the commands available at a privilege level
username <i>username password password</i>	Creates a username entry in the local database
username <i>username privilege number</i>	Assigns a privilege level to <i>username</i>
Switch command: set port security <i>mod_num/port_num...enable mac address</i>	Creates port security using MAC address

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; they are designed, however, to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to these questions can be found in Appendix A, on page 477.

- 1 Define an access policy.

- 2 What is the access layer defined as?

- 3 Is HTTP access normally enabled on a Cisco router? What is the main purpose of using HTTP?

- 4 Name at least two components relating to controlling access to network devices.

- 5 What way of accessing a network device requires a password?

6 What feature of the Cisco IOS protects a console connection left unattended?

7 What does the **access-class** command do when applied to a virtual terminal configuration?

8 What VLAN is the default VLAN for a Catalyst switch and why is it a good idea to change this?

9 What does port security do on a Catalyst series switch?

10 What is the range of numerical representation of a standard IP access list? An extended access list?

11 Should a standard or an extended access list be used when filtering a particular host?

12 When implementing route filtering, what type of access list is used—a standard or an extended access list?

13 In general, what type of policies should be implemented in the core layer?

14 Which physical access method of a Cisco router should be disabled if not used?

15 What is the virtual terminal connection commonly called?

16 What does the banner do?

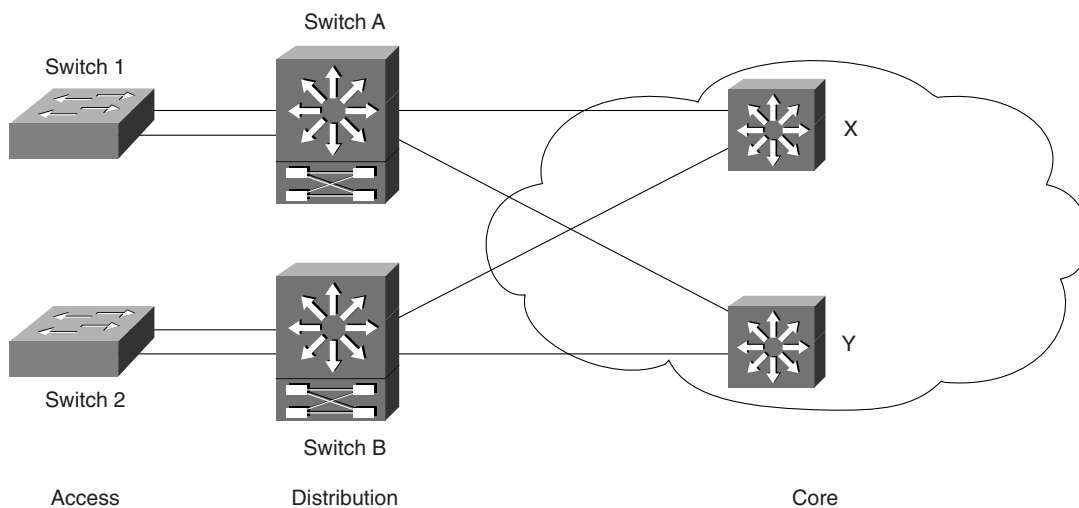
17 Why is it important to have physical security for a network device?

18 What does the Cisco command **login local** do on a router?

Scenarios

Please refer to the Scenario Figure 12-8 below as a reference to Scenario 12-1 and Scenario 12-2.

Figure 12-8 Scenario 12-1 and 12-2 Network



Scenario 12-1

Given the network depicted in Figure 12-8, answer the following questions related to this scenario.

- 1 Assume you are connected to the console port of the RSM on Switch A. Establish a console login with a password of **san-jose**.
- 2 While still connected to the console, establish a Telnet login with a password of **san-fran**.
- 3 Assume that a management VLAN (VLAN1) exists on Switch 1 and Switch A. Further, a workstation is connected to VLAN1 off of Switch 1. Set up an access list on the RSM on Switch A to allow only the workstation to Telnet to Switch A. Assume the workstation has the IP address of 192.168.1.12.
- 4 Following the configuration in Exercise 3 for this scenario, add HTTP access to the RSM on Switch A. Assume local authentication with a username of **web** and password of **cisco**.
- 5 Configure Switch 1 such that the aforementioned workstation is the only one allowed to be connected on port 4/5. The workstation has a MAC address of 00-00-0e-12-34-56.

Scenario 12-2

- 1 Set a banner message upon login to Switch B. It should read, “Unauthorized access will be prosecuted.”
- 2 Set an extended access list 101 such that only SMTP traffic is allowed to and from the RSM Switch B on Interface VLAN 100.
- 3 Switch B has a VLAN 200 that connects to core Router Y. The RSM on Switch B is running EIGRP with a process ID of 225. Construct a distribute list that allows only routes from 172.16.100.0 to traverse into the core.
- 4 Construct a new privilege level on Switch 2 that allows the user to log in as the operator with password of **cisco**. This privilege level allows only one thing—to show the startup configuration.

Scenarios Answers

Scenario 12-1 Answers

- 1 The console login should look something like the configuration that follows:

```
RSM(config)#line console 0
RSM(config-line)#login
RSM(config-line)#password san-jose
```

- 2 The Telnet or vty login statements look very similar to that of the console. The correct answer is as follows:

```
RSM(config)#line vty 0 4
RSM(config-line)#login
RSM(config-line)#password san-fran
```

- 3 The correct configuration is as follows:

```
RSM(config)#access list 1 permit 192.168.1.12
RSM(config)#line vty 0 4
RSM(config-line)#access-class 1 in
```

- 4 The correct configuration is as follows:

```
RSM(config)#access list 1 permit 192.168.1.12
RSM(config)#line vty 0 4
RSM(config-line)#access-class 1 in
RSM(config)#ip http server
RSM(config)#ip http access-class 1 in
RSM(config)#ip http authentication local
RSM(config)#username web password cisco
```

- 5 This feature is implemented on Switch 1 and designates only one particular MAC address access to the port:

```
Switch1(enable)#set port security enable 4/5 00-00-0e-12-34-56
```

Scenario 12-2 Answers

- 1 The correct answer is as follows:

```
Switch1(enable)#set banner motd "Unauthorized access will be prosecuted!"
```

- 2 The correct configuration is as follows:

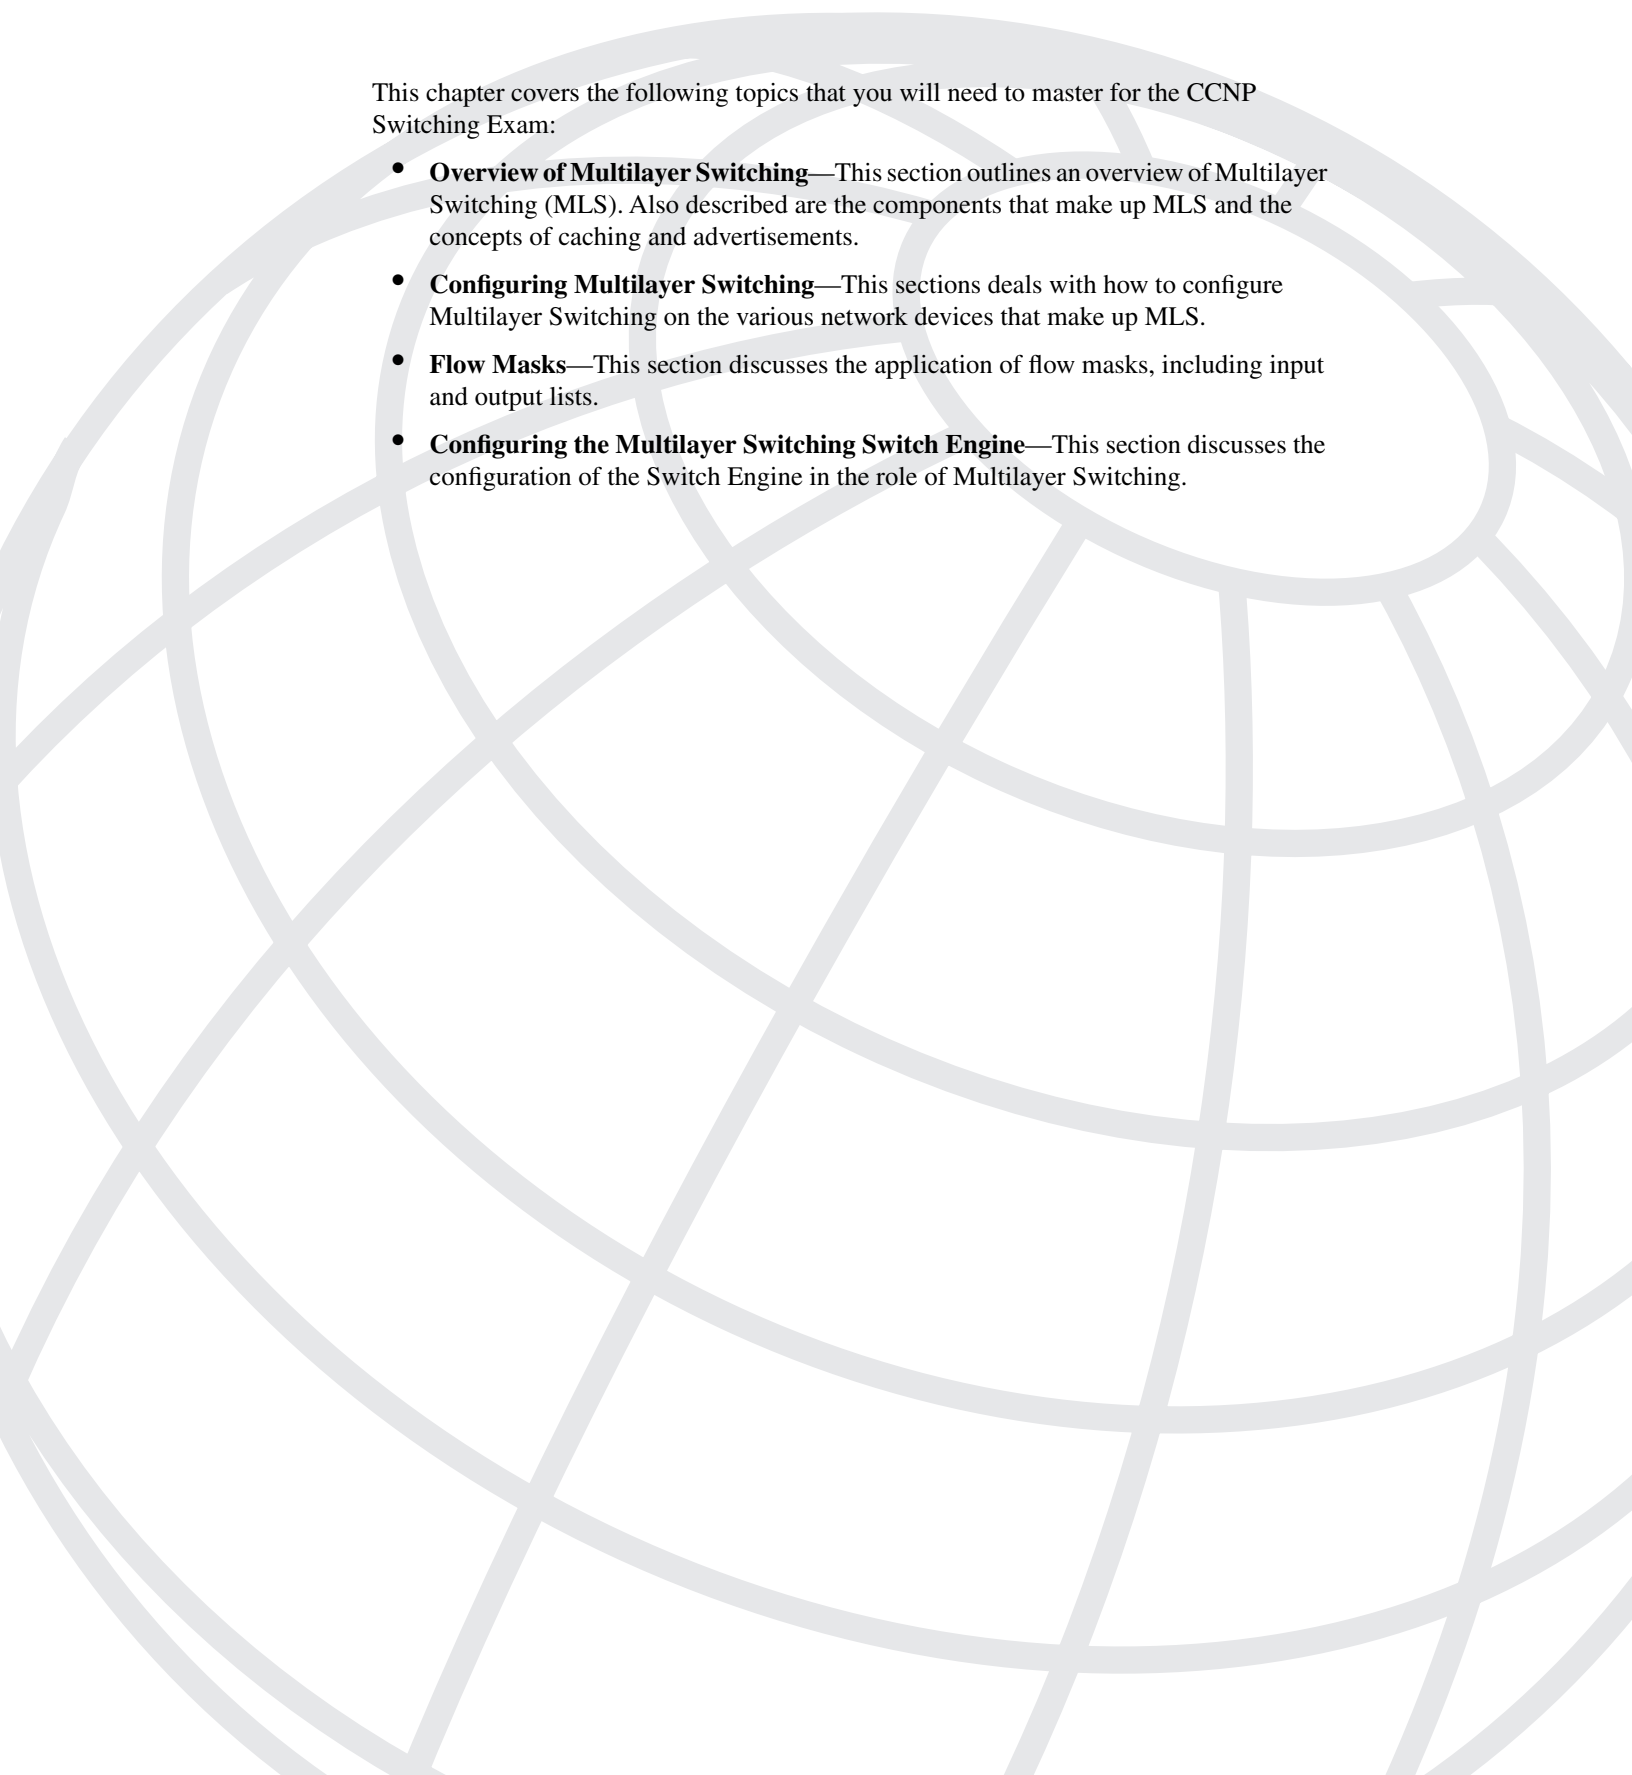

```
interface VLAN100
 access-group 101 out
!
access list 101 permit tcp any any eq smtp
```

- 3 The correct configuration is as follows:

```
router eigrp 225
network 172.16.0.0
!
distribute-list 5 out VLAN200
access-list 5 permit 172.16.100.0 0.0.0.255
```

- 4 The configuration that accomplishes the goal for this exercise is as follows:

```
privilege configure level 3 username
privilege exec level 3 show run
enable secret level 3 cisco
username operator password cisco
```



This chapter covers the following topics that you will need to master for the CCNP Switching Exam:

- **Monitoring Cisco Switches**—This section covers the methods available and commands used for monitoring Cisco switches.
- **General Troubleshooting Model**—This section reviews a general model for troubleshooting network devices, including Cisco switches.
- **Troubleshooting Cisco Switches with show Commands**—This section discusses and defines the various commands that can be used to troubleshooting Cisco switches.
- **Physical Layer Troubleshooting**—This section discusses the tools involved in troubleshooting the physical layer.

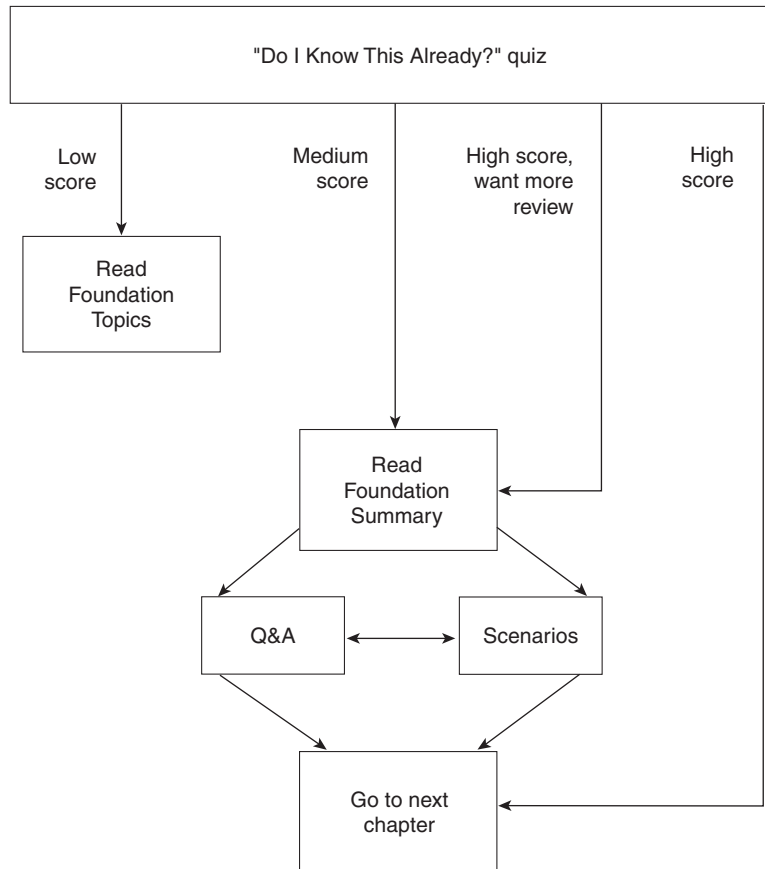
Monitoring and Troubleshooting

We've examined a lot of aspects of the Cisco switching world, but now it's time to check out the issues of monitoring and troubleshooting the environment. We will look at what methods are available to monitor the switching environment as well as some of the commands associated with them. Additionally, we will discuss a general model for troubleshooting the Cisco switched environment and the associated commands.

How to Best Use This Chapter

By taking the following steps, you can make better use of your study time:

- Keep your notes and answers for all your work with this book in one place for easy reference.
- Take the “Do I Know This Already?” quiz and write down your answers. Studies show retention is significantly increased through writing facts and concepts down, even if you never look at the information again.
- Use the diagram in Figure 13-1 to guide you to the next step.

Figure 13-1 *How To Use This Chapter*

“Do I Know This Already?” Quiz

The purpose of the “Do I Know This Already?” quiz is to help you decide what parts of this chapter to use. If you already intend to read the entire chapter, you do not necessarily need to answer these questions now.

The 12-question quiz helps you make good choices of how to spend your limited study time. Use the scoresheet in Table 13-1 to record your score.

Table 13-1 *Scoresheet for Quiz*

Quizlet Number	Foundation Topics Section Covering These Questions	Questions	Score
1	Monitoring Cisco Switches	1–2	
2	Monitoring Commands	3–4, 6	
3	General Troubleshooting Model	5	
4	Troubleshooting Commands	7–10	
5	Physical Layer Troubleshooting	11–12	
All questions		1–12	

1 What is the main method of out-of-band management for Cisco switches?

2 What is an application that uses SNMP to perform in-band management?

3 CDP operates at what layer of the OSI model?

4 What is the command to verify that RMON is enabled on the switch?

5 Using a troubleshooting model, what step is generally taken after ascertaining all the facts?

6 What is the default value for the read-write community string?

7 How many simultaneous Telnet sessions are supported on a Cisco switch?

8 What command shows information on the modules installed in a particular switch?

9 What command shows the contents of memory displaying MAC addresses and their associated ports?

10 What command would I use if I wanted to see information on Spanning Tree?

11 What physical layer troubleshooting tool is similar to “sonar” in that it bounces a signal to determine length.

12 What is the name of the tool that decodes the various protocols in captured packets?

The answers to the “Do I Know This Already?” quiz are found in Appendix A, “Answers to the ‘Do I Know This Already?’ Quizzes and Q&A,” on page 477. The suggested choices for your next step are as follows:

- **8 or fewer overall score**—Read the chapter. This reading includes the “Foundation Topics,” the “Foundation Summary,” Q&A, and scenarios at the end of the chapter.
- **9–12 overall score**—Begin with the “Foundation Summary,” and then follow with the Q&A and scenarios at the end of the chapter.

Foundation Topics

Monitoring Cisco Switches

You can monitor and manage your Catalyst switches in a number of different ways. One way is primarily through a console port using either the command-line interface (CLI) or other methods for performing network management functions, such as Cisco Discovery Protocol (CDP), Embedded Remote Monitoring (RMON), or Switched Port Analyzer (SPAN). The console port is an EIA/TIA-232 DCE interface to which you can connect a console terminal or modem. The type of connector, however, used depends on the hardware. On a Catalyst 5000 with Supervisor I or II, a rollover cable is used with the above hardware. On a Supervisor III or a Catalyst 6000, a straight through cable is used in conjunction with a modular plug. Other kinds of switches may be different.

Through the console port, you can directly access the CLI or configure a Serial Line Internet Protocol (SLIP) interface to access such network management functions as Telnet, **ping**, and SNMP. An IP address can be assigned to the Cisco switch for management purposes. Once the address is in place, you can direct Telnet to access the IP address of the switch to reach the CLI. You can also use the IP address of the switch to access an SNMP agent, such as CiscoWorks 2000.

NOTE This chapter contains references and commands that are based on Cisco Switch-Based IOS.


Out-of-Band Management

Out-of-band management access for Cisco switches in general is performed via the following methods:

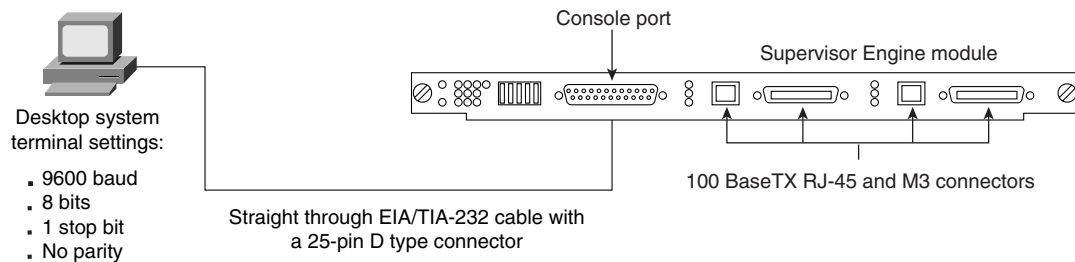
- Console Port Connection
- SLIP

Console Port Connection

The console port is the local (out-of-band) console terminal connection to the switch—a DB-25 female connector shown in Figure 13-2. Other switches may require different console cables in conjunction with modular plugs instead of the DB-25.

Figure 13-2 *The Catalyst 5000 Console Port*

To use the console port, connect via a straight-through cable, an EIA/TIA-232 terminal (configured for 9600 baud, no parity, eight data bits, and one stop bit), modem, or network management workstation, as shown in Figure 13-3. As noted above, however, this may vary depending on the type of switch used.

Figure 13-3 *Attaching to the Console Port*

The console port enables you to perform the following functions:

- Configure the switch with a command-line interface.
- Monitor network statistics and errors.
- Configure SNMP agent parameters.
- Download software updates to the switch or distribute software images residing in Flash memory to attached devices.

Serial Line Internet Protocol (SLIP)

You can access the Cisco switch command line using Serial Line Internet Protocol (SLIP). This protocol is a version of Internet Protocol (IP) that runs over serial links allowing IP communications through the console port.


Configuring SLIP on the Console Port

Catalyst series switches support out-of-band management through the use of a modem attached to the console port. This out-of-band connection works in conjunction with SLIP. The out-of-band connection can be used to:

- Establish a Telnet session that provides access to the Cisco switch CLI.
- Use the Telnet Server feature.
- Establish an SNMP management session that provides the capability to use an SNMP-based management platform such as the CiscoWorks 2000 solution.

To establish an out-of-band connection on a Cisco switch, connect a 100 percent Hayes-compatible modem by means of a straight-through cable with a 25 pin D type connector as shown in Figure 13-4. The modem should be configured for auto answer mode.

Figure 13-4 Out-of-band Management Using SLIP



Use the SLIP (sl0) interface for point-to-point SLIP connections between the switch and an IP host.

CAUTION

You *must* use the console port for the SLIP connection. When the SLIP connection is enabled and SLIP is attached on the console port, an EIA/TIA-232 terminal cannot connect via the console port. If you are connected to the switch CLI through the console port and you enter the **slip attach** command, you will lose the console port connection. Use Telnet to access the switch, enter privileged mode, and enter the **slip detach** command to restore the console port connection.

To enable and attach SLIP on the console port, perform the following sequence of tasks:

- Step 1** Access the switch from a remote host with Telnet via the **telnet** *{host_name | ip_addr}* command.
- Step 2** Enter privileged mode on the switch via the **enable** command.
- Step 3** Set the console port SLIP address and the destination address of the attached host via the **set interface s10 slip_addr dest_addr** command.
- Step 4** Enable SLIP for the console port via the **slip attach** command.
- Step 5** Verify the SLIP interface configuration via the **show interface** command.

Example 13-1 shows how to configure SLIP on the console port and verify the configuration:

Example 13-1 *Configuring SLIP on the Console Port and Verifying the Configuration*

```

sparc20% telnet 172.20.52.71
Trying 172.20.52.71 ...
Connected to 172.20.52.71.
Escape character is '^]'.

Cisco Systems Console
Enter password:
Console> enable
Enter password:
Console> (enable) set interface s10 10.1.1.1 10.1.1.2
Interface s10 slip and destination address set.
Console> (enable) slip attach
Console Port now running SLIP.
Console> (enable) show interface
s10: flags=51<UP,POINTOPOINT,RUNNING>
      slip 10.1.1.1 dest 10.1.1.2
sc0: flags=63<UP,BROADCAST,RUNNING>
      vlan 523 inet 172.20.52.71 netmask 255.255.255.224 broadcast 172.20.52.95
Console> (enable)

```

In-Band Management

The following protocols are used to perform in-band management of a Cisco switch:

- Simple Network Management Protocol (SNMP)
- Telnet
- Cisco Discovery Protocol (CDP)

SNMP

Simple Network Management Protocol (SNMP) is an application layer protocol designed to facilitate the exchange of management information between network devices. The SNMP system consists of three parts—SNMP manager, SNMP agent, and Management Information Base (MIB).

Instead of defining a large set of commands, SNMP places all operations in a **get-request**, **get-next-request**, and **set-request** format. For example, an SNMP manager can get a value from an SNMP agent or store a value into that SNMP agent. The SNMP manager can be part of a network management system (NMS), and the SNMP agent can reside on a networking device such as a switch. The SNMP agent can respond to MIB-related queries being sent by the NMS.

The following list details the basic functions supported by SNMP agents:

- **Accessing a MIB variable using the get-request or get-next-request format**—This function is initiated by the SNMP agent as a result of a request for the value of a MIB variable from a network management station. The SNMP agent gets the value of a MIB variable by accessing information stored in the MIB and then responds.
- **Setting a MIB variable**—This function is also initiated by the SNMP agent as a result of a message from a network management station. The SNMP agent requests that the value of a MIB variable be changed.
- **SNMP trap**—This function is used to notify a network management station that an extraordinary event has occurred at an agent. When a trap condition occurs, the SNMP agent sends an SNMP agent trap message to each of the network management stations as specified in the trap receiver table.

To configure SNMP on your switch, perform the following steps:

- Step 1** Configure the SNMP community strings via the **set snmp community {read-only | read-write | read-write-all} community_string** command.
- Step 2** Assign a trap receiver address and community via the **set snmp trap rcvr_address rcvr_community** command. If you enter incorrect information, enter the **clear snmp trap** command to delete the entry. Then re-enter the **set snmp trap** command.
- Step 3** If desired, configure the switch so that it issues an authentication trap via the **set snmp trap enable** command.

The **set snmp** Command Options

The syntax for the **set snmp community** command, used to configure SNMP community strings, is as follows:

```
set snmp community {read-only | read-write | read-write-all} [community_string]
```

The keywords for the **set snmp community** command are as follows:

- **read-only**—Keyword to assign read-only access to the specified SNMP community.
- **read-write**—Keyword to assign read-write access to the specified SNMP community.
- **read-write-all**—Keyword to assign read-write access to the specified SNMP community. The read-write-all offers access to the community strings themselves.
- *community_string*—An optional parameter. This is the name of the SNMP community. The default SNMP community strings are as follows:
 - **read-only** (public)
 - **read-write** (private)
 - **read-write-all** (secret)

Example 13-2 demonstrates some sample output after entering the **set snmp community** command exercising all three *community_string* options.

Example 13-2 *set snmp community Command Output*

```

Console> (enable) set snmp community read-only public
SNMP read-only community string set.

Console> (enable) set snmp community read-write private
SNMP read-write community string set.

Console> (enable) set snmp community read-write-all secret
SNMP read-write-all community string set.

```

As demonstrated in Example 13-3, to view the options of the **set snmp** command enter the command at the CLI in enable mode.

Example 13-3 *Viewing set snmp Command Options*

```

Console> (enable) set snmp

Set snmp commands:
-----
set snmp community      Set SNMP community string
set snmp help           Show this message
set snmp rmon           Set SNMP RMON
set snmp trap           Set SNMP trap information

```

An IP permit trap is sent when unauthorized access based on the IP permit list is attempted. The **set snmp trap** command is a privileged mode switch command used to enable or disable the different SNMP traps on the system or to add an entry into the SNMP authentication trap receiver table. The default configuration has SNMP traps disabled. Use the **show snmp**

command to verify the appropriate traps were configured. The syntax for the **set snmp trap** command is as follows:

```
set snmp trap {enable | disable} [all | module | chassis | bridge | repeater | auth |
vtp | ippermit | vmps | config | entity | stpx]
set snmp trap rcvr_addr rcvr_community
```

Table 13-2 documents the keywords and arguments for the **set snmp trap** command

Table 13-2 set snmp trap *Command Keywords/Arguments*

Command Keyword/Argument	Definition
enable	Keyword to activate SNMP traps.
disable	Keyword to deactivate SNMP traps.
all	Optional keyword to specify all trap types.
module	Optional keyword to specify the moduleUp and moduleDown traps from the CISCO-STACK-MIB.
chassis	Optional keyword to specify the ciscoSyslogMIB trap from the CISCO-SYSLOG-MIB.
bridge	Optional keyword to specify the newRoot and topologyChange traps from RFC 1493 (the BRIDGE-MIB).
repeater	Optional keyword to specify the rptrHealth, rptrGroupChange, and rptrResetEvent traps from RFC 1516 (the SNMP-REPEATER-MIB).
auth	Optional keyword to specify the authenticationFailure trap from RFC 1157.
vtp	Optional keyword to specify the VTP from the CISCO-VTP-MIB.
ippermit	Optional keyword to specify the IP Permit Denied access from the CISCO-STACK-MIB.
vmps	Optional keyword to specify the vmVmpsChange trap from the CISCO-VLAN-MEMBERSHIP-MIB.
config	Optional keyword to specify the sysConfigChange trap from the CISCO-STACK-MIB.
entity	Optional keyword to specify the entityMIB trap from the ENTITY-MIB.
stpx	Optional keyword to specify the STPX trap.
<i>rcvr_addr</i>	IP address or IP alias of the system to receive SNMP traps.
<i>rcvr_community</i>	Community string to use when sending authentication traps.

Example 13-4 shows how to enable SNMP chassis traps:

Example 13-4 *Enabling SNMP Chassis Traps*

```
Console> (enable) set snmp trap enable chassis
SNMP chassis alarm traps enabled.
Console> (enable)
```

Example 13-5 shows how to enable all SNMP traps:

Example 13-5 *Enabling All SNMP Traps*

```
Console> (enable) set snmp trap enable
All SNMP traps enabled.
Console> (enable)
```

Example 13-6 shows how to disable SNMP chassis traps:

Example 13-6 *Disabling SNMP Chassis Traps*

```
Console> (enable) set snmp trap disable chassis
SNMP chassis alarm traps disabled.
Console> (enable)
```

Example 13-7 shows how to add an entry in the SNMP trap receiver table:

Example 13-7 *Adding an Entry in the SNMP Trap Receiver Table*

```
Console> (enable) set snmp trap 192.122.173.42 public
SNMP trap receiver added.
Console> (enable)
```

SNMP Verification

To verify SNMP settings, enter the **show snmp** command. After entering this command, you will see the output in Example 13-8.

Example 13-8 *show snmp Command Output*

```
Console> show snmp
RMON: Enabled
Traps Enabled: Chassis
Port Traps Enabled: None
Community-Access      Community-String
-----
read-only              public
Trap-Rec-Address      Trap-Rec-Community
-----
192.122.173.42       public
Console>
```

Telnet Client Access

Remote, in-band SNMP management is possible through any LAN or ATM interface assigned to the same VLAN as the Supervisor module's NMP IP address. In-band connections can be used to establish Telnet sessions to the Cisco switch CLI or SNMP management sessions on an SNMP-based management platform, such as CiscoWorks 2000.

Cisco switches provide outgoing Telnet functionality from the CLI; this feature allows a network manager to use Telnet from the CLI of the switch to other devices on the network. Using Telnet, a network manager can maintain a connection to a Cisco switch while also connecting to another switch or router. Cisco switches support up to eight simultaneous Telnet sessions. Telnet sessions disconnect automatically after remaining idle for a configurable time period. To access the switch through a Telnet session, you must first set the IP address for the switch.

NOTE

Outgoing Telnet is allowed from "enable" access mode on the Catalyst 4000, 5000, and 6000 series switches. The syntax below demonstrates an attempt to Telnet from user EXEC mode.

```
6500-1> telnet
Unknown command "telnet". Use 'help' for more info.
6500-1>
```

To access the switch from a remote host with Telnet, perform these steps:

- Step 1** From the remote host, enter the **telnet** command and the name or IP address of the switch you want to access. The syntax for this command is: **telnet** {*hostname* | *ip_addr*}.
- Step 2** At the prompt, enter the *<password>* for the CLI. If no password has been configured, press **Enter**.
- Step 3** Enter the necessary commands to complete your desired tasks.
- Step 4** When finished, exit the Telnet session via the **quit** command.

After entering the **telnet** command, you will see the display in Example 13-9.

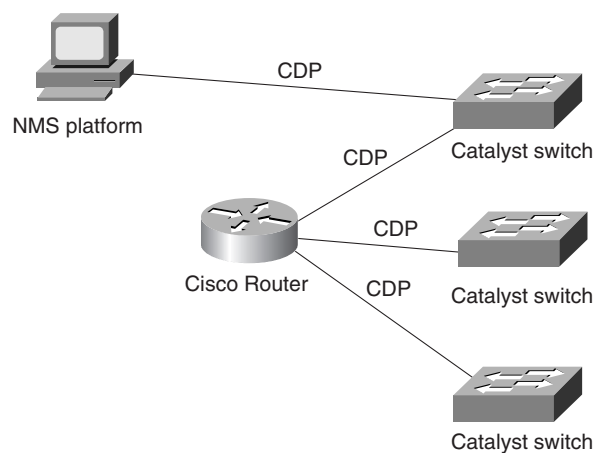
Example 13-9 telnet Command Output

```
host% telnet cat5000-1.cisco.com
Trying 172.16.44.30 ...
Connected to cat5000-1.
Escape character is '^]'.
Cisco Systems Console
Enter password: <password>
Console>
```

Cisco Discovery Protocol (CDP)

Cisco Discovery Protocol (CDP) is media- and protocol-independent and runs on all Cisco-manufactured equipment including routers, bridges, access and communication servers, and switches. With CDP, network management applications can retrieve the device type and the SNMP-agent address of neighboring devices (see Figure 13-5). Applications are now enabled to send SNMP queries to neighboring devices.

Figure 13-5 *A Typical Cisco Network Environment with CDP Enabled*



CDP meets a need created by the existence of lower-level, virtually transparent protocols. CDP enables network management applications to dynamically discover Cisco devices that are neighbors of already known devices, neighbors running lower-layer transparent protocols in particular. CDP runs on all media that support the Subnetwork Access Protocol (SNAP). CDP runs over the data link layer only, not the network layer. Therefore, two systems that support different network layer protocols can learn about each other. Cached CDP information is available to network management applications. Cisco devices never forward a CDP packet. When new information is received, old information is discarded.

Example 13-10 shows how to display CDP information about neighboring systems:

Example 13-10 *Displaying CDP Information about Neighboring Systems*

```

Console> show cdp neighbor 4

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
Port    Device-ID                Port-ID    Platform    Capability
-----
4/1     001905905                4/1       WS-C5000    TS
4/1     062000101 (CAT3)        9         WS-C1201    SI
  
```

continues

Example 13-10 *Displaying CDP Information about Neighboring Systems (Continued)*

```

4/1      069000022          8/1      WS-C5500      TS
4/1      069000040          4/2      WS-C5500      TS
Console>

An explanation of the output screen is shown below:
Port - Port that the CDP information was learned on.
Device-ID - Serial number of the device (and name if configured)
Port-ID - Port at the remote device
Platform - Cisco product number
Capability - Capability of the device (see Capability Codes listed at top of output

```

Embedded Remote Monitoring

Cisco switches provide support for the Embedded Remote Monitoring (RMON) of Ethernet and Fast Ethernet ports. Embedded RMON provides you with visibility into network activity. It enables you to access and remotely monitor the RMON specification RFC 1757 groupings of statistics, historical information, alarms, and events for any port through SNMP or the TrafficDirector Management application.

The RMON feature monitors network traffic at the data link layer of the OSI model without requiring a dedicated monitoring probe or network analyzer. RMON enables a network manager to analyze network traffic patterns, set up proactive alarms to detect problems before they affect users, identify heavy network users as candidates to move to dedicated or higher speed ports, and perform trend analysis for long-term planning.

The statistics group of the RMON specification maintains utilization and error statistics for the switch that is monitored. Statistics include information about:

- Collisions.
- Cyclic redundancy checks (CRC) and alignment.
- Undersized or oversized packets.
- Jabber.
- Fragments.
- Broadcast, multicast, and unicast messages.
- Bandwidth utilization.

The history group takes periodic samples from the statistics section and stores them for later retrieval. This includes information such as utilization, error counts, and packet counts.

A system network administrator uses the alarm group to set a sampling interval and threshold for any RMON recorded item. Examples of alarm settings include absolute or relative values, rising or falling thresholds of utilization, packet counts, and CRC errors.

The event group allows events (generated traps) to be logged, printed, and provided to a network manager. The time and date is recorded with each logged event. Network managers use the event group to create customized reports based on alarm types.

Extended RMON capabilities are provided through the use of a Cisco SwitchProbe connected to the switch's SPAN port. Refer to the section, "Switched Port Analyzer," for additional information.

To configure a Cisco switch for RMON, activate SNMP remote monitoring support via the **set snmp rmon enable** command. After entering the **set snmp rmon enable** command, you will see the display in Example 13-11.

Example 13-11 `set snmp rmon enable` Command Output

```
Console> (enable) set snmp rmon enable  
SNMP RMON support enabled.
```


Switched Port Analyzer

Cisco switches have a Switched Port Analyzer (SPAN) feature enables you to monitor traffic on any port for analysis by a network analyzer device or RMON probe. This feature also provides RMON2 statistics on all nine RMON groups and all seven layers of the OSI model. Enhanced SPAN (E-SPAN) enables you to monitor traffic from multiple ports with the same VLAN to a port for analysis.

The SPAN redirects traffic from an Ethernet, Fast Ethernet, or Fiber Distributed Data Interface (FDDI) port or VLAN to an Ethernet or Fast Ethernet monitor port for analysis and troubleshooting. You can monitor a single port or VLAN using a dedicated analyzer such as a Network Associates Sniffer, or an RMON probe, such as a Cisco SwitchProbe. Figure 13-6 is an example of the SPAN feature on the Catalyst 5000 series switch.

A more recent feature is called R-SPAN, which allows for the monitoring of a remote switch's traffic. R-SPAN is available on the Catalyst 6000 series switches. It can only be used in a switched network solely of Catalyst 6000 switches. In other words, no other type of switches, Cisco or otherwise, can be in the direct path between the two switches.

In this configuration, all traffic on Ethernet port 5 is mirrored onto the configured SPAN port Ethernet 10. The network analyzer located on Ethernet 10 can see network traffic on Ethernet 5 without being physically attached to it.

Figure 13-6 SPAN Configuration on a Catalyst 5000 Series Switch

Example 13-12 shows how to display SPAN information.

Example 13-12 Displaying SPAN Information

```

Console> show span
Status      : enabled
Admin Source: VLAN 1
Oper Source : None
Destination : Port 1/1
Direction  : transmit/receive
Console>

```

The following list defines the **show span** command output fields in Example 13-12.

- **Admin Source**—Source port or VLAN for SPAN information.
- **Oper Source**—Operator port or VLAN for SPAN information.
- **Destination**—Destination port for SPAN information.
- **Direction**—Status of whether transmit, receive, or transmit/receive information is monitored.
- **Status**—Status of whether SPAN is enabled or disabled.

CiscoWorks 2000

CiscoWorks 2000 is an integrated management solution for Cisco networks. For our purposes here, we are strictly concerned with the LAN Management Solution, which is just part of the overall architecture of CiscoWorks 2000.

CiscoWorks 2000 provides configuration, administration, monitoring, and troubleshooting tools for the campus. This includes topology maps, configuration services, and important system, device, and performance information. CiscoWorks 2000 can be integrated with popular SNMP management platforms, such as HP OpenView, for seamless management of complex networks. Additionally, CiscoWorks 2000 solutions can be used independently of these SNMP management applications and do not require these services to be fully functional.

Table 13-3 defines the various features of CiscoWorks 2000 LAN Management in greater detail.

Table 13-3 *CiscoWorks 2000 LAN Management Features*

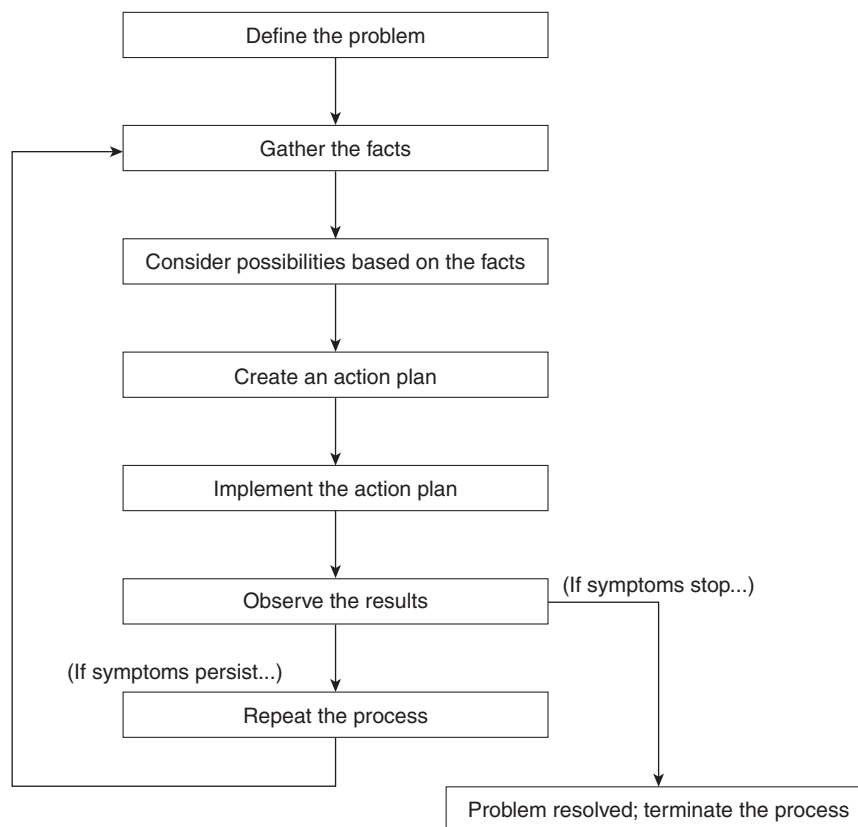
CiscoWorks 2000 Feature	What It Is/Does
Campus Bundle for ATM and LANE	This product is an updated version of the former ATM Director. The Campus Bundle offers network discovery and display, ATM and LANE configuration, user tracking, LAN/WAN traffic, and performance management capabilities on a device and network-wide basis.
CiscoView	A graphical management application providing dynamic status, statistics, and comprehensive configuration information for local or remote Cisco internetworking products. CiscoView displays a physical view of a device backplane, with graphs and color-coding for at-a-glance status and to display performance and other statistics. In addition, CiscoView has the ability to modify configurations such as trap, IP route, virtual LAN (VLAN), and bridge configurations.
Campus Manager	Campus Manager features include: intelligent discovery and display of large Layer 2 networks on browser-accessible topology maps; configuration of VLAN/LANE and ATM services and assignment of switch ports to those services link and device status display based upon SNMP polling; identification of Layer 2 configuration discrepancies; diagnostic tools for connectivity related problems between end stations, and Layer 2 and Layer 3 devices; automatic location and correlation of information on users by media access control (MAC), IP address, NT or NetWare Directory Services (NDS) login or UNIX hostname, with their physical connections to the switched network.
TrafficDirector	Offers graphical reporting and analysis of RMON collected traffic data both from RMON enabled Catalyst switches and from external SwitchProbes, which are also available from Cisco.
Resource Manager Essentials	A suite of Web-based applications offering network management solutions for Cisco switches, access servers, and routers. The suite consists of Inventory Manager, Change Audit, Device Configuration Manager, Software Image Manager, Availability Manager Syslog Analyzer, and Cisco Management Connection.

General Troubleshooting Model

When you're troubleshooting a network environment, a systematic approach works best. Define the specific symptoms, identify all potential problems that could be causing the symptoms, and then systematically eliminate each potential problem (from most likely to least likely) until the symptoms disappear.

Figure 13-7 illustrates the process flow for the general problem-solving model. This process flow is not a rigid outline for troubleshooting an internetwork; it is a foundation from which you can build a problem solving process to suit your particular environment.

Figure 13-7 *General Problem Solving Model*



The following steps detail the problem solving process outlined in Figure 13-7:

- Step 1** When analyzing a network problem, make a clear problem statement. You should define the problem in terms of a set of symptoms and potential causes. To do this, identify the general

symptoms and then ascertain what kinds of problems (causes) could result in these symptoms. For example, hosts might not be responding to service requests from clients (a symptom). Possible causes might be a misconfigured host, bad interface cards, or missing router configuration commands.

- Step 2** Gather the facts you need to help isolate possible causes. Ask questions of affected users, network administrators, managers, and other key people. Collect information from sources such as network management systems, protocol analyzer traces, output from router diagnostic commands, or software release notes.
- Step 3** Consider possible problems based on the facts you gathered. Using the facts you gathered, you could eliminate potential problems from your list. For example, depending on the data, you might be able to eliminate hardware as a problem allowing you to focus on software problems. At every opportunity, try to narrow the number of potential problems so you can create an efficient plan of action.
- Step 4** Create an action plan based on the remaining potential problems. Begin with the most likely problem and devise a plan in which only *one* variable is manipulated. This approach allows you to reproduce a given solution to a specific problem. If you alter more than one variable simultaneously, you might solve the problem, but identifying the specific change that eliminated the symptom becomes more difficult.
- Step 5** Implement the action plan, performing each step carefully while testing to see if the symptom disappears.
- Step 6** Whenever you change a variable, be sure to gather results. Generally, you should use the same method of gathering facts that you used in Step 2. Analyze the results to determine whether the problem has been resolved. If it has been resolved, then the process is complete.
- Step 7** If the problem has not been resolved, you must create an action plan based on the next most likely problem in your list. Return to Step 4 and reiterate the process until the problem is solved. Make sure to undo any “fixes” you made in implementing your action plan. Remember that you want to change only one variable at a time.

Troubleshooting with show Commands

Enter the **show system** command to display the power supply, fan, temperature alarm, system, and modem status; the number of days, hours, minutes, and seconds since the last system restart; the baud rate; the MAC address range; and the system name, location, and contact. Example 13-13 demonstrates typical information displayed by entering the **show system** command.

Example 13-13 show system Command Output

```

Console> show system
PS-Status Fan-Status Temp-Alarm Sys-Status Uptime d,h:m:s
-----
ok         ok         off        ok         27,17:05:50
Modem      Baud      MAC-Address-Range
-----
disabled  9600 00-04-0b-a0-04-1f to 00-04-0b-a0-05-56
System Name          System Location          System Contact
-----
WBU-Catalyst-5000 5      Closet 202 1/F          Luis x5529

```

Table 13-4 documents **show** commands that can assist in troubleshooting hardware, configuration, or network problems in a switched network environment.

Table 13-4 show Commands to Aid in Troubleshooting

Command	What the Output Displays
show arp	The contents of the ARP table and aging time.
show atm	The ATM interfaces, traffic, VC and VLAN information and status.
show cam dynamic	The dynamic CAM table.
show config	The current system configuration.
show fddi	The settings of the FDDI/CDDI module.
show flash	The Flash code names, version numbers, and sizes.
show interface	The Supervisor module network interface information.
show ip route	The IP route information.
show log	The system or module error log.
show mac	The MAC counters for all the installed modules.
show module	Module status and information (hardware/firmware/software).
show netstat	Statistics for the various TCP/IP stack protocols and state of active network connections.
show port	The port status and counters for all installed modules.
show spantree	The Spanning Tree information for the VLANs, including port states.

Table 13-4 *show* Commands to Aid in Troubleshooting (Continued)

show system	The status of the power supply, fan, temperature alarm, system, and uptime.
show test	The results of diagnostic tests on the specified modules.
show trunk	The ISL/Dot1Q information including trunking status (trunking/nontrunking).
show vlan	The virtual LAN type, status and assigned modules and ports.

Physical Layer Troubleshooting

The most common network problems can be traced to cable problems. The following questions will help determine whether there is a UTP cable problem.

- Are the cables the correct type for this installation?
 - Category 3 cabling can only support 10BaseT. Was a Category 3 cable installed instead of a category 5?
 - For Category 5 cabling, was the cable installed correctly? Severe bends in a Category 5 cable can cause a 10/100-Mbps interface to run at 10 Mbps. Some devices do not handle auto negotiation correctly. Check whether a 10/100-Mbps connection is connected at 10 Mbps instead of 100 Mbps.
- Is the cable a crossover or straight-through? Which type should it be? Compare the RJ-45 connector wiring at both ends of the cable, including all wiring closet connections.
- Is the punchdown wiring correct? Are there missing, loose, or broken wires on the punch-down block?
- One of the first ways to determine whether a cable is installed correctly is to check the devices' port link integrity LED on both ends of the cable. Each device transmits a link integrity pulse to the other device.
 - If the link LED is not on, try another port.
 - Is the other device is powered up? Is the link LED lit on the other device?
- Is there a broken wire at either end of the cable?
 - Cables that are installed too tightly with a tie wrap may have broken wires in the connectors.
 - Cables that are pulled through plenum can have broken wires and exhibit intermittent open circuit conditions.

Troubleshooting Ethernet

This section provides troubleshooting procedures for common Ethernet media problems. Table 13-5 outlines problems commonly encountered on Ethernet networks and offers general guidelines for solving those problems.

Table 13-5 *Media Problems: Ethernet*

Media Problem	Suggested Actions
Excessive noise	<p>Step 1 Use the show interfaces ethernet EXEC command to determine the status of the router's Ethernet interfaces. The presence of many CRC errors but not many collisions is an indication of excessive noise.</p> <p>Step 2 Check cables to determine whether any are damaged.</p> <p>Step 3 Look for badly spaced taps causing reflections.</p> <p>Step 4 If you are using 100BaseTX, make sure you are using Category 5 cabling and not another type, such as category 3.</p>
Excessive collisions	<p>Step 1 Use the show interfaces ethernet command to check the rate of collisions. The total number of collisions with respect to the total number of output packets should be around 0.1 percent or less.</p> <p>Step 2 Use a time domain reflectometer (TDR) to find any unterminated Ethernet cables.</p> <p>Step 3 Look for a jabbering transceiver attached to a host. (This might require host-by-host inspection or the use of a protocol analyzer.)</p> <p>Full-duplex links should never have collisions.</p>
Excessive runt frames	<p>In a shared Ethernet environment, runt frames are almost always caused by collisions. If the collision rate is high, refer to the problem "Excessive collisions" earlier in this table.</p> <p>If runt frames occur when collisions are not high or in a switched Ethernet environment, then they are the result of underruns or bad software on a network interface card.</p> <p>Use a protocol analyzer to try to determine the source address of the runt frames.</p>
Late collisions ¹	<p>Step 1 Use a protocol analyzer to check for late collisions. Late collisions should never occur in a properly designed Ethernet network. They usually occur when Ethernet cables are too long or when there are too many repeaters in the network.</p> <p>Step 2 Check the diameter of the network and make sure it is within specification.</p>

Table 13-5 *Media Problems: Ethernet (Continued)*

No link integrity on 10BaseT, 100BaseT4, or 100BaseTX	Step 1	Make sure you are not using 100BaseT4 when only two pairs of wire are available. 100BaseT4 requires four pairs.
	Step 2	Check for 10BaseT, 100BaseT4, or 100BaseTX mismatch (for example, a card different than the port on a hub).
	Step 3	Determine whether there is cross-connect (for example, be sure straight-through cables are not being used between a station and the hub).
	Step 4	Check for excessive noise (see the problem “Excessive noise” previously in this table).

¹A late collision is a collision that occurs beyond the first 64 bytes of an Ethernet frame.

Network Testing

One of the most useful and important troubleshooting aids when performing network testing is the **ping** command. Enter the **ping** command to send Internet Control Message Protocol (ICMP) echo request packets to another node on the network to confirm the connection to that node. Enter **Ctrl-C** to stop pinging.

```
ping -s host [packet_size] [packet_count]
```

The syntax descriptions for each parameter are as follows:

- **-s**—Causes **ping** to send one datagram per second, printing one line of output for every response received. The **ping** command does not return any output when no response is received.
- *host*—The IP address or IP alias of the host.
- *packet_size*—This optional parameter represents the number of bytes in a packet, from 1 to 2,000 bytes, with a default of 56 bytes. The actual packet size is eight bytes larger because the switch adds header information.
- *packet_count*—This optional parameter represents the number of packets to send.

Following are sample results of the **ping** command:

- **Normal response**—The normal response occurs in one to ten seconds depending on network traffic.
- **Destination does not respond**—If the host does not respond, a “no answer” message appears in ten seconds.
- **Destination unreachable**—The gateway given in the route table for this destination indicates that the destination is unreachable.
- **Network or host unreachable**—The switch found no corresponding entry in the route table.

In Example 13-14, a host with IP alias **elvis** is **pinged** a single time, then **pinged** once per second until **Ctrl-C** is entered to stop **pinging**:

Example 13-14 ping Example

```

Console> ping elvis
elvis is alive
Console> ping -s elvis
ping elvis: 56 data bytes
64 bytes from elvis: icmp_seq=0. time=11 ms
64 bytes from elvis: icmp_seq=1. time=8 ms
64 bytes from elvis: icmp_seq=2. time=8 ms
64 bytes from elvis: icmp_seq=3. time=7 ms
64 bytes from elvis: icmp_seq=4. time=11 ms
64 bytes from elvis: icmp_seq=5. time=7 ms
64 bytes from elvis: icmp_seq=6. time=7 ms
^C
----elvis PING Statistics----
7 packets transmitted, 7 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 7/8/11
Console>

```

If when using the **ping** command you do not get a response the following items should be evaluated:

- Is there a cable problem?
- Can the testing terminal (or workstation) communicate with any other devices?
- Can any other devices communicate with the switch?
- Are there port errors on the switch? (Check with the **show port** command.)
- Is there normal port traffic? (Check with the **show mac** command.)
- Are the port and interface sc0 in the same VLAN? (Check with the **show port** and **show int** commands.) The consequence is that you won't be able to **ping** each other.

Traceroute

The traceroute command can be very useful in troubleshooting by determining where along a particular network path a particular problem might be. For instance if connectivity cannot be totally achieved on a multihop route, using traceroute can determine where the packets stop.

```

traceroute [-n] [-w wait_time] [-i initial_ttl] [-m max_ttl] [-p dest_port] [-q
nqueries] [-t tos] host [data_size]

```

Use the **traceroute** command to display a hop-by-hop path through an IP network from the switch to a specific destination host. Table 13-6 documents the traceroute command parameters.

Table 13-6 *traceroute Command Parameters*

Parameter	Description
-n	(Optional) Prevents traceroute from performing a DNS lookup for each hop on the path. Only numerical IP addresses are printed.
-w wait time	(Optional) Specifies the amount of time (in seconds) that traceroute will wait for an ICMP response message. The allowed range for wait time is 1 to 300 seconds; the default is five seconds.
-i initial ttl	(Optional) Causes traceroute to send ICMP datagrams with a TTL value equal to initial_ttl instead of the default TTL of 1. This causes traceroute to skip processing for hosts that are less than initial_ttl hops away.
-m max ttl	(Optional) Specifies the maximum TTL value for outgoing ICMP datagrams. The allowed range for max_ttl is 1 to 255; the default value is 30.
-p dest port	(Optional) Specifies the base UDP destination port number used in traceroute datagrams. This value increments each time a datagram is sent. The allowed range for dest_port is 1 to 65535; the default base port is 33434. Use this option in the unlikely event that the destination host is listening to a port in the default traceroute port range.
-q nqueries	(Optional) Specifies the number of datagrams to send for each TTL value. The allowed range for nqueries is 1 to 1000; the default is three.
-t tos	(Optional) Specifies the TOS to be set in the IP header of the outgoing datagrams. The allowed range for tos is 0 to 255; the default is 0. Use this option to see if different types of service cause routes to change.
<i>host</i>	IP alias or IP address in dot notation (a.b.c.d) of the destination host.
<i>data size</i>	(Optional) Number of bytes, in addition to the default of 40 bytes, of the outgoing datagrams. The allowed range is 0 to 1420; the default is 0.

Network Test Equipment

In many situations, third-party diagnostic tools can be more useful than commands that are integrated into the router. For example, enabling a processor-intensive **debug** command can be disastrous in an environment experiencing excessively high traffic levels. However, attaching a network analyzer to the suspect network is less intrusive and is more likely to yield useful information without interrupting the operation of the router.

The following are some typical third-party troubleshooting tools used for troubleshooting internetworks:

- Volt-ohm Meters, digital multimeters, and cable testers.
- Time domain reflectometers (TDRs) and optical time domain reflectometers (OTDRs).
- Breakout Boxes, Fox Boxes, and bit/block error rate testers (BERTs/BLERTs).
- Network monitors.
- Network analyzers.

Volt-Ohm Meters, Digital Multimeters, and Cable Testers

Volt-ohm meters and digital multimeters are at the lower end of the spectrum of cable testing tools. These devices measure parameters such as AC and DC voltage, current, resistance, capacitance, and cable continuity. They are used to check physical connectivity.

Cable testers (scanners) also enable you to check physical connectivity. Cable testers are available for shielded twisted-pair (STP), unshielded twisted-pair (UTP), 10BaseT, coaxial cables, and twinax cables. A given cable tester might be able to perform any of the following functions:

- Test and report on cable conditions, including near-end crosstalk (NEXT), attenuation, and noise.
- Perform time domain reflectometer (TDR), traffic monitoring, and wire map functions.
- Display media access control (MAC) layer information about LAN traffic, provide statistics such as network utilization and packet error rates, and perform limited protocol testing (for example, TCP/IP tests such as **ping**).

Similar testing equipment is available for fiber-optic cable. Due to the relatively high cost of fiber cable and its installation, fiber-optic cable should be tested both before installation (on-the-reel testing) and after installation. Continuity testing of the fiber requires either a visible light source or a reflectometer. Light sources capable of providing light at the three predominant wavelengths (850 nanometers (nm), 1300 nm, and 1550 nm) are used with power meters that can measure the same wavelengths and test attenuation and return-loss in the fiber.

TDRs and OTDRs

At the top end of the cable-testing spectrum are time domain reflectometers (TDRs). These devices can quickly locate open and short circuits, crimps, kinks, sharp bends, impedance mismatches, and other defects in metallic cables.

A TDR works by bouncing a signal off the end of the cable. Opens, shorts, and other problems reflect the signal back at different amplitudes depending on the problem. A TDR measures how much time it takes for the signal to reflect and calculates the distance to a fault in the cable. TDRs can also be used to measure the length of a cable. Some TDRs can also calculate the propagation rate based on a configured cable length.

Fiber-optic measurement is performed by an optical time domain reflectometer (OTDR). OTDRs can accurately measure the length of the fiber, locate cable breaks, measure the fiber attenuation, and measure splice or connector losses. An OTDR can be used to take the “signature” of a particular installation, noting attenuation and splice losses. This baseline measurement can then be compared with future signatures when a problem in the system is suspected.

Breakout Boxes, Fox Boxes, and BERTs/BLERTs

Breakout boxes, fox boxes, and bit/block error rate testers are digital interface testing tools used to measure the digital signals present at PCs, printers, modems, CSU/DSUs, and other peripheral interfaces. These devices can monitor data line conditions, analyze and trap data, and diagnose problems common to data communication systems. Traffic from data terminal equipment (DTE) through data communications equipment (DCE) can be examined to help isolate problems, identify bit patterns, and ensure that the proper cabling has been installed. These devices, however, cannot test media signals such as Ethernet, Token Ring, or FDDI.

Network Monitors

Network monitors continuously track packets crossing a network, providing an accurate picture of network activity at any moment or a historical record of network activity over a period of time. They do not decode the contents of frames. Monitors are useful for baselining, in which the activity on a network is sampled over a period of time to establish a normal performance profile, or baseline.

Monitors collect information such as packet sizes, the number of packets, error packets, overall usage of a connection, the number of hosts and their MAC addresses, and details about communications between hosts and other devices. This data can be used to create profiles of LAN traffic as well as assist in locating traffic overloads, planning for network expansion, detecting intruders, establishing baseline performance, and distributing traffic more efficiently.

Network Analyzers

A network analyzer (also called a protocol analyzer) decodes the various protocol layers in a recorded frame and presents them as readable abbreviations or summaries, detailing which layer is involved (physical, data link, and so forth) and what function each byte or byte content serves.

Most network analyzers can perform many of the following functions:

- Filter traffic that meets certain criteria so that, for example, all traffic to and from a particular device can be captured.
- Time-stamp captured data.
- Present protocol layers in an easily readable form.
- Generate frames and transmit them onto the network.
- Incorporate an “expert” system in which the analyzer uses a set of rules, combined with information about the network configuration and operation, to diagnose and solve or offer potential solutions to network problems.

Foundation Summary

The Foundation Summary is a collection of tables and figures that provides a convenient review of many key concepts in this chapter. For those of you already comfortable with the topics in this chapter, this summary could help you recall a few details. For those of you who just read this chapter, this review should help solidify some key facts. For any of you doing your final prep before the exam, these tables and figures will hopefully be a convenient way to review the day before the exam.

Table 13-7 *show* Commands to Aid in Troubleshooting

Command	What The Output Displays
show arp	The contents of the ARP table and aging time.
show atm	The ATM interfaces, traffic, VC and VLAN information and status.
show cam dynamic	The dynamic CAM table.
show config	The current system configuration.
show fddi	The settings of the FDDI/CDDI module.
show flash	The Flash code names, version numbers, and sizes.
show interface	The Supervisor module network interface information.
show ip route	The IP route information.
show log	The system or module error log.
show mac	The MAC counters for all the installed modules.
show module	Module status and information (hardware/firmware/software).
show netstat	Statistics for the various TCP/IP stack protocols and state of active network connections.
show port	The port status and counters for all installed modules.
show spantree	The Spanning Tree information for the VLANs, including port states.
show system	The status of the power supply, fan, temperature alarm, system, and uptime.
show test	The results of diagnostic tests on the specified modules.
show trunk	The ISL/Dot1Q information including trunking status (trunking/nontrunking).
show vlan	The virtual LAN type, status and assigned modules/ports.

Table 13-8 *CiscoWorks 2000 LAN Management Features*

CiscoWorks 2000 Feature	What It Is/Does
Campus Bundle for ATM and LANE	This product is an updated version of the former ATM Director. The Campus Bundle offers network discovery and display, ATM and LANE configuration, user tracking, LAN/WAN traffic, and performance management capabilities on a device and network-wide basis.
CiscoView	A graphical management application providing dynamic status, statistics, and comprehensive configuration information for local or remote Cisco internetworking products. CiscoView displays a physical view of a device backplane, with graphs and color-coding for at-a-glance status and to display performance and other statistics. In addition, CiscoView has the ability to modify configurations such as trap, IP route, virtual LAN (VLAN), and bridge configurations.
Campus Manager	Campus Manager features include: intelligent discovery and display of large Layer 2 networks on browser-accessible topology maps; configuration of VLAN/LANE and ATM services and assignment of switch ports to those services link and device status display based upon SNMP polling; identification of Layer 2 configuration discrepancies; diagnostic tools for connectivity related problems between end stations, and Layer 2 and Layer 3 devices; automatic location and correlation of information on users by media access control (MAC), IP address, NT or NetWare Directory Services (NDS) login or UNIX hostname, with their physical connections to the switched network.
TrafficDirector	Offers graphical reporting and analysis of RMON collected traffic data both from RMON enabled Catalyst switches and from external SwitchProbes, which are also available from Cisco.
Resource Manager Essentials	A suite of Web-based applications offering network management solutions for Cisco switches, access servers, and routers. The suite consists of Inventory Manager, Change Audit, Device Configuration Manager, Software Image Manager, Availability Manager Syslog Analyzer, and Cisco Management Connection.

Q&A

The questions and scenarios in this book are more difficult than what you should experience on the actual exam. The questions do not attempt to cover more breadth or depth than the exam; however, they are designed to make sure that you know the answer. Rather than allowing you to derive the answer from clues hidden inside the question itself, the questions challenge your understanding and recall of the subject. Questions from the “Do I Know This Already?” quiz from the beginning of the chapter are repeated here to ensure that you have mastered the chapter’s topic areas. Hopefully, these questions will help limit the number of exam questions on which you narrow your choices to two options and then guess.

The answers to these questions can be found in Appendix A, on page 477.

- 1 What is the main method of out-of-band management for Cisco switches?

- 2 What is an application that uses SNMP to perform in-band management?

- 3 CDP operates at what layer of the OSI model?

- 4 What is the command to verify that RMON is enabled on the switch?

- 5 Using a troubleshooting model, what step is generally taken after ascertaining all the facts?

6 What is the default value for the read-write community string?

7 How many simultaneous Telnet sessions are supported on a Cisco switch?

8 What command shows information on the modules installed in a particular switch?

9 Explain the function of the **show flash** command.

10 What command displays content addressable memory?

11 What is the command to display CDP information about neighboring systems?

12 What command would I use if I wanted to see information on spanning tree for VLAN1?

13 What physical layer troubleshooting tool is similar to “sonar” in that it bounces a signal to determine length.

14 What is the name of the tool that decodes the various protocols in captured packets?

15 Explain the function of the **show mac** command.

16 Explain the function of the **show config** command.

17 What command is used to display errors?

18 Explain the function of the **show port** command.

Scenarios

Scenario 13-1

For the following scenario, refer to the **show cdp neighbor** command output in Example 13-15.

Example 13-15 Scenario 13-1: *show cdp neighbor* Command Output

```
Cat9> show cdp neighbor

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
Port    Device-ID                Port-ID    Platform    Capability
-----
3/1     001906906(Cat2)           2/2        WS-C5000    TS
3/2     060002103(Cat1)           2/1        WS-C5500    TS
2/1     064000022(Top Cat)        3/1        WS-C5500    TS
2/1     064020070(Cat3)           4/1        WS-C5500    TS
Cat9>
```

- 1 What can you deduce about the topology of this switched network?
- 2 Are there any routers in this particular setup?
- 3 How many devices are downstream from the switch named Cat9?
- 4 Compare the output in Example 13-15 with the output in Example 13-16 that was captured 10 minutes later. What happened?

Example 13-16 Scenario 13-1: *show cdp neighbor* Command Output for Comparison

```
Cat9> show cdp neighbor

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
Port    Device-ID                Port-ID    Platform    Capability
-----
3/1     001906906(Cat2)           2/1        WS-C5000    TS
3/2     060002103(Cat1)           2/2        WS-C5500    TS
2/1     064000022(Top Cat)        3/1        WS-C5500    TS
2/1     064020070(Cat3)           4/1        WS-C5500    TS
Cat9>
```

Scenario 13-2

I have a VLAN 10 that has 12 ports assigned to it. I am experiencing problems on port 3/3, which is assigned to VLAN 10. You have physical access to the switch and a protocol analyzer. What steps would you take to assess the situation?

Scenarios Answers

Scenario 13-1 Answers

- 1 If we look at the output in Example 13-15, we can deduce that there are two Ethernet based interfaces that connected to Cat2 and Cat1 due to the fact that they were learned on two different ports (3/1 and 3/2). If we look at the other two switches, we see that they are both learned through port 2/1. This infers that this is an ATM LANE interface.
- 2 Looking at the capability field of the output in Example 13-15, we see that there are no devices capable of routing directly connected.
- 3 We don't really know how many could be downstream, just the ones that are directly connected, which is four switches.
- 4 If we compare the output in Example 13-15 to the output in Example 13-16, we see that the ports changed at the remote ends of both Cat1 and Cat2. This means that these switches were reconfigured to connect via different ports.

Scenario 13-2 Answers

Because you have access to the switch, the first step would be to inspect the cable connected to port 3/3. Assuming all is well with the cable, the next step would be to check the port status. Port status can be assessed by the following command:

```
Console>show port 3/3
```

Assuming you still don't have the answer to what's wrong, the next step would be to connect the protocol analyzer to an open port, let's say this is 4/1. You then need to activate SPAN to observe traffic on port 3/3. The command to do this is shown below:

```
Console>set span 4/1 3/3
```

You should be able to assess the situation by observing the traffic going to and from port 3/3 using the protocol analyzer.



Scenarios for Final Preparation

This chapter presents three scenarios that can be used to review most of the concepts contained in this book. The scenarios are designed to assist you in final preparation for the Cisco Switching Exam. Case studies are presented with network diagrams and questions covering many switching topics.

This reading emphasizes an overall understanding of switching concepts, configuration commands, and network operation. Although the Cisco Switching Exam may not contain scenarios of this type, this chapter will better prepare you by thinking about the “bigger picture” of a network and how each switching topic can be applied.

Scenario 14-1

Refer to the network diagram in Figure 14-1 and complete the following tasks. Assume that each of the Catalysts has a MAC address formed from its one letter name (Catalyst A = aa-aa-aa-aa-aa-aa, Catalyst B = bb-bb-bb-bb-bb-bb, and so on).

- 1 Assume that all Catalyst switches have the default VTP configurations. Which Catalyst will become a VTP server?
- 2 Catalyst A is configured with the following commands:

```
set vtp domain alpha
set vlan 101 VLAN101
set vlan 102 VLAN102
```

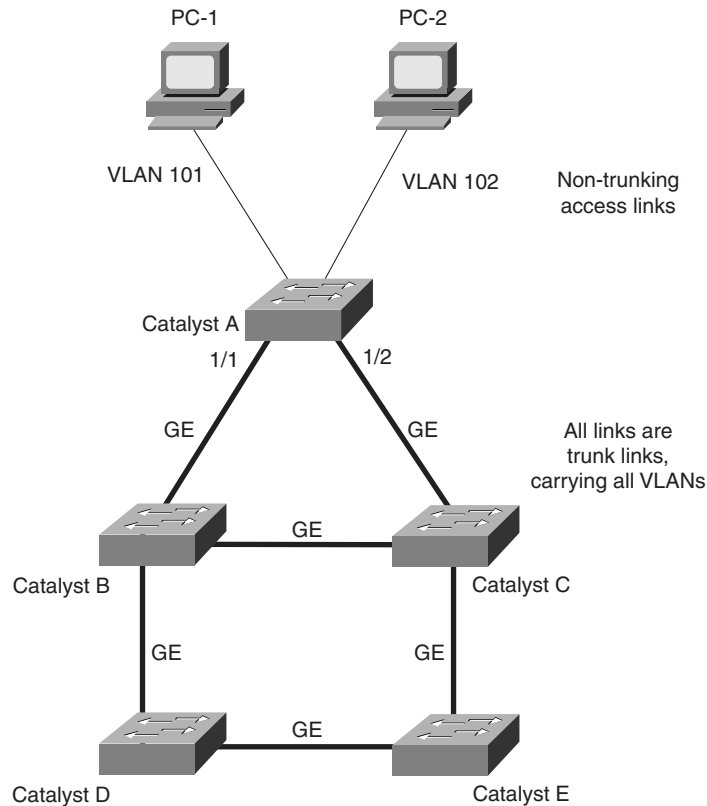
Which other Catalyst will learn of the new VLANs and VTP configuration?

- 3 To make sure VTP information is passed over trunk links, which VLAN(s) should be enabled over the trunks between Catalyst switches?
- 4 Configure Catalyst A to trunk only VLANs 101 and 102 over its 1/1 Gigabit Ethernet link, using ISL. Let Catalyst B take the active trunk negotiation role over the link. (Port 1/2 will be configured similarly.)
- 5 Suppose Catalyst A has the following commands added to it:

```
set vtp pruning enable
set vtp pruneeligible 101-110
```

Will broadcasts on VLAN 101 be sent across the trunk link between Catalyst A and Catalyst B?

Figure 14-1 Network Diagram for Scenario 14-1




- 6 Which Catalyst will become the root bridge of the Spanning Tree domain (assume only VLANs 101 and 102 are present on all switches)?
- 7 What command can force Catalyst B to become the root bridge for only VLANs 101 and 102?
- 8 Using STP commands, how can Catalyst A port 1/1 be put in the forwarding state to carry VLAN 101 traffic, while port 1/2 is put in the forwarding state for VLAN 102? (This also means that port 1/1 will be blocking for VLAN 102, and port 1/2 will be blocking for VLAN 101.) In effect, this setup achieves some load balancing by separating the two VLANs to pass over separate uplinks.
- 9 Suppose Catalyst B is the root bridge for all VLANs. What commands on Catalyst A could be used to achieve the same load balancing across ports 1/1 and 1/2 as in the previous question? (VLAN 101 should be transported over 1/1 and VLAN 102 over 1/2.)
- 10 What Spanning-Tree Protocol feature can be used to minimize the initialization delay on the Catalyst A ports where PCs are connected?
- 11 From Figure 14-1, where should the *UplinkFast* feature be enabled? Where should *BackboneFast* be enabled?

- 12 On Catalyst A, what command can be used to assign an IP address of 10.1.101.1 255.255.255.0 to VLAN 101 for Telnet purposes?
- 13 Catalyst E has an IP address of 10.1.254.1 255.255.255.0 assigned to its sc0 interface. However, neither Catalyst C nor Catalyst D can ping Catalyst E successfully. Running the **show cdp neighbor detail** command shows that Catalyst E is indeed alive and sending CDP information. As well, Catalyst E is shown to be using 10.1.254.1 as its IP address. What are some possible causes for the ping failure?
- 14 For this network, which is more appropriate for improved multicast performance and handling—IGMP snooping or CGMP?
- 15 Assume that a server is connected to Catalyst E port 3/3. What command can be used to monitor traffic transmitted and received on the server port with a network analyzer connected to Catalyst E port 3/8?

Scenario 14-2

Refer to the network diagram in Figure 14-2 and complete the following tasks. Assume that each of the Catalysts has a MAC address formed from its one letter name (Catalyst A = aa-aa-aa-aa-aa-aa, Catalyst B = bb-bb-bb-bb-bb-bb, and so on.)

Figure 14-2 Network Diagram for Scenario 14-2




- 1 The link between Catalyst A and Catalyst C should be a Gigabit EtherChannel using trunking for all available VLANs. What protocol is used to negotiate the EtherChannel link?
- 2 If Catalyst A has its EtherChannel defined on ports 2/1 and 2/2 in auto mode, what is necessary to bring the link into successful EtherChannel operation?
- 3 Assume that only two hosts are connected to Catalyst A. Host PC-1 (MAC address 11-11-11-11-11-11, IP address 10.1.101.1) communicates only with Server-1 (a1-a1-a1-a1-a1-a1, 10.1.101.200). Host PC-2 (22-22-22-22-22-22, 10.1.102.1) only talks with Server-2 (a2-a2-a2-a2-a2-a2, 10.1.102.200). Is the traffic balanced across the links of the EtherChannel trunk, assuming no router is involved?
- 4 Catalyst C is a Catalyst 5000 with an RSM. What commands are needed on the RSM to configure an interface for routing on VLAN 101, using the network 10.1.101.0 with subnet mask 255.255.255.0?
- 5 What is needed for the RSM in Catalyst C to exchange routing information with the RSM in Catalyst D?
- 6 Assume that all Catalysts are using VTP to configure VLANs in a single domain. VLANs 101 and 102 are the only access layer VLANs in the network and must be present on all switches. Where will LANE components be placed? How many of each will be required?
- 7 Suppose that Catalyst C will function as a LECS. How can Catalyst D obtain the LECS NSAP address? How can you (the network administrator) obtain the address?
- 8 What command on Catalyst B will allow Telnet access to the switch only from hosts on the 192.168.17.0 255.255.255.0 network?
- 9 Assume that all switches and route processors are using the default multicast configurations. Where in the network will multicast traffic originating from PC-1 on Catalyst A (VLAN 101) be seen?
- 10 What is needed to limit multicast traffic to only those ports that explicitly join multicast groups, using CGMP with PIM dense mode? Assume this same need on both VLAN 101 and 102.
- 11 Network administrators want to have tight control over hosts moving around within the network. Catalyst B needs to have port level security enabled on all 24 ports of its Module 2 line card. Only one host is connected per port, so the default behavior of shutting the port down is acceptable. What command would be necessary to do this?

Scenario 14-3

Refer to the network diagram in Figure 14-3 and complete the following tasks. Assume that each of the Catalysts has a MAC address formed from its one letter name (Catalyst A = aa-aa-aa-aa-aa-aa, Catalyst B = bb-bb-bb-bb-bb-bb, and so forth).

Figure 14-3 Network Diagram for Scenario 14-3



- 1 In what layer of the hierarchical campus design are Catalysts A and B located? Catalysts C and D? What type of core is present?
- 2 Load balancing is desirable across the two trunk links from Catalyst A to Catalysts C and D, and on the two trunk links from Catalyst B to Catalysts C and D. Given that the Access Layer functions at Layer 2 and the Core at Layer 3 (using MLS), what is the appropriate way to accomplish load balancing?
- 3 Suppose the PCs in VLAN 101 use address 192.168.101.1 as their default gateway. The PCs in VLAN 102 use 192.168.102.1. What commands are necessary to configure HSRP on the Catalyst C RSM so that it becomes the active router for VLAN 101 and the standby router for VLAN 102? If a failed router interface is restored, control should be passed back to it from the HSRP standby router. (You may use IP addresses 192.168.101.2 and 192.168.102.2 if needed.)

- 4 MLS is to be implemented on the network. What switch commands are needed on Catalyst C and D to allow MLS operation with the respective RSMs?
- 5 MLS is configured on both RSMs illustrated in Figure 14-3 for all supported VLANs. How many XTAG entries does Catalyst C have? (This situation assumes that Catalyst C has a NFFC module to process MLS information.)
- 6 What type of access list should be implemented on an RSM to achieve the specified type of MLS flow mask:
 - a. Destination-Source flow mask
 - b. Destination flow mask
 - c. Full flow mask
- 7 Using MLS, some degree of packet filtering is desired. For example, users in the 192.168.191.0 255.255.255.0 network should be allowed to use only HTTP (www) traffic into VLAN 180. What access list should be used to implement this filter and how should it be applied on an RSM?
- 8 What command can be used to view the current MLS cache table, along with outbound switch ports?
- 9 Suppose PC-1 needs to talk to PC-3. The first packet travels through Catalyst A, C, and on to B and PC-3. Then a reply is sent from PC-3 to Catalyst B, D, then to A and back to PC-1. (The paths described here are hypothetical and might not be experienced on an actual network.) Which Catalysts have identified *candidate* MLS entries and which ones have identified *enable* MLS entries?
- 10 Answer the following MLS timer questions:
 - a. What is the interval used to age out incomplete MLS shortcut entries?
 - b. How long are normal, complete shortcut entries cached by default?
 - c. For short-term flows (DNS, pings, and so on), what is the default fast aging time?
- 11 Port level security is desired on Catalyst B port 3/1, where 24 users are connected via an Ethernet hub. Rather than have the switch port shut down upon a security violation, network administrators want only the hosts in violation to be rejected. What command can be used to accomplish this?

Scenarios Answers

Scenario 14-1 Answers

- 1 All of the Catalysts default to server mode, for the VTP domain “null.”
- 2 All other Catalysts will learn of the changes. Because each server defaults to the domain “null,” each one will learn of the new domain “alpha,” join it as a server, and create VLANs 101 and 102.
- 3 VLAN 1 is always used to pass VTP information and should not be disabled on a trunk. VLAN 1 is also used for many switch-related protocols, including DTP, PAgP, and CDP. In fact, on Ethernet trunks you cannot remove VLAN 1 from the list of VLANs to trunk, unless Catalyst software Release 5.4 or greater is used. In Release 5.4 or greater, VLAN 1 can be removed from a trunk. Only the user traffic on VLAN 1 is actually removed, while the switch-related protocols still pass with VLAN 1 tagging.
- 4 The following commands could be used to accomplish this task:

```
set trunk 1/1 auto isl
clear trunk 1/1 2-1005
set trunk 1/1 101,102
```

- 5 Yes, broadcasts will be sent across the trunk link. Although Catalyst B has no PCs in VLAN 101, Catalyst B does have other trunk links that need to transport VLAN 101 (with all other VLANs). Therefore, Catalyst B informs A via VTP that it has a leaf node requiring VLAN 101.
- 6 Catalyst A will become the root for both VLAN 101 and VLAN 102 because it has the lowest MAC address and the lowest Bridge ID. (Remember that a separate Spanning-Tree Protocol runs for each VLAN—PVST.)
- 7 The following commands adjust the bridge priority to 100, much less (more desirable) than the default 32768 on the other Catalysts:

```
set spantree priority 100 101
set spantree priority 100 102
```

Another alternative is to use the following commands:

```
set spantree root 101
set spantree root 102
```

- 8 For this network, only one choice is available: deterministic root bridge placement. The root bridge for VLAN 101 must be Catalyst B (to cause forwarding on A port 1/1), and the root bridge for VLAN 102 must be Catalyst C (to cause forwarding on A port 1/2). The following commands can be used:

Catalyst B:

```
set spantree root 101
set spantree secondary 102
```

Catalyst C:

```
set spantree root 102
set spantree secondary 101
```

- 9 Because Catalyst B is the root bridge, all its links are Root Ports by definition (each has a Root Path Cost of 0). Therefore, the link from B to A (A's port 1/1), the link from B to D, and the link from B to C will all be in the forwarding state. From A's perspective, the Root Path Cost of the link from A to B (0+4) is lower than that of the link from A to C (0+4+4). Therefore, port 1/1 on A will become A's Root Port. On the link from A to C, port 1/2 on A will become the Designated Port for the shared segment, because A has the lowest BridgeID. C will place its link to A in the blocking state.

Adjusting the port cost for A's port 1/1 is necessary to make it undesirable only for VLAN 102. The command **set spantree portvlancost 1/1 1000 102** will make the Root Path Cost of port 1/1 much higher (0+1000) than that of port 1/2 (0+4+4). Port 1/2 will enter the forwarding state for VLAN 102, while port 1/1 will be blocking. (Notice that the port cost for 1/1 only needs to be one greater than the port cost for 1/2 to achieve the same result.)

- 10 **set spantree portfast *module/port* enable**
- 11 UplinkFast should only be enabled on Catalyst A—the leaf node of the network with multiple uplinks. BackboneFast, when used, should be enabled on every switch in the network.
- 12 **set interface sc0 10.1.101.1 255.255.255.0 101**
- 13 Catalyst E could have its sc0 interface assigned to the wrong VLAN. Because the diagram does not have a router, passing traffic between VLANs is impossible.

The VLAN assigned to Catalyst E's sc0 interface might not be trunked to Catalyst C or Catalyst D. Because all links between switches are trunk links in this network, the VLAN would have to be assigned to the trunks for connectivity.

VTP might not be configured correctly between the switches, such that additional VLANs are not being defined in common. For example, Catalyst E could have its sc0 assigned to VLAN 53, and Catalyst C's sc0 assigned to VLAN 53. Suppose that Catalyst C is in VTP transparent mode and Catalyst E is in VTP client mode. Without any other VTP servers, Catalyst E will not make VLAN 53 active.

Remember how useful the **show cdp neighbor** command is for troubleshooting. CDP uses multicasts at Layer 2 to exchange information and will work successfully even if IP addresses or subnet masks are misconfigured.

- 14 With IGMP snooping, a switch can listen to IGMP activity for itself without requiring router intervention.

CGMP is a much more efficient and scalable technology, especially when Layer 3 switches are in use. However, CGMP requires a router to process IGMP events and inform the switch.

15 set span 3/3 3/8 both

Scenario 14-2 Answers

- 1 Port Aggregation Protocol (PAgP)
- 2 Catalyst C must have its EtherChannel in *desirable* mode to initiate an EtherChannel negotiation with Catalyst A. In addition, all links in the EtherChannel bundle must be set to trunking mode with identical configurations.
- 3 To determine which link is used within an EtherChannel bundle, recall that the source and destination MAC addresses (by default) are exclusive-OR'd. Because this bundle has two-links, only the least significant bit of the XOR is used to index the EtherChannel link.

For the PC-1 to Server-1 conversation:

11-11-11-11-11-11 XOR a1-a1-a1-a1-a1-a1

Least significant bit:1 XOR1 = 0 (first link of bundle)

For the PC-2 to Server-2 conversation:

22-22-22-22-22-22 XOR a2-a2-a2-a2-a2-a2

Least significant bit:0 XOR0 = 0 (first link of bundle)

Therefore, both conversations will use the same link within the EtherChannel bundle. The EtherChannel will not be balanced.

- 4 **interface vlan 101**
ip address 10.1.101.1 255.255.255.0
no shutdown
- 5 An IP routing protocol (RIP, IGRP, EIGRP, OSPF, IS-IS, and so forth) is needed on both RSMs; LANE connectivity between the RSMs, and a common ELAN and LECs defined on both Catalyst C and Catalyst D.
- 6 First, recall that Catalyst switches use VLAN 1 to transport VTP information. By default, VLAN 1 exists on a Catalyst. However, LANE has no default ELAN configuration, so each ELAN must be defined. To support VLAN 1, an ELAN 1 must be configured to pass VTP across the ATM LANE cloud.
 - The LECs can be placed on either Catalyst C or D.
 - An LES/BUS pair for ELAN 1 on either C or D.

- An LES/BUS pair for ELAN 101 on either C or D.
- An LES/BUS pair for ELAN 102 on either C or D.
- An LEC for ELAN 1 on both C and D.
- An LEC for ELAN 101 on both C and D.
- An LEC for ELAN 102 on both C and D.

Grand total: 1 LECS, 3 LES/BUS pairs, and 6 LECs.

- 7 Catalyst D can use ILMI to request the LECS address from an ATM switch, hardcode the LECS address in its configuration, or use the well-known LECS NSAP address. ILMI is the most flexible and recommended. You may get the LECS NSAP address from Catalyst C itself, using the **show lane default** command.
- 8 **set ip permit 192.168.17.0 255.255.255.0**
- 9 By default, a switch must forward broadcast and multicast frames out all available ports on a VLAN. Therefore, the multicast traffic will be seen on all VLAN 101 ports on Catalyst A. In addition, Catalyst C and Catalyst D will bridge the multicast traffic over ELAN 101 via LANE. Finally, all VLAN 101 ports on Catalyst B will also forward the multicasts.
- 10 CGMP configuration is needed on both switches and routers. On a router, the following commands can be used:

```
ip multicast-routing
interface vlan 101
ip pim dense-mode
ip cgmp
interface vlan 102
ip pim dense-mode
ip cgmp
```

On a switch, only the following command is needed:

```
set cgmp enable
```

- 11 **set port security 2/1-24 enable**

Scenario 14-3 Answers

- 1 A and B are in the Access Layer, while C and D could be considered to be in the Distribution Layer or a combination Distribution/Core Layer. In the latter case, this scenario would be called a Collapsed Core.

- 2 Spanning Tree is ineffective in this network because the Layer 2 VLANs terminate at the Layer 3 route processors on Catalysts C and D. In other words, both links from each Access Layer switch are forwarding for all VLANs. Therefore we must exploit the Layer 3 technology to gain load balancing.

To do this, HSRP would be used on the route processors in Catalysts C and D. Catalyst C's RSM should become the primary default gateway for VLANs 101 and 201, while Catalyst D's RSM becomes the primary default gateway for VLANs 102 and 202. In this fashion, VLAN 101 traffic will travel over the A-C link; VLAN 201 traffic over the B-C link; VLAN 102 over the A-D link; and VLAN 202 over the B-D link. Also keep in mind that the RSMs should be configured as backup or secondary routers for the non-primary VLANs, should a primary router fail.

- 3 HSRP load balancing can be configured with the following RSM commands:

```
interface vlan 101
ip address 192.168.101.2 255.255.255.0
standby 101 priority 110
standby 101 preempt
standby 101 ip 192.168.101.1
interface vlan 102
standby 102 priority 100
standby 102 preempt
standby 102 ip 192.168.102.1
```

The default gateway address shared between the RSMs is configured as 192.168.101.1 for VLAN 101 and 192.168.102.1 for VLAN 102. In VLAN 101, the virtual interface has an IP address of 192.168.101.2. Two HSRP groups are defined, one for each VLAN.

Interface VLAN 101 will be the active router for VLAN 101 due to its higher priority of 110 (over a default of 100 on Catalyst D's RSM). If control is passed to the standby router, this router will assume control again through the use of the *preempt* command. For VLAN 102, the roles are reversed. With its lower priority of 100, this router will become the standby router in group 102. (Catalyst D's RSM will be configured with priority 110 for VLAN 102 to take the active router role.)

4 set mls enable

MLS is enabled by default on capable Catalyst switches. This default enables the switch to listen for MLSP messages from an MLS-capable router.

- 5 Catalyst C will only have two XTAG entries—one per MLS router. (Each XTAG entry contains all the MAC addresses for a single router's interfaces.)
- 6 a. Standard access list
b. No access list
c. Extended access list

```
7 access-list 101 permit tcp 192.168.191.0 0.0.0.255 any eq www
access-list 101 deny ip 192.168.191.0 0.0.0.255 any
access-list 101 permit ip any any
interface vlan 180
ip access-group 101 out
```

The first line of the access list allows HTTP traffic specifically from the 192.168.191.0 network. The second line denies any other traffic from this network. The third line permits any other traffic from any other network. The access list is applied to outbound traffic on the VLAN 180 interface, so that only traffic that the access list permits will be allowed into VLAN 180.

8 show mls entry

9 Candidate MLS packets were identified on Catalyst C, for the packet from PC-1 to PC-3. Candidates were also identified on Catalyst D, for the return packet from PC-3 to PC-1. However, no single Catalyst saw both candidate and enable packets because the path changed between the initial and return packets. Therefore, no Catalyst has identified enable packets, and no complete MLS entries exist.

- 10** a. Incomplete aging time is 5 seconds.
b. Normal aging time is 256 seconds.
c. Fast aging time is indefinite (set to 0 seconds by default).

```
11 set port security 3/1 enable
set port security 3/1 maximum 24
set port security 3/1 violation restrict
```

The first command line enables port level security on port 3/1. The second line configures port security to learn up to 24 MAC addresses dynamically on that port. The last line configures the switch to restrict any MAC addresses found to be in violation. The port will stay up, allowing the other users to communicate.



Answers to the “Do I Know This Already?” Quizzes and Q&A Sections

Answers to Chapter 2 “Do I Know This Already?” Quiz

- 1 *Describe the differences between Layer 2, Layer 3, and Layer 4 switching.*

In Layer 2 switching, frames are forwarded based on the Layer 2 source and destination MAC addresses. In Layer 3 switching, network layer source and destination addresses (IP, IPX, and so forth) are used. In Layer 4 switching, some application information is taken into account along with Layer 3 addresses. For IP, this information includes the port numbers from such protocol types as UDP and TCP.

- 2 *What is multilayer switching (MLS)?*

MLS forwards traffic using information from Layer 2, Layer 3, and Layer 4—all in hardware at wire speed.

- 3 *What is the 20/80 rule of networking?*

The 20/80 rule states that 20 percent of network traffic on a LAN segment will stay on that segment. The remaining 80 percent must go across the network core, either to enterprise servers or to the Internet.

- 4 *What is a collision domain? Where does it exist in a switched LAN?*

A collision domain is a network segment where shared media access is supported. Devices on the shared media must compete for access when transmitting data. In a switched network, the collision domain is restricted to a single switch port and does not extend across the switch.

- 5 *What is a broadcast domain? Where does it exist in a switched LAN?*

A broadcast domain is where broadcast frames propagate in a network. Basically, a broadcast domain covers an area where Layer 2 devices are located and terminates at the boundary of a Layer 3 device. In a switched network, the broadcast domain extends to all switch ports. This is because a switch forwards broadcasts out of all available ports in a VLAN.

6 *What is a VLAN, and why is it used?*

A VLAN (virtual LAN) is a group of switch ports that communicate as if they were attached to a single shared-media LAN segment. VLANs can be extended across buildings or backbones as long as the VLAN is connected end-to-end through trunking or physical connections. A VLAN is a broadcast domain that is used to segment networks for ease of management and better performance.

7 *In which OSI layer do devices in the distribution layer typically operate?*

Layer 3 devices are typically used in the distribution layer.

8 *How many layers are required in the hierarchical campus network design model?*

A campus network design could have two or three distinct layers, depending on the campus size. In smaller networks, the distribution and core layers are often merged into one layer with the access layer. In larger networks, all three layers can exist: access, distribution, and core.

9 *Which Cisco switch products should be used in the distribution layer of a campus network?*

Cisco Catalyst 2926G, 2948G-L3, 4908G-L3, 5000/5500, and 6000/6500 switch families all can be used in the distribution layer because they each offer Layer 3 switching functionality.

10 *When might a Catalyst 5000 be selected for use in a wiring closet? What attributes make it a good choice?*

A wiring closet indicates an access layer switch where high port density and low per-port cost are important. The Catalyst 5000 offers these attributes for a large user community served by the wiring closet. In addition, the 5000 family of switches offers access for a wide variety of network media (Ethernet, Token Ring, and so on).

11 *What building blocks are used to build a scalable campus network?*

The switch block and core block are both used to build a scalable network. In addition, the server block, WAN block, and mainframe block can be added to the design in certain networks.

12 *What are two types of core or backbone designs?*

Collapsed core and dual core.

Answers to Chapter 2 Q&A Section

1 *Where is the most appropriate place to connect a block of enterprise servers? Why?*

A block of enterprise servers should be connected into the core, just as switch blocks are. This maximizes connectivity from the servers to all other devices in the network. In effect, all users will see the same number of switch “hops” to access a server. Connecting into the core also provides maximum scalability as more server blocks can be added in the future.

2 *Describe the differences between Layer 2, Layer 3, and Layer 4 switching.*

In Layer 2 switching, frames are forwarded based on the Layer 2 source and destination MAC addresses. In Layer 3, network layer source and destination addresses (IP, IPX, and so on) are used. In Layer 4 switching, application information is taken into account along with Layer 3 addresses. For IP, this information includes protocol types such as UDP and TCP.

3 *What problems occur as switch blocks are added to a Layer 2 core design?*

Layer 3 router peering limits and equal-cost path limits can both be exceeded as the number of switch blocks increases.

4 *What is multilayer switching (MLS)?*

MLS forwards traffic using information from Layer 2, Layer 3, and Layer 4—all in hardware at wire speed.

5 *How can redundancy be provided at the switch and core block layers? (Consider physical means, as well as functional methods using protocols, algorithms, and so on.)*

In a switch block, redundancy can be provided through two distribution switches. Each access switch can be linked to both distribution switches for fault tolerance. The Spanning-Tree Algorithm will keep one of the two links blocked at all times. In the core block, a dual core can be used with two core switches. Each distribution switch has dual links, with one link to each core switch. Here, the redundant links can stay active for load sharing and redundancy. Instead of Spanning Tree, the routing protocols running on the Layer 3 distribution switches handle path selection and load balancing.

6 *What is the 20/80 rule of networking?*

The 20/80 rule states that 20 percent of network traffic on a LAN segment will stay on that segment. The remaining 80 percent must go across the network core either to enterprise servers or to the Internet.

7 *What factors should be considered when sizing a switch block?*

Traffic types, flows, and patterns, as well as the size and number of common workgroups. Additionally, the Layer 3 switching capacity in the distribution layer should be sized according to the amount of traffic crossing the core from one subnet or VLAN to another.

8 *What is a collision domain? Where does it exist in a switched LAN?*

A collision domain is a network segment where shared media access is supported. Devices on the shared media must compete for access when transmitting data. In a switched network, the collision domain is restricted to a single switch port and does not extend across the switch.

9 *What are the signs of an oversized switch block?*

The distribution switches begin to become bottlenecks in handling the volume of interVLAN traffic. Access list processing in the distribution can also become a rate-limiting factor. Broadcast and multicast traffic forwarding can slow down the Layer 2 and Layer 3 switches in the switch block.

10 *What is a broadcast domain? Where does it exist in a switched LAN?*

A broadcast domain is the extent of a network where broadcast frames propagate. Basically, a broadcast domain covers an area where Layer 2 devices are located and terminates at the boundary of a Layer 3 device. In a switched network, the broadcast domain extends to all switch ports. This is because a switch forwards broadcasts out all available ports in a VLAN.

11 *What are the attributes and issues of having a collapsed core block?*

Attributes: Cost savings (no separate high-end core switches) and design simplicity.
Issues: Scalability becomes limited.

12 *What is a VLAN, and why is it used?*

A VLAN (virtual LAN) is a group of switch ports that communicate as if they were attached to a single shared-media LAN segment. VLANs can be extended across buildings or backbones, as long as the VLAN is connected end-to-end through trunking or physical connections. A VLAN is a broadcast domain. VLANs are used to segment networks for ease of management and better performance.

13 *When would a Layer 3 core block be desirable or necessary?*

A Layer 3 core block would be desirable in a very large campus network. Router peering problems are overcome for networks with a large number of switch blocks.

14 *In which OSI layer do devices in the distribution layer usually operate?*

Layer 3 devices are typically used in the distribution layer.

- 15** *What is network segmentation? When is it necessary? How is it done in a campus network design?*

Segmentation is the process of dividing up a LAN into smaller, discrete collision domains. If a large percentage of collisions is observed on a LAN, segmentation is appropriate. In a campus network design, segmentation occurs at each switch port. A similar form of segmentation involves reducing the size of broadcast domains. Forming switch blocks (broadcast domains) that terminate at Layer 3 devices in the distribution layer does this.

- 16** *How many layers are required in the hierarchical campus network design model?*

A campus network design could have two or three distinct layers, depending on the campus size. In smaller networks, the distribution and core layers are often merged into one layer.

- 17** *How many switches are sufficient in a core block design?*

Usually two switches are sufficient in the core block, offering load sharing and redundancy. More core switches can be added as the size of the network and core traffic flow dictates.

- 18** *Which Cisco switch products should be used in the distribution layer of a campus network?*

Cisco Catalyst 2926G, 2948G-L3, 4908G-L3, 5000, and 6000 switch families can all be used in the distribution layer because they all offer Layer 3 switching functionality.

- 19** *List three methods used for Layer 3 switching in Cisco products.*

Multilayer Switching (MLS) with an integrated switch/router (Catalyst 5000 with RSM, Catalyst 6000 with MSFC); MLS with a switch and an external router (Catalyst 5000 with NFFC); and Cisco Express Forwarding (CEF) in hardware ASICs (Catalyst 2926G, 2948G-L3, 4908G-L3, 8500).

- 20** *When might a Catalyst 5000 be selected for use in a wiring closet? What attributes make it a good choice?*

A wiring closet indicates an access layer switch, where high port density and low per-port cost are important. The Catalyst 5000 offers these attributes for a large user community served by the wiring closet. In addition, the 5000 family of switches offers access for a wide variety of network media (Ethernet, Token Ring, and so on).

- 21** *Which Cisco Catalyst switches can be used in the access layer? (Consider the most important attributes of access layer switches.)*

Actually, most any Catalyst switch family can be used in the access layer. Low per-port cost and high port density are the two most important factors to consider. The choice of Catalyst series depends on the size of the workgroup or switch block. For example, the

Catalyst 1900, 2820, 2900XL, or 3500XL could be chosen for small groups (< 50) of users. A Catalyst 4000 might be selected for medium sized groups (< 100), and Catalyst 5000 or even 6000 are good candidates for large groups (> 100).

22 *What building blocks are used to build a scalable campus network?*

The switch block and core block are both used to build a scalable network. In addition, the server block, WAN block, and mainframe block can be added to the design in certain networks.

23 *Which Cisco switch family has the most scalable performance?*

The Catalyst 6000 series offers the most scalable performance, with a backplane bandwidth of 32 to 256 Gbps and multilayer switching rate of 15 to 150 Mpps.

24 *What are two types of core or backbone designs?*

Collapsed core and dual core.

25 *Which Cisco switch family is the most flexible for network media and translation?*

The Catalyst 5000 is the most flexible of all LAN switch families. The large choice of media modules allows this switch to offer “any-to-any” switching.

Answers to Chapter 3 “Do I Know This Already?” Quiz

1 *What are the different Ethernet technologies and their associated IEEE standards?*

Ethernet (10 Mbps, IEEE 802.3), Fast Ethernet (100 Mbps, IEEE 802.3u), and Gigabit Ethernet (1000 Mbps, IEEE 802.3z)

2 *What benefits result with switched Ethernet over shared Ethernet?*

Switched Ethernet ports receive dedicated bandwidth, have a reduced collision domain, and have increased performance due to segmentation or fewer users per port.

3 *At what layer are traditional 10 Mbps Ethernet, Fast Ethernet, and Gigabit Ethernet different?*

All Ethernet technologies share a common data link layer (IEEE 802.3), but have different physical layers.

4 *Describe Cisco’s EtherChannel technology.*

EtherChannel is a proprietary technique to aggregate or bundle several Fast Ethernet or Gigabit Ethernet links together. The resulting EtherChannel link appears to be a single physical connection. The benefits are load sharing across the bundled links, redundancy between the bundled links, and higher performance. Fast EtherChannel supports up to eight full-duplex links, or up to 1600 Mbps throughput. Gigabit EtherChannel also supports up to eight full-duplex links, or up to 16000 Mbps throughput.

- 5** *In a campus network, where is Fast Ethernet typically used? Where is Gigabit Ethernet typically used?*

Fast Ethernet is typically used for links between access layer and distribution layer devices, and between end users and the access layer devices. Gigabit Ethernet is typically used between all layers—between access and distribution layer devices, between distribution and core layer devices, and for the links between core layer devices.

- 6** *What is the maximum length of a Category 5 100BaseTX cable?*

A Category 5 100TX cable can be up to 100 meters in length.

- 7** *Name a type of Token Ring segmentation.*

Source-route bridging, source-route transparent bridging, or source-route switching.

- 8** *What part of a Token Ring frame specifies the exact path the frame should take to reach its destination?*

The Routing Information Field (RIF).

- 9** *What is the purpose of a Gigabit Interface Converter (GBIC)?*

A GBIC is used as a modular media-independent connection for Gigabit Ethernet. A switch with a GBIC port will accept GBIC modules that support various network media types. Changing network media cabling only requires a low-cost change of the GBIC module.

- 10** *What must be done to a switch before Telnet access is allowed?*

An IP address must be assigned to the management interface on the switch, and the management interface must be assigned to a VLAN.

- 11** *What type of user interface or command set does the Catalyst 5000 family of switches support? What type is the Catalyst 3500XL?*

Catalyst 5000 supports the CLI-based user interface, while the 3500XL supports the IOS-based user interface.

- 12** *What protocol is used by a Catalyst switch to learn about neighboring routers and switches?*

Cisco Discovery Protocol (CDP).

- 13** *What port speeds can be assigned to a Fast Ethernet switch port?*

Fast Ethernet ports support 10 Mbps, 100 Mbps, and Auto, for autonegotiation of speed and duplex mode.

- 14** *What port speeds can be assigned to a Token Ring switch port?*

Token Ring ports support 4 Mbps, 16 Mbps, and Auto for autosensing the speed of an existing ring or device.

Answers to Chapter 3 Q&A Section

- 1 *What are the different Ethernet technologies and their associated IEEE standards?*

Ethernet (10 Mbps, IEEE 802.3), Fast Ethernet (100 Mbps, IEEE 802.3u), and Gigabit Ethernet (1000 Mbps, IEEE 802.3z)

- 2 *What benefits result with switched Ethernet over shared Ethernet?*

Switched Ethernet ports receive dedicated bandwidth, have a reduced collision domain, and have an increased performance due to segmentation or fewer users per port.

- 3 *When a 10/100 Ethernet link is autonegotiating, which will be chosen if both stations can support the same capabilities—10BaseT full duplex, 100BaseTX half duplex, or 100BaseTX full duplex?*

100BaseTX full-duplex will be chosen because it has the highest autonegotiation priority and is common to both end stations.

- 4 *At what layer are traditional 10-Mbps Ethernet, Fast Ethernet, and Gigabit Ethernet different?*

All Ethernet technologies share a common data link layer (IEEE 802.3) but have different physical layers.

- 5 *Describe Cisco’s EtherChannel technology.*

EtherChannel is a proprietary technique to aggregate or bundle several Fast Ethernet or Gigabit Ethernet links together. The resulting EtherChannel link appears to be a single physical connection. The benefits are load sharing across the bundled links, redundancy between the bundled links, and higher performance. Fast EtherChannel supports up to 8 full-duplex links, or up to 1600-Mbps throughput. Gigabit EtherChannel also supports up to 8 full-duplex links or up to 16000-Mbps throughput.

- 6 *A switch port is being configured as shown below. What command is needed next to set the port to full-duplex mode?*

```
Switch(config-if)# interface FastEthernet 0/13  
Switch(config-if)#
```

Enter the command `duplex full` at the prompt. This command is an IOS-based switch, as evidenced by the interface configuration mode prompt shown.

- 7 *In a campus network, where is Fast Ethernet typically used? Where is Gigabit Ethernet typically used?*

Fast Ethernet is typically used for links between access layer and distribution layer devices, and between end users and the access layer devices. Gigabit Ethernet is typically used between distribution layer and core layer devices and for the links between core layer devices.

- 8 *What is the maximum length of a Category 5 100BaseTX cable?*

A Category 5 100TX cable can be up to 100 meters in length.

- 9 *A CLI-based switch port has been configured for 100 Mbps full-duplex mode, but a link cannot be established. What are some commands that could be used to investigate and correct the problem?*

To see the current state of one or more ports, you could use the **show port** command. This command would show the speed and duplex modes of the ports, as well as whether a link has been established. One reason the link is not established could be that the port is shutdown or disabled. To enable the port, use the **set port enable module/number** command. Because the port has been set to 100-Mbps full-duplex mode, it is possible that the end station can only support 10 Mbps at half-duplex. Therefore, the port should be set for autonegotiate mode with the **set port speed module/number auto** command. Otherwise, the port could be set to a fixed speed and mode that would match the end station.

- 10 *Name a type of Token Ring segmentation.*

Source-route bridging, source-route transparent bridging, or source-route switching.

- 11 *What part of a Token Ring frame specifies the exact path the frame should take to reach its destination?*

The Routing Information Field (RIF).

- 12 *What switch command will set the enable-mode password on an IOS-based switch? A CLI-based switch?*

For IOS-based switches, use **enable password level 15 password** to set the enable password. On CLI-based switches, use the **set enablepass password** command.

- 13 *What is the purpose of a GBIC?*

A GBIC is used as a modular media-independent connection for Gigabit Ethernet. A switch with a GBIC port will accept GBIC modules that support various network media types. Changing network media cabling only requires a low-cost change of the GBIC module.

- 14 *What CLI-based commands will allow Telnet and ping access to a switch management interface at 192.168.200.10, subnet mask 255.255.255.0, on VLAN 5? Now add a command to allow access between the switch and devices located off the local VLAN 5 subnet, using a router at 192.168.200.1.*

To set the IP address of the management interface, use the following command: **set interface sc0 192.168.200.10 255.255.255.0 192.168.200.255** and **set interface sc0 5** to assign the management interface to VLAN 5. To allow access to devices off VLAN 5, you should add the following command: **set ip route default 192.168.200.1**.

- 15 *What must be done to a switch before Telnet access is allowed?*

An IP address must be assigned to the management interface on the switch, and the management interface must be assigned to a VLAN.

- 16 *What factors determine the choice of a distribution layer switch, its access-to-distribution layer link media, and its Layer 3 processor?*

The distribution switch links into the core must be able to support the traffic load from all connected access layer switches and their end users. Therefore, the distribution switch should have the port density and port media required to support the access layer. In addition, the Layer 3 capabilities of the switch should be able to support a packet switching rate equal to or greater than the rate of packets coming from the access layer and destined across the core.

- 17 *What type of user interface or command set does the Catalyst 5000 family of switches support? What type is the Catalyst 3500XL?*

Catalyst 5000 supports the CLI-based user interface, while the 3500XL supports the IOS-based user interface.

- 18 *What protocol is used by a Catalyst switch to learn about neighboring routers and switches?*

Cisco Discovery Protocol (CDP).

- 19 *What switch command can be used to find the IP addresses of nearby Cisco switches on a network?*

On either style of user interface, the **show cdp neighbor detail** command will show all known neighboring Cisco switches and their management IP addresses.

- 20 *What port speeds can be assigned to a Fast Ethernet switch port?*

Fast Ethernet ports support 10 Mbps, 100 Mbps, and Auto, for autonegotiation of speed and mode.

- 21 *What is the purpose of switch clustering? Can clustered switches share switching loads with each other?*

Switch clustering is used to logically group up to 16 switches together as a unit, for management purposes. After a command switch is chosen and configured, other switches can be discovered and added into the cluster. The entire cluster can be monitored and managed from a web interface or from the command-line interface. Switch clustering has nothing to do with load sharing, as each switch still functions independently.

- 22 *What port speeds can be assigned to a Token Ring switch port?*

Token Ring ports support 4 Mbps, 16 Mbps, and Auto, for autosensing the speed of an existing ring or device.

Answers to Chapter 4 “Do I Know This Already?” Quiz

1 *What is a VLAN? When is it used?*

A VLAN is a group of devices on the same broadcast domain, as a logical subnet or segment. VLANs can span switch ports, switches within a switch block, or closets and buildings. VLANs are used to group users and devices into common workgroups across geographical areas. VLANs help provide segmentation, security, and problem isolation.

2 *What are two types of VLANs, in terms of spanning areas of the campus network?*

End-to-end (spans entire campus network) and local (spans local geographic area).

3 *Generally speaking, what must be configured (both switch and end user device) for a port-based VLAN?*

Only the switch needs configuring for a port-based VLAN; a port must be assigned to an existing VLAN. The end user device needs no configuration because the VLAN just appears to be a normal network segment.

4 *What are the components of a Token Ring VLAN?*

TrBRF (Token Ring Bridge Relay Function) and TrCRF (Token Ring Concentrator Relay Function).

5 *What is a trunk link?*

A trunk link is a connection between two devices that transports traffic from multiple VLANs. Each frame is identified with its source VLAN during its trip across the trunk link.

6 *What methods of VLAN frame identification can be used on a Catalyst switch?*

ISL and IEEE 802.1Q (ATM LANE and IEEE 802.10 can also be used, but are not discussed in this chapter.)

7 *What is the purpose of Dynamic Trunking Protocol (DTP)?*

DTP allows negotiation of a common trunking method between endpoints of a trunk link.

8 *What VTP modes can a Catalyst switch be configured for? Can VLANs be created in each of the modes?*

Server, Client, and Transparent modes. VLANs can be created in Server mode, but not in Client mode. In Transparent mode, VLANs can be created, but only on the local switch; they are not advertised to other switches.

- 9 *How many VTP management domains can a Catalyst switch participate in? How many VTP servers can a management domain have?*

A switch can be a member of only one VTP management domain. A VTP domain must have at least one server. There can be more than one server, for redundancy, but it is recommended to have no more than two.

- 10 *What conditions must exist for two Catalyst switches to be in the same VTP management domains?*

- (a) Both switches must have the same VTP domain name defined and enabled.
- (b) Both switches must be adjacent on a trunk link.
- (c) Trunking must be enabled and active between the adjacent switches on the trunk link.

- 11 *What is the purpose of VTP pruning?*

VTP pruning is used to control unnecessary flooding of traffic across trunk links.

- 12 *Which VLAN numbers are never eligible for VTP pruning? Why?*

VLAN numbers 1 and 1001–1005 are ineligible for pruning. VLAN 1 is reserved as the management VLAN, while VLANs 1002–1005 are reserved as the default FDDI and Token Ring function VLANs. At press time, VLAN 1001 has no special purpose, but is reserved and cannot be pruned.

Answers to Chapter 4 Q&A Section

- 1 *What is a VLAN? When is it used?*

A VLAN is a group of devices on the same broadcast domain, as a logical subnet or segment. VLANs can span switch ports, switches within a switch block, or closets and buildings. VLANs are used to group users and devices into common workgroups across geographical areas. VLANs help provide segmentation, security, and problem isolation.

- 2 *When a VLAN is configured on a Catalyst switch port, in how much of the campus network will the VLAN number be unique and significant?*

The VLAN number will be significant in the local switch. If trunking is enabled, the VLAN number will be significant across the entire trunking domain. In other words, the VLAN will be transported to every switch that has a trunk link supporting that VLAN.

- 3 *What are two types of VLANs, in terms of spanning areas of the campus network?*

End-to-end (spans entire campus network) and local (spans local geographic area).

- 4 *What is the Catalyst CLI-based switch command to configure ports 4/11 and 5/1 through 5/24 for VLAN 2?*

```
set vlan 2 4/11,5/1-24
```

- 5 *Generally speaking, what must be configured (both switch and end user device) for a port-based VLAN?*

Only the switch needs configuring for a port-based VLAN. A port must be assigned to an existing VLAN. The end user device needs no configuration because the VLAN just appears to be a normal network segment.

- 6 *What is the default VLAN on all ports of a Catalyst switch?*

VLAN 1

- 7 *What are the components of a Token Ring VLAN?*

TrBRF (Token Ring Bridge Relay Function) and TrCRF (Token Ring Concentrator Relay Function).

- 8 *What is a trunk link?*

A trunk link is a connection between two switches that transports traffic from multiple VLANs. Each frame is identified with its source VLAN during its trip across the trunk link.

- 9 *What methods of Ethernet VLAN frame identification can be used on a Catalyst switch?*

ISL and IEEE 802.1Q.

- 10 *What is the difference between these two trunking methods? How many bytes are added to trunked frames for VLAN identification in each method?*

ISL uses encapsulation and adds a 26-byte header and a 4-byte trailer. 802.1Q adds a 4-byte tag field within existing frames, without encapsulation.

- 11 *What is the purpose of Dynamic Trunking Protocol (DTP)?*

DTP allows negotiation of a common trunking method between endpoints of a trunk link.

- 12 *What CLI-based commands are needed to configure a Catalyst switch trunk port 1/1 to transport only VLANs 100, 200–205, and 300 using IEEE 802.1Q? (Assume that trunking is enabled and active on the port already.)*

```
clear trunk 1/1 1-1000
```

```
set trunk 1/1 on 100,200-205,300 dot1q
```

Other options could be used, as long as the neighboring switch is configured to match or negotiate with this one. Remember that the default VLANs 1-1000 must be cleared first, because you need to trunk only the VLANs listed.

- 13** *What VTP modes can a Catalyst switch be configured for? Can VLANs be created in each of the modes?*

Server, Client, and Transparent modes. VLANs can be created in Server mode. VLANs cannot be created in Client mode. In Transparent mode, VLANs can be created, but only on the local switch. VLANs are not advertised to other switches.

- 14** *Two neighboring switch trunk ports are set to auto mode with ISL trunking mode. What will the resulting trunk mode become?*

Trunking will not be established at all. Both switches are in the passive auto state and are each waiting to be asked to start the trunking mode. Instead, the link will remain an access link on both switches.

- 15** *How many VTP management domains can a Catalyst switch participate in? How many VTP servers can a management domain have?*

A switch can be a member of only one VTP management domain. A domain must have at least one server. There can be more than one server, for redundancy, but it is recommended to have no more than two.

- 16** *What CLI-based command can be used on a Catalyst switch to verify exactly what VLANs will be transported over a trunk link?*

show trunk

- 17** *What conditions must exist for two Catalyst switches to be in the same VTP management domains?*

Both switches must have the same VTP domain name defined and enabled.

Both switches must be adjacent on a trunk link.

Trunking must be enabled and active between the adjacent switches on the trunk link.

- 18** *What are the types of VTP messages or advertisements used by Catalyst switches? What field in these messages determines if a switch should use and record VLAN data in the messages?*

The VTP message types are Advertisement Requests, Summary Advertisements, Subset Advertisements, and VLAN Membership Advertisements (an extension to version 1 for VTP pruning). The Configuration Revision Number is used to determine if the VTP data is newer and should be used.

- 19** *What CLI-based command can be used to configure a Catalyst switch to become a VTP server for the domain “engineering”? The domain should be secured with the password “secret123.”*

set vtp domain engineering password secret123

set vtp mode server

20 *What is the purpose of VTP pruning?*

VTP pruning is used to control unnecessary flooding of traffic across trunk links.

21 *Which VLAN numbers are never eligible for VTP pruning? Why?*

VLAN numbers 1 and 1001–1005. VLAN 1 is reserved as the management VLAN, while VLANs 1002–1005 are reserved as the default FDDI and Token Ring function VLANs.

22 *What commands can be used to make only VLANs 300 and 400 eligible for VTP pruning?*

clear vtp pruneeligible 2-1000

set vtp pruneeligible 300,400

Remember to clear the pruning eligibility of the default VLANs first because you need to make only VLANs 300 and 400 eligible. VLAN 1 is always ineligible and 2–1000 are eligible by default.

23 *What are the steps needed to establish Token Ring switching with VLANs?*

Step 1 Define a TrBRF

Step 2 Define a TrCRF and assign it to a parent TrBRF

Step 3 Assign switch ports to the TrCRF

Answers to Chapter 5 “Do I Know This Already?” Quiz

1 *What is EtherChannel? What types of switch links can it be used with?*

EtherChannel is a method for aggregating multiple physical Ethernet ports into a single logical link. EtherChannel can be used with full-duplex Fast Ethernet or Gigabit Ethernet links.

2 *How is traffic distributed over an EtherChannel?*

Traffic is distributed according to addresses contained in frames passing through the switch—not according to port loads or equal distribution across the individual ports in a bundle. Switches use an XOR computation of source, destination, or both addresses of either MAC or IP, depending on the switch capabilities.

3 *What is PAgP used for?*

PAgP is a protocol that is used to dynamically and to automatically configure an EtherChannel over multiple physical ports.

4 *What is a bridging loop? Why is it bad?*

A bridging loop is a path through a bridged or switched network that provides connectivity in a loop. Broadcast, multicast, and unknown unicast frames introduced into the loop are propagated by each switch, causing the frames to circulate around the loop. Network bandwidth and CPU resources can be completely absorbed by the increasing amount of broadcast traffic. Bridging loops can be ended by breaking the loop connectivity.

5 *Name two types of Spanning-Tree Protocol messages used to communicate between bridges.*

Configuration BPDUs and Topology Change Notification BPDUs. Configuration BPDUs are used to inform bridges of global STP parameters and are used to form the Spanning Tree topology. Topology Change Notification BPDUs are used to inform bridges that a link state has changed potentially impacting the Spanning Tree topology.

6 *What criteria are used to select the following:*

- a. *Root Bridge*
- b. *Root Port*
- c. *Designated Port*
- d. *Redundant (or Secondary) Root Bridges*
 - a. Lowest Bridge ID (Bridge priority, MAC address)
 - b. Lowest Root Path Cost
 - c. Lowest Root Path Cost on a shared segment
 - d. Next-to-lowest Bridge ID

If a tie occurs, then this sequence of parameters are used to decide:

1. Lowest Bridge ID
2. Lowest Root Path Cost
3. Lowest Sender Bridge ID
4. Lowest Port ID

7 *What conditions cause a STP topology change? What effect does this have on STP and the network?*

A topology change occurs when a port moves to the Forwarding state or from Forwarding or Learning to the Blocking state. During a topology change, addresses are aged out in Forward Delay seconds while active stations are not aged out of the bridging table. The

STP is not recomputed; TCN BPDUs are sent throughout the network notifying other switches of the topology change. Only the port where the topology change is occurring is affected, by moving through the STP states.

- 8** *What is the single most important design decision to be made in a network running STP?*

Root Bridge placement.

- 9** *Where should the Root Bridge be located in a switched network?*

It should be located as close to the center of the network as possible. For example, in a hierarchical design, the Root Bridge should be located in the Distribution layer.

- 10** *What happens to a port that is neither a Root Port nor a Designated Port?*

That port is placed in the Blocking state so that no bridging loops form from it.

- 11** *How is the Root Path Cost calculated for a switch port?*

The Root Path cost is a cumulative value incremented as Configuration BPDUs are passed from switch to switch. For each instance of STP (one per VLAN), a switch adds the Port Cost of its local port to the current Root Path cost value as a BPDU is received.

- 12** *What is the maximum number of Root Ports that a Catalyst switch can have?*

Only one per instance of STP (one per VLAN). Each instance of STP works to find a single path from each switch back to the Root Bridge.

- 13** *What mechanism is used to set the STP timer values for all switches in a network?*

The timers are set on the Root Bridge, and the values are propagated to all other switches by including them in Configuration BPDUs.

- 14** *What parameters can be tuned to influence the selection of a port as a Root or Designated Port?*

Port Cost and Port Priority.

- 15** *What technology can be useful to decrease the amount of time STP keeps an end user's workstation in the Blocking state when it powers up?*

PortFast, which moves the port immediately into the Forwarding state.

- 16** *Where should the UplinkFast feature be used in a switched network?*

Only on switches that are leaf-nodes in the Spanning-Tree topology, such as the Access Layer.

Answers to Chapter 5 Q&A Section

- 1 *What is EtherChannel? What types of switch links can it be used with?*

EtherChannel is a method for aggregating multiple physical Ethernet ports into a single logical link. EtherChannel can be used with full-duplex Fast Ethernet or Gigabit Ethernet links.

- 2 *How does an EtherChannel distribute broadcasts and multicasts?*

Broadcasts and multicasts are sent across only one port of the bundle and are not distributed across the EtherChannel.

- 3 *How is traffic distributed over an EtherChannel?*

Traffic is distributed according to addresses contained in frames passing through the switch—not according to port loads or equal distribution across the individual ports in a bundle. Switches use an XOR computation of source, destination, or both addresses of either MAC or IP, depending on the switch capabilities.

- 4 *What CLI-based switch command could be used to configure a 4-port EtherChannel on switch ports 3/1, 3/2, 3/3, and 3/4? The switch should use PAgP to actively negotiate the EtherChannel.*

set port channel 3/1-4 mode desirable

- 5 *What is PAgP used for?*

PAgP is a protocol that is used to dynamically and automatically configure an EtherChannel over multiple physical ports.

- 6 *What happens if one port of an EtherChannel is unplugged or goes dead? What happens when that port is reconnected?*

Traffic on the disconnected port will be moved to the next available link in the EtherChannel bundle. When the port is reconnected, traffic will not automatically move back to the original port of the bundle. Rather, new traffic will be learned and applied to the restored link.

- 7 *What is a bridging loop? Why is it bad?*

A bridging loop is a path through a bridged or switched network that provides connectivity in a loop. Broadcast or multicast frames introduced into the loop are propagated by each switch causing the frames to circulate around and around the loop. Network bandwidth and CPU resources can be completely absorbed by the increasing amount of broadcast traffic. Bridging loops can be ended by breaking the loop connectivity.

8 Put the following STP port states in chronological order:

- a. Learning
- b. Forwarding
- c. Listening
- d. Blocking
- d. Blocking
- c. Listening
- a. Learning
- b. Forwarding

9 Name two types of Spanning-Tree Protocol messages used to communicate between bridges.

Configuration BPDUs and Topology Change Notification BPDUs.

10 What commands can be used to configure a CLI-based switch as the Root Bridge on VLAN 10, assuming that the other switches are using the default STP values?

The following command can be used to directly set the Bridge Priority value:

```
set spantree priority 100 10
```

Any priority value less than the default 32768 will promote the switch to Root Bridge status. A value of 100 is easy to use and remember.

Another method is to use the following command, which directly sets the switch to Root Bridge and offers an easy user interface with no values to remember or compare:

```
set spantree root 10
```

11 Using your Root Bridge answer from question 10, what commands can be used to configure another CLI-based switch as a secondary or backup Root Bridge on VLAN 10?

Using the example answer for question 10, the Root Bridge has been given a priority of 100. A secondary Root Bridge can be configured by using a Bridge Priority value that is greater than 100 and less than the default 32768. In this example answer, a secondary priority value of 200 is used, because it too is easy to remember:

```
set spantree priority 200 10
```

Another method is to directly configure the switch as a secondary Root Bridge by using this command:

```
set spantree root secondary 10
```

12 *What criteria are used to select the following:*

- a. Root Bridge*
 - b. Root Port*
 - c. Designated Port*
 - d. Redundant (or secondary) Root Bridges*
- a. Lowest Bridge ID (Bridge priority, MAC address)
 - b. Lowest Root Path Cost
 - c. Lowest Root Path Cost on a shared segment
 - d. Next-to-lowest Bridge ID

If a tie occurs, then these parameters are used to decide:

- 1. Lowest Bridge ID
- 2. Lowest Root Path Cost
- 3. Lowest Sender Bridge ID
- 4. Lowest Port ID

13 *Which of the following switches will become the Root Bridge, given the information in the table below? Which switch will become the secondary Root Bridge, in the event that the Root Bridge fails?*

Switch Name	Bridge Priority	MAC Address	Port Costs
Catalyst A	32768	00-d0-10-34-26-a0	All are 19
Catalyst B	32768	00-d0-10-34-24-a0	All are 4
Catalyst C	32767	00-d0-10-34-27-a0	All are 19
Catalyst D	32769	00-d0-10-34-24-a1	All are 19

The Root Bridge will be Catalyst C, because its Bridge Priority has the lowest value. The Bridge Priority is more significant because it is stored in the upper bits of the Bridge ID field. If Catalyst C fails in its duty as Root Bridge, then Catalyst B will take over as the secondary Root Bridge. Because Catalyst B has the default Bridge Priority (32768), along with another switch, the lowest MAC address will be the deciding factor.

- 14 *What conditions cause a STP topology change? What effect does this have on STP and the network?*

A topology change occurs when a port moves to the Forwarding state or from Forwarding or Learning to the Blocking state. During a topology change, addresses are aged out in Forward Delay seconds, while active stations are not aged out of the bridging table. The STP is not recomputed; TCN BPDUs are sent throughout the network notifying other switches of the topology change. Only the port where the topology change is occurring is affected by moving through the STP states.

- 15 *A Root Bridge has been elected in a switched network. Suppose a new switch is installed with a lower Bridge ID than the existing Root Bridge. What will happen?*

After the new switch comes up, a Root Bridge election will take place. This will occur at the next Hello time when the new switch announces itself as root. It will become the Root Bridge because it has the lowest Bridge ID, and the Spanning Tree topology will be recomputed. Where switch ports change state as a result of the election and topology change, outages will occur until the Forwarding state starts again.

- 16 *What is the single-most important design decision to be made in a network running STP?*

Root Bridge placement.

- 17 *Where should the Root Bridge be located in a switched network?*

The Root Bridge should be placed as close to the center of the network as possible. For example, in a hierarchical design, the Root Bridge should be located in the Distribution layer.

- 18 *Suppose a switch receives Configuration BPDUs on two of its ports. Both ports are assigned to the same VLAN. Each of the BPDUs announces Catalyst A as the Root Bridge. Can the switch use both of these ports as Root Ports? Why?*

No. By definition, only one Root Port is selected on each switch. The port with the lowest Root Path Cost will be chosen.

- 19 *What happens to a port that is neither a Root Port nor a Designated Port?*

That port is placed in the Blocking state so that no bridging loops form from it.

- 20 *Suppose you need to troubleshoot your Spanning Tree topology and operation. What commands and information can you use on a switch to find information about the current STP topology?*

Because it won't be obvious which switch in your network is the Root Bridge, you can begin on any switch with the **show spantree** vlan command. This command will show the current Root Bridge ID, Root Port, Designated Port, costs, timers, and port states. Consider the following example:

```

Switch> show spantree 11
VLAN 11
Spanning tree enabled

Designated Root          00-50-a2-8d-58-00
Designated Root Priority  32768
Designated Root Cost     4
Designated Root Port     1/2
Root Max Age 20 sec      Hello Time 2 sec      Forward Delay 15 sec

Bridge ID MAC ADDR       00-50-bd-19-6c-00
Bridge ID Priority        32768
Bridge Max Age 20 sec    Hello Time 2 sec      Forward Delay 15 sec

Port,Vlan Vlan  Port-State      Cost  Priority  Fast-Start  Group-method
-----
 1/1     11  not-connected      4     32     disabled
 1/2     11  forwarding         4     32     disabled
Switch>

```

From this information, you can discover the MAC address of the Root Bridge. Sometimes it can be helpful to “walk” the Spanning Tree topology to find the Root Bridge. To do this, you would need to Telnet to the next upstream switch to get one hop closer to the Root. Find the Root Port, listed above. Then use the **show cdp neighbor [module/port] detail** command to find the IP address of the upstream neighbor. Using this methodology, you can trace every hop to the Root Bridge.

21 *How is the Root Path Cost calculated for a switch port?*

The Root Path cost is a cumulative value that is incremented as Configuration BPDUs are passed from switch to switch. A switch adds the Port Cost of its local port to the current Root Path Cost value as a BPDU is received.

22 *What conditions can cause ports on the Root Bridge of a network to move into the Blocking state? (Assume that all switch connections are to other switches. No crossover cables are used to connect two ports together on the same switch.)*

None. Every active port on the Root Bridge becomes a Root Port, since the Root Path Cost is zero. By this definition, the ports can never be in the Blocking state.

23 *What is the maximum number of Root Ports that a Catalyst switch can have?*

Only one. STP works to find a single path from each switch back to the Root Bridge.

24 *What mechanism is used to set the STP timer values for all switches in a network?*

The timers are set on the Root Bridge, and the values are propagated to all other switches by including them in Configuration BPDUs.

- 25 *What parameters can be tuned to influence the selection of a port as a Root or Designated Port?*

Port Cost and Port Priority.

- 26 *What CLI-based command can be used to enable fast STP convergence for a single workstation on switch port 3/7?*

set spantree portfast 3/7 enable

- 27 *What technology can be useful to decrease the amount of time STP keeps an end user’s workstation in the Blocking state when it powers up?*

PortFast, which moves the port immediately into the Forwarding state.

- 28 *What happens if the STP Hello Time is decreased to one second in an effort to speed up STP convergence? What happens if the Hello Time is increased to ten seconds?*

Setting the Hello Timer to one second doubles the amount of Configuration BPDUs that are sent by a switch, as compared to the default 2 second timer. While this does share BPDU information more often, it really doesn’t help the long convergence delay when a port comes up. The significant delays come from the Forward Delay timer, which is used to move a port through the Listening and Learning states. By default, this process takes 30 seconds and is unaffected by the Hello Timer.

- 29 *Where should the UplinkFast feature be used in a switched network?*

Only on switches that are leaf-nodes in the Spanning Tree topology, such as the Access layer.

- 30 *What CLI-based switch command can be used to safely adjust the STP timers on the Root Bridge in VLAN 7? Assume that the network consists of Catalyst A, B, and C, all connected to each other in a triangle fashion.*

Because the three switches form a triangle loop, one link will eventually be placed in the Blocking state. Therefore, the maximum distance across the network is 3 switch hops. This value can be used to define the network diameter to safely adjust the STP timers for faster convergence:

```
set spantree root dia 3
```

Answers to Chapter 6 “Do I Know This Already?” Quiz

- 1 *What is the basic unit of ATM data? What is its basic format (header, payload, and so forth)?*

The basic ATM data unit is the cell. An ATM cell consists of a 5-byte header and a 48-byte payload.

2 *What is an ATM edge device? What Cisco devices can be used?*

An ATM edge device interfaces native ATM to other media. For example, Cisco Catalyst switches (5000 and 6000) can be used to bridge between LAN ports and an ATM LANE module. Also, Cisco routers (4500/4700, 7500, for example) can bridge between any LAN or WAN media and ATM LANE. Some video CODECs can also function as edge devices, converting raw video to ATM cells.

3 *What type of addressing is used to identify ATM devices?*

The Network Service Access Point (NSAP) address uniquely identifies all ATM devices. The NSAP is a 20-byte value.

4 *What information is carried within each ATM data unit to specify how to get from the source to the destination?*

The VPI/VCI pair references the virtual connection between source and destination. The path identified by the VPI/VCI pair is used to relay cells on ATM switches.

5 *What is LANE used for?*

LAN Emulation (LANE) is an ATM standard that enables an ATM network to function as a LAN backbone. A LAN is emulated using LANE components to provide packet transport, broadcast and multicast support, and address management (MAC addresses to ATM addresses). LANE can also be used to trunk multiple LANs across an ATM network.

6 *What are the functional components of LANE?*

LANE is built with the LECS, LES, BUS, and LEC components.

7 *What is the difference between a VLAN and an ELAN?*

A VLAN (virtual LAN) is a logical network segment (broadcast domain) used on LAN switches. An ELAN (emulated LAN) is a logical network segment (also a broadcast domain) created between ATM LANE devices. VLANs and ELANs are separate until a Catalyst LANE module bridges them.

8 *When is an LE_ARP used?*

LE_ARP is used to resolve between a MAC address and an NSAP address.

9 *For each of the LANE components, how many are necessary for LANE operation?*

LECS: one; LES: one per ELAN; BUS: one per ELAN; LEC: one at each device where ELAN connectivity is needed.

10 *Which LANE component provides connectivity between a VLAN and an ELAN on a switch?*

The LEC performs this function as it is configured with both VLAN and ELAN identifiers.

11 *What NSAP addresses must be configured into the LECS database?*

The LES NSAP address for each ELAN must be configured into the LECS. LANE clients will connect to the LECS and ask for their respective LES address.

12 *What is SSRP? Which LANE components can be configured for SSRP?*

SSRP (Simple Server Redundancy Protocol) is a method to provide multiple or redundant LANE components in a network. Because the LANE 1.0 specification only allows for one instance of each component, SSRP was developed to allow redundancy. Redundant LECS and LES/BUS components can be configured. However, redundant LECs in an ELAN are not allowed.

13 *What Catalyst switch command can be used to view the current status of each LANE component (LECS, LES, BUS, and LEC)?*

LECS: **show lane database.**

LES: **show lane server.**

BUS: **show lane bus.**

LEC: **show lane client.**

Answers to Chapter 6 Q&A Section

1 *What is the basic unit of ATM data? What is its basic format (header, payload, etc.)?*

The basic ATM data unit is the cell. An ATM cell consists of a 5-byte header and a 48-byte payload.

2 *What process allows an IP packet to be transported within ATM cells?*

Segmentation and Reassembly (SAR) provided by the ATM Adaptation Layer (AAL).

3 *What is an ATM edge device? What Cisco devices can be used?*

An ATM edge device interfaces native ATM to another media. For example, Cisco Catalyst switches (5000 and 6000) can be used to bridge between LAN ports and an ATM LANE module. Also Cisco routers (4500/4700, 7500, and so on) can bridge between any LAN or WAN media and ATM LANE. Some video CODECs can also function as edge devices, converting raw video to ATM cells.

4 *What types of VCs can be built with ATM?*

Permanent virtual circuits (PVCs) and switched virtual circuits (SVCs); of those, both point-to-point and point-to-multipoint circuits can be configured.

5 *List the hierarchy of ATM VCs.*

Virtual Channel (VC), then Virtual Path (VP), then Transmission Path.

6 *What type of addressing is used to identify ATM devices?*

The network service access point (NSAP) address uniquely identifies all ATM devices. The NSAP is a 20-byte value.

7 *What are the three basic components of a NSAP address?*

A 20-byte NSAP address is made up of a Prefix (13 bytes), an ESI (6 bytes), and a Selector (1 byte).

8 *What information is carried within each ATM data unit to specify how to get from the source to the destination?*

The VPI/VCI pair references the virtual connection between source and destination. The path identified by the VPI/VCI pair is used to relay cells on ATM switches.

9 *Name two inherent ATM protocols used by ATM switches to communicate with other ATM devices.*

ILMI (switch to other devices) and PNNI (switch to switch).

10 *What is LANE used for?*

LAN Emulation (LANE) is an ATM standard that enables an ATM network to function as a LAN backbone. A LAN is emulated using LANE components to provide packet transport, broadcast and multicast support, and address management (MAC addresses to ATM addresses). LANE can also be used to trunk multiple LANs across an ATM network.

11 *Where should the LECS NSAP address be configured in a LANE network? Where should the LES NSAP address be configured?*

The LECS address should be configured in the ATM switch(es) so that the address can be provided to other devices via ILMI. The LES address should be configured into the LECS database.

12 *What are the functional components of LANE?*

LANE is built with the LECS, LES, BUS, and LEC components.

13 *What is the effect of a failed LECS on a LANE network? A failed LES? A failed BUS? A failed LEC?*

A failed LECS will not provide LES addresses to any new LANE clients that initialize. Existing LECs will continue to operate but new LECs will not be able to join an ELAN.

A failed LES cannot provide MAC to NSAP address resolution for its ELAN. Therefore, LECs will not be able to get a reply to LE_ARP requests.

A failed BUS cannot provide broadcast or multicast functionality to its ELAN. This situation has a major impact on ELAN operation, as broadcasts submitted to the BUS go unprocessed.

Lastly, a failed LEC cannot provide the bridging function between a VLAN and an ELAN. Therefore, the VLAN portion will continue to operate but will be segmented from the ELAN. Other LECs in the ELAN can communicate, but the failed LEC will not be a part of the ELAN.

- 14 *When a LEC initializes, what LANE component does it contact first? Why?*

A LEC attempts to contact the ATM switch first to learn the LECS NSAP address via ILMI. Then the LEC tries to contact the LECS directly.

- 15 *What is the difference between a VLAN and an ELAN?*

A VLAN (virtual LAN) is a logical network segment (broadcast domain) used on LAN switches. An ELAN (emulated LAN) is a logical network segment (also a broadcast domain) created between ATM LANE devices. VLANs and ELANs are separate until a Catalyst LANE module with a LEC interconnects them. Then the ELAN functions as a means to extend a VLAN across the ATM network.

- 16 *When is an LE_ARP used?*

LE_ARP is used to resolve between a MAC address and an NSAP address.

- 17 *The following NSAP addresses were obtained from **show lane default**. Match the LANE components to their NSAP addresses:*

```
47.00001606288000300000000F.0064A28D5EA0.**
47.00001606288000300000000F.0064A28D5EA1.**
47.00001606288000300000000F.0064A28D5EA2.**
47.00001606288000300000000F.0064A28D5EA3.00

47.00001606288000300000000F.0064A28D5EA2.** : BUS
47.00001606288000300000000F.0064A28D5EA3.00 : LECS
47.00001606288000300000000F.0064A28D5EA1.** : LES
47.00001606288000300000000F.0064A28D5EA0.** : LEC
```

- 18 *How many of each LANE component are necessary for LANE operation?*

LECS: one; LES: one; BUS: one; LEC: one at each device where ELAN connectivity is needed.

- 19 *A network consists of two Catalyst switches connected by an ATM switch. Four VLANs exist on each Catalyst and are trunked between all Catalyst switches. How many LECS components are present? How many LES components? How many BUS components? How many LECs?*

Only 1 LECS is used for all ELANs. Four LESs are needed—one for each ELAN. Four BUSs are needed—one for each ELAN. Eight LECs are needed—one for each ELAN present on each switch (2 switches x 4 ELANs).

- 20 Which LANE component performs bridging between a VLAN and an ELAN?

The LEC performs this function, as it is configured with both VLAN and ELAN identifiers.

- 21 Given an NSAP address of 47.00001606288000300000000F.0064A28D5EA1.0A, what LANE component does this represent? On which subinterface of ATM 0 is this assigned?

This address represents a LES because the rightmost digit of the ESI portion (MAC address) is a “1”. The LES has been configured on interface ATM 0.10. The subinterface number is given in the Selector portion of the NSAP address (0A).

- 22 What NSAP addresses must be configured into the LECS database?

The LES NSAP address for each ELAN must be configured into the LECS. LANE clients will connect to the LECS and ask for their respective LES address.

- 23 A LANE database (mylane-db) has already been created on a Catalyst switch. The switch contains a single ATM interface (ATM 0) and subinterfaces ATM 0.1 and ATM 0.2. What commands can be used to enable the LECS with this database on the ATM link?

```
interface atm 0
 lane config database mylane-db
 lane config auto-config-atm-address
```

Remember that the LECS must be configured on the major ATM interface, not a subinterface.

- 24 If a switch supports four VLANs, how many LECs will be needed? How many ATM subinterfaces will be required?

Four LECs will be needed, one for each VLAN; four ATM subinterfaces will be needed, one for each LEC.

- 25 What is SSRP? Which LANE components can be configured for SSRP?

SSRP (Simple Server Redundancy Protocol) is a method to provide multiple or redundant LANE components in a network. Because the LANE 1.0 specification only allows for one instance of each component, SSRP was developed to allow redundancy. Redundant LECS and LES/BUS components can be configured. However, redundant LECs in an ELAN are not allowed.

- 26 How many ATM subinterfaces are required to place both an LES and a BUS on a Catalyst switch?

Only one subinterface is needed. The LES and BUS are co-located on the switch and share a single subinterface. (This also means that the Selector byte in the NSAP addresses will be the same between LES and BUS, but the ESI portion will be different.)

- 27 *What Catalyst switch command can be used to view the current status of each LANE component (LECS, LES, BUS, and LEC)?*

BUS: **show lane bus**

LEC: **show lane client**

LECS: **show lane database**

LES: **show lane server**

- 28 *Is it possible to place all LANE components (LECS, LES, BUS, and LEC) on a single Catalyst switch? If only one VLAN is required, what is the minimum number of interfaces required?*

Yes it is possible, though the LEC will need another LEC to communicate with in the ELAN. With only one VLAN (ELAN), one major ATM interface is needed for the LECS, and only one ATM subinterface is needed for the remaining components. The LES, BUS, and LEC can all be configured on one subinterface.

Answers to Chapter 7 “Do I Know This Already?” Quiz

- 1 *Where can a router be placed in relation to switches for interVLAN routing?*

External to the switches or internal (integrated) to a switch.

- 2 *What types of links can be used to interconnect switches and an external router? How many VLANs can be carried on each?*

Links can be used with one VLAN per physical link, using any supported media. Trunk links can also be used to carry multiple VLANs over a single link, using such media as Fast Ethernet, Gigabit Ethernet, ATM LANE, and FDDI.

- 3 *What trunking methods can a router support?*

IEEE 802.1Q (Ethernet), ISL (Ethernet), LANE (ATM), and 802.10 (FDDI).

- 4 *What is the difference between interVLAN routing and multilayer switching (MLS)?*

InterVLAN routing uses a route processor to forward data between VLANs. It also requires that each packet crossing a VLAN boundary be processed by a routing decision on the router. MLS also uses a route processor to forward data between VLANs, but the routing decision is only required for the first packet exchange between two nodes. From that point on, data is forwarded by the switching engine and does not require further routing.

- 5 *What Catalyst commands can be used to locate and then connect to an internal route processor?*

First, the **show module** command should be used to get a list of installed modules in the switch. Then, the **session slot-number** command can be used to start a Telnet session to the route processor module.

- 6 *What should be configured on a route processor to dynamically determine routing paths to remote networks?*

Routing protocols, such as RIP, IGRP, EIGRP, OSPF, BGP, and so on.

- 7 *If a router is used to route between VLANs, what additional information is needed so that traffic will actually be routed?*

A default gateway is required on the end stations so that the stations can forward packets that are destined to a different VLAN or network to the router.

- 8 *Suppose a router connects four VLANs of a switched network, providing interVLAN routing. If the router is then configured for bridging to support nonroutable protocols, how would the network be affected?*

The whole idea of using VLANs in the network is to achieve segmentation between networks and workgroups. Each VLAN is an independent broadcast domain. Enabling bridging between VLANs does permit nonroutable protocols to pass between VLANs. However, the nature of bridging means that broadcasts must be forwarded to all bridge ports. The end result is that the entire four-VLAN network becomes one broadcast domain.

Answers to Chapter 7 Q&A Section

- 1 *Which of the following modules performs interVLAN routing?*

- a. Catalyst 5000 Supervisor III
- b. Catalyst 5000 RSFC
- c. Catalyst 5000 LANE Module
- d. Catalyst 5000 Gigabit EtherChannel Module
- e. Catalyst 5000 RSM

Both (b) RSFC and (e) RSM are route processors for the Catalyst 5000 family.

- 2 *Where can a router be placed in relation to switches for InterVLAN routing?*

External to the switches or internal (integrated) to a switch.

3 *How many links are needed to connect a router to four VLANs on a switch?*

Only one link is required if trunking is used. If trunking is not used, four links (one per VLAN) are required.

4 *What types of link can be used to interconnect switches and an external router? How many VLANs can be carried on each?*

Links can be used with one VLAN per physical link, using any supported media. Trunk links can also be used to carry multiple VLANs over a single link, using such media as Fast Ethernet, Gigabit Ethernet, ATM LANE, and FDDI.

5 *Which Catalyst route processor module uses four internal Gigabit Ethernet links to interface with the switch backplane?*

The MultiLayer Switch Module (MSM) uses four Gigabit Ethernet links internally to interface with the Catalyst 6000/6500 switching fabric.

6 *What trunking methods can a router support?*

IEEE 802.1Q (Ethernet), ISL (Ethernet), LANE (ATM), and 802.10 (FDDI).

7 *Which is better: one link per VLAN or a single trunk link supporting all VLANs? Why?*

One link per VLAN offers VLAN transport using native Ethernet technology. Therefore, no extra trunking encapsulation processing is needed, subsequently increasing efficiency. Because many links would be needed for many VLANs, this method is not scalable. Trunking is much more scalable by supporting a large number of VLANs over a single link. Although requiring extra processing of the trunking encapsulation method, (Inter-Switch Link [ISL]) which is performed efficiently in hardware ASICs.

8 *What is the difference between InterVLAN routing and MultiLayer Switching (MLS)?*

InterVLAN routing uses a route processor to forward data between VLANs. It also requires that each packet crossing a VLAN boundary be processed by a routing decision on the router. MLS also uses a route processor to forward data between VLANs, but the routing decision is only required for the first packet exchange between two nodes. From that point on, data is forwarded by the switching engine and does not require further routing.

9 *What commands are needed to assign IP address 10.1.2.247 255.255.0.0 to an RSM interface for VLAN 72?*

```
interface vlan 72  
ip address 10.1.2.247 255.255.0.0
```

- 10 *What Catalyst commands can be used to locate and then connect to an internal route processor?*

First, the **show module** command should be used to get a list of installed modules in the switch. Then, the **session slot-number** command can be used to start a telnet session to the route processor module.

- 11 *What commands can be used to configure IP address 10.10.10.1 255.0.0.0 on an external router’s VLAN 101 interface, assuming FastEthernet 0/0 is being used as a trunk link?*

```
interface fastethernet 0/0.1  
encapsulation isl 101  
ip address 10.10.10.1 255.0.0.0
```

- 12 *What should be configured on a route processor to dynamically determine routing paths to remote networks?*

Routing protocols, such as RIP, IGRP, EIGRP, OSPF, and BGP.

- 13 *A router with an ATM LANE module is used to route traffic between six VLANs located on each of 10 Catalyst switches. Each switch has a LANE module configured for the matching VLANs and ELANs. All switches are connected by an ATM switch. What LANE components are needed on the router to support interVLAN routing? How many of each component should be configured?*

At a minimum, the router will need one LEC configured per ELAN (VLAN), for a total of six LECs. Somewhere in the network, one LECS (supporting all ELANs) and six LES/BUS pairs (one per ELAN) must be present.

- 14 *If a router is used to route between VLANs, what additional information is needed so that traffic will actually be routed?*

A default gateway is required so that end stations can forward packets, destined to a different VLAN or network, to the router.

- 15 *Suppose a router connects four VLANs of a switched network, providing interVLAN routing. If the router is then configured for bridging to support non-routable protocols, how would the network be affected?*

The whole idea of using VLANs in the network is to achieve segmentation between networks and workgroups. Each VLAN is an independent broadcast domain. Enabling bridging between VLANs does permit non-routable protocols to pass between VLANs. However, the nature of bridging means that broadcasts must be forwarded to all bridge ports. The end result is that the entire four-VLAN network becomes one broadcast domain.

- 16 A switch has its sc0 interface configured with an IP address of 10.10.1.17 255.255.0.0 on VLAN 10. A router is connected to the switched network and has interfaces at 10.10.1.1 and 10.11.1.1. What must be configured to allow someone connected to the switch console to ping host 10.11.1.100 on VLAN 11?

On the switch, a default gateway pointing to the router must be configured with the **set ip route default 10.10.1.1** command.

Answers to Chapter 8 “Do I Know This Already?” Quiz

- 1 What devices make up the basis for Layer 3 switching as it relates in a Cisco environment?

Catalyst switches

- 2 What device is the definition of a Multilayer Switch Engine (MLS-SE)?

The Multilayer Switch Engine is a Supervisor III card in a Catalyst switch with a Netflow Feature Card (NFFC) enabled on it. On a Catalyst 6000, the PFC/MSFC combination can also perform MLS.

- 3 What devices can be used as a Multilayer Switch Route Processor (MLS-RP)?

A Route Switch Module (RSM) and any Cisco router that supports MLS in software (Typically, a 75xx, 72xx, 45xx, 47xx, or 85xx series router). Please note that the Catalyst 6000 can be an MLS-RP through the use of the internal MSFC.

- 4 What is the command for enabling MLS on an RP?

In global configuration mode, enter the following command:

```
Router(config)#mls rp ip
```

- 5 What two things are required to make an interface on an RP MLS-enabled?

The interface must be identified with a VLAN ID and also the VTP Domain that it belongs to. Note that on a Catalyst 6000, MLS is on by default.

- 6 What command is used to verify the MLS configuration for an MLS-RP ?

```
Router#show mls rp
```

- 7 What are the three types of flow masks modes supported on a MLS-SE?

Destination-IP, Source-Destination-IP, and Full-Flow.

- 8 What is the command to add an input access list to a MLS flow?

```
Router(config)#mls rp ip input-acl
```

- 9 *When using an external RP to a switch, is this configured automatically or manually?*

Only internal RSMs are automatically configured as RPs. Any external router must be manually configured on the switch.

- 10 *What is the command to enable multilayer switching for a Catalyst switch?*

Switch(enable)#**set mls enable**

Answers to Chapter 8 Q&A Section

- 1 *What devices are the basis for Layer 3 switching as it relates in a Cisco environment?*

Catalyst switches.

- 2 *What device is the definition of a Multilayer Switch Engine (MLS-SE)?*

The Multilayer Switch Engine is a Supervisor III card in a Catalyst switch with a Netflow Feature Card (NFFC) enabled on it.

- 3 *What devices can be used as a Multilayer Switch Route Processor (MLS-RP)?*

A Route Switch Module (RSM) or any Cisco router that supports MLS in software (Typically, a 75xx, 72xx, 45xx, 47xx, or 85xx series router).

- 4 *What is the command for enabling MLS on an RP?*

In global configuration mode, enter the following command:

Router(config)#**mls rp ip**

- 5 *What two things are required to make an interface on an RP MLS-enabled?*

The interface must be identified with a VLAN ID and also the VTP Domain that it belongs to.

- 6 *What command is used to verify the MLS configuration for an MLS-RP ?*

Router#**show mls rp**

- 7 *What are the three types of flow masks modes supported on a MLS-SE?*

Destination-IP, Source-Destination-IP, and IP-Flow.

- 8 *What is the command to add an input access list to a MLS flow?*

Router(config)#**mls rp ip input-acl**

- 9 *When using an external RP to a switch, is this configured automatically or manually?*

Only RSMs are automatically configured as RPs, any external router must be manually configured on the switch.

- 10 *What is the command to enable Multilayer Switching for a Catalyst switch?*
Switch(enable)#**set mls enable**
- 11 *Assuming that MLS is running, what effect does the command **clear ip route** do on an MLS-RP?*
This command clears the MLS cache.
- 12 *What three components are required in a Cisco implementation of MLS?*
The MLS-SE, the MLS-RP, and MLSP.
- 13 *Define a Destination-IP flow mask.*
The MLS-SE maintains one MLS entry for each destination IP address. All flows to a given destination IP address use this MLS entry.
- 14 *What is the command to display MLS entries in the cache?*
Switch (enable)#**show mls entry**

Answers to Chapter 9 “Do I Know This Already?” Quiz

- 1 *What is the name of the protocol that allows a set of routers that are working together to form one virtual-router?*
Hot Standby Router Protocol or HSRP.
- 2 *What is the minimum number of routers needed to perform HSRP?*
The minimum number of routers needed is at least two. One functions as an active router and one as a standby.
- 3 *In a properly functioning virtual router, what happens when the active router fails?*
In a properly functioning HSRP environment, packets will still be routed in the event of a failed router.
- 4 *How many standby groups can exist on any one LAN?*
In any one LAN, up to 255 standby groups can exist.
- 5 *Name the six states that an HSRP configured router can be in.*
Initial, Learn, Listen, Speak, Standby, and Active.
- 6 *When configuring HSRP on a particular router interface, if the standby group is not explicitly configured, what standby group does the interface fall into by default?*
Standby group 0 is the default setting when configuring an interface for HSRP. You may, however, override this by configuring a setting of your own choosing.

7 *What command is used to display the HSRP virtual router IP and MAC address?*

The Cisco command **show standby** will display the HSRP virtual router IP and MAC address.

8 *Which router in an HSRP group becomes the forwarding router and how is it determined?*

The router that becomes the forwarding router in an HSRP group is the one with the highest priority. The priority is determined by what has been configured. The default value is 100, but the priority can be any number between 0 and 255. If the priorities are equal, then the highest IP address takes priority.

9 *In the command **standby 35 priority 90**, what does the “35” stand for?*

The “35” indicates the standby group number.

10 *An HSRP router exchanges Hello messages with other HSRP routers. What is contained in the hello message?*

The HSRP Hello message contains the hellotime and the holdtime values, in addition to the priority, group number, password, and virtual router.

11 *What does the term tracking imply in an HSRP environment?*

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.

12 *What command would allow you to debug HSRP?*

The command **debug standby** would allow you to debug HSRP activity.

13 *What does the **preempt** command do in the HSRP environment?*

The **preempt** command is used to ensure that after an active router has failed, it will resume its active router role once the router has recovered. That is assuming that the router has higher priority AND preempt is configured. Without this feature, the “new” active router will remain the active router indefinitely.

14 *What command or commands enable the preempt role in an HSRP-enabled network?*

The command **standby 35 preempt** will enable the preempt role on an interface, where “35” in this example is the group number.

Answers to Chapter 9 Q&A Section

- 1 *What is the name of the protocol that allows a set of routers that are working together to form one virtual router?*
Hot Standby router Protocol or HSRP.
- 2 *What problem makes HSRP necessary?*
The fact that there isn't a dynamic protocol to discover new default gateways for hosts in the event of failure.
- 3 *What is the minimum number of routers needed to perform HSRP?*
The minimum number of routers needed is at least two. One functions as an active router and one as a standby.
- 4 *What is the RFC that pertains to HSRP?*
RFC 2281 covers Hot Standby router Protocol.
- 5 *In a properly functioning virtual router, what happens when one of the routers fails?*
In a properly functioning HSRP environment, packets will still be routed in the event of a failed router.
- 6 *How many standby groups can exist on any one LAN?*
In any one LAN, up to 255 standby groups can exist.
- 7 *What constitutes an HSRP group?*
An HSRP group consists of an active router, a standby router, and the virtual router.
- 8 *What is the role of the active router?*
The active router is responsible for forwarding packets sent to the virtual router.
- 9 *Name the six states that an HSRP configured router can be in.*
Initial, Learn, Listen, Speak, Standby, and Active.
- 10 *When configuring HSRP on a particular router interface, if the standby group is not explicitly configured, what standby group does the interface fall into by default?*
Standby group 0 is the default setting when configuring an interface for HSRP. You may, however, override this by configuring a setting of your own choosing.
- 11 *Assume you are using five VLANs within your network and want to implement HSRP. How many HSRP groups would you need to create?*
Five; when using multiple VLANs in an HSRP implementation, a separate HSRP group is created for each VLAN.

12 *What command is used to display the HSRP virtual router IP and MAC address?*

The Cisco command **show standby** will display the HSRP virtual router IP and MAC address.

13 *Which router in an HSRP group becomes the forwarding router and how is it determined?*

The router that becomes the forwarding router in an HSRP group is the one with the highest priority. The priority is determined by what has been configured. The default value is 100, but the priority can be any number between 0 and 255.

14 *In the command **standby 35 priority 90**, what does the “35” stand for?*

The “35” indicates the standby group number.

15 *An HSRP router exchanges Hello messages with other HSRP routers. What is contained in the Hello message?*

The HSRP Hello message contains the hellotime and the holdtime values, in addition to the priority, group number, password, and virtual router.

16 *What does the term tracking imply in an HSRP environment?*

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.

17 *What command would allow you to debug HSRP?*

The command **debug standby** would allow you to debug HSRP activity.

18 *What does the **preempt** command do in the HSRP environment?*

The **preempt** command is used to ensure that after an active router has failed, that it will resume its active router role once the router has recovered. Without this feature the “new” active router will remain the active router indefinitely.

19 *What command or commands enable the preempt role in an HSRP-enabled network?*

The command **standby 35 preempt** will enable the preempt role on an interface, where “35” in this example is the group number.

20 *What command shows the status of an HSRP router?*

The command **show standby** is used to display status.

Answers to Chapter 10 “Do I Know This Already?” Quiz

- 1 *Name the three types of traffic available in today’s multimedia environment?*

Unicast traffic, broadcast traffic, and multicast traffic.

- 2 *What Layer 4 protocol is used to carry multicast traffic?*

The transport layer protocol UDP is used to carry multicast traffic. UDP is a simpler, more efficient protocol because there is no flow control, reliability, or error recovery added to IP.

- 3 *What Class of IP address is used in a multicast environment?*

IP multicast is Class D.

- 4 *Describe the makeup of the Class D multicast address by octet or bits.*

The first octet or 8 bits starts with the binary 1110 or decimal 224 and carries through to decimal 239. The last 28 bits or 3 octets make up the multicast group address or id. These series of bits are unstructured, unlike traditional Class A, B, or C addresses.

- 5 *What is the name of the protocol used to report their multicast group membership with neighboring multicast routers?*

Internet Group Management Protocol (IGMP).

- 6 *What is the special name assigned to the one multicast router that performs Host Membership Queries to determine which groups have members?*

The router is called a querier.

- 7 *What does a host send to the multicast group address to join a group?*

The host can send a Host Membership Report to join a multicast group.

- 8 *Which type of routing involves transmitting packets from one source to one source?*

Unicast routing.

- 9 *Define a distribution tree.*

A distribution tree is constructed by routers and specifies a unique forwarding path between the subnet of the source and the subnets containing members of the multicast group.

- 10 *Name the two types of distribution trees.*

Source specific trees and shared, or center-specific, trees.

- 11 *Name the three types of dense mode routing protocols.*

Distance Vector Multicast Routing Protocol (or DVMRP), Multicast Open Shortest Path First (or MOSPF), and Protocol Independent Multicast Dense Mode (or PIM DM).

12 *Name the two types of sparse mode routing protocols.*

Core-Based Trees (CBT) and Protocol Independent Multicast Sparse Mode (PIM SM).

13 *Which multicast routing protocol is widely used on the MBONE?*

DVMRP.

Answers to Chapter 10 Q&A Section

1 *Name the three types of traffic available in today’s multimedia environment?*

Unicast traffic, broadcast traffic, and multicast traffic.

2 *What Layer 4 protocol is used to carry multicast traffic?*

The transport layer protocol UDP is used to carry multicast traffic. UDP is a simpler, more efficient protocol because there is no flow control, reliability, or error recovery added to IP.

3 *What Class of IP address is used in a multicast environment?*

IP multicast is Class D.

4 *Describe the makeup of the Class D multicast address by octet or bits.*

The first octet or 8 bits starts with binary 1110 or decimal 224 and carries through to decimal 239. The last 28 bits or 3 octets make up the multicast group address or id. These series of bits are unstructured, unlike traditional Class A, B, or C addresses.

5 *What is the name of the protocol used to report their multicast group membership with neighboring multicast routers?*

Internet Group Management Protocol (IGMP).

6 *What is the special name assigned to the one multicast router that performs Host Membership Queries to determine which groups have members?*

The router is called a querier.

7 *What does a host send to the all-router group address of 224.1.1.1 to join a group?*

The host can send a Host Membership Query to join a multicast group.

8 *Which type of routing involves transmitting packets from one source to one source?*

Unicast routing.

9 *Define a distribution tree.*

A distribution tree is constructed by routers and specifies a unique forwarding path between the subnet of the source and the subnets containing members of the multicast group.

10 *Name the two types of distribution trees.*

Source specific trees and shared, or center-specific, trees.

11 *Name the three types of dense mode routing protocols.*

Distance Vector Multicast Routing Protocol (DVMRP), Multicast Open Shortest Path First (MOSPF), and Protocol Independent Multicast Dense Mode (PIM DM).

12 *Name the two types of sparse mode routing protocols.*

Core-Based Trees (CBT) and Protocol Independent Multicast Sparse Mode (PIM SM).

13 *Which multicast routing protocol is widely used on the MBONE?*

DVMRP.

14 *Name three characteristics of IP multicasting?*

All the characteristics of IP multicasting are:

- Facilitates transmission of an IP datagram to a multicast group comprised of zero or more hosts identified by a single IP destination address.
- Delivers a multicast datagram to all members of the multicast group with the same “best-effort” reliability as regular unicast IP datagrams.
- Supports dynamic membership of a multicast group.
- Supports all multicast groups regardless of the location or number of members.
- Supports the membership of a single host in one or more multicast groups.
- Upholds multiple data streams at the application level for a single group address.
- Supports a single group address for multiple applications on a host.

15 *Certain traditional routing protocols use multicasts to carry routing information. Name one routing protocol and the multicast address it uses.*

The traditional routing protocols of OSPF and RIP2 use multicast addressing. Their addresses are 224.0.0.5 and 224.0.0.6 for OSPF and 224.0.0.9 for RIP2.

16 *What is the name of the Cisco specific protocol that is used with routers and switches to configure the multicast forwarding table to represent group membership?*

Cisco Group Membership Protocol (CGMP).

- 17 *What is the algorithm used in a source specific distribution tree?*
Reverse Path Forwarding (RPF).
- 18 *What is used to manage the scope of multicast delivery?*
Time-to-live (TTL).
- 19 *What two characteristics describe when Sparse Mode PIM is most useful?*
When traffic is intermittent and there are few receivers in a group.
- 20 *MOSPF is best suited to which type of environment?*
When there are only a few pairs (source, group) active at any time.

Answers to Chapter 11 “Do I Know This Already?” Quiz

- 1 *Which Internet Request for Comment (RFC) deals with multicasts?*
RFC 1112 is titled Host Extensions for IP Multicasting and was the original specification. RFC 2236, titled “Internet Group Management Protocol, Version 2,” is the most recent.
- 2 *What is the name of the industry standard protocol that deals with multicast groups? The Cisco proprietary protocol?*
IGMP is the industry standard, and CGMP is the Cisco proprietary protocol.
- 3 *What command enables multicast routing on a Cisco router?*
At the enable prompt type **ip multicast-routing**.
- 4 *What command is issued to enable PIM in Sparse Mode?*
At the enable prompt type **ip pim sparse-mode**.
- 5 *What is the default type of IGMP used in a Cisco router?*
IGMP version 2 is the default in all routers running IOS version 11.3(2)T or greater.
- 6 *What command would you use to display all multicast packets received and transmitted on a router?*
You would need to activate multicast debugging. The command syntax is `debug ip mpacket`.
- 7 *What is the status of CGMP in default mode on a Cisco router?*
CGMP is disabled in default mode on a Cisco router.

8 *How is CGMP enabled on a Cisco router? On a Catalyst switch?*

CGMP is enabled on a Cisco router by typing **ip cgmp**. On a Catalyst switch, you type **set cgmp enable**.

9 *What is the purpose of CGMP leave?*

CGMP leave is the capability of the switch to detect IGMP Version 2 leave messages. If there are no CGMP join messages within a certain timeframe, the port is pruned from the multicast group. This condition makes optimal use of bandwidth.

Answers to Chapter 11 Q&A Section

1 *Which Internet Request for Comment (RFC) deals with multicasts?*

RFC 1112 is titled Host Extensions for IP Multicasting and was the original specification. RFC 2236, titled “Internet Group Management Protocol, Version 2,” is the most recent.

2 *What is the name of the industry standard protocol that deals with multicast groups? The Cisco proprietary protocol?*

IGMP is the industry standard, and CGMP is the Cisco proprietary protocol.

3 *What command enables multicast routing on a Cisco router?*

At the enable prompt type **ip multicast-routing**.

4 *What command is issued to enable PIM in Sparse Mode?*

At the enable prompt type **ip pim sparse-mode**.

5 *What is the default type of IGMP used in a Cisco router?*

IGMP version 2 is the default in all routers running IOS version 11.3(2)T or greater.

6 *What command would you use to display all multicast packets received and transmitted on a router?*

You would need to activate multicast debugging. The command syntax is **debug ip mpacket**.

7 *What is the status of CGMP in default mode on a Cisco router?*

CGMP is disabled in default mode on a Cisco router.

8 *How is CGMP enabled on a Cisco router? On a Catalyst switch?*

CGMP is enabled on a Cisco router by typing **ip cgmp**. On a Catalyst switch, you type **set cgmp enable**.

9 *What is the purpose of CGMP leave?*

CGMP leave is the capability of the switch to detect IGMP Version 2 leave messages. If there are no CGMP join messages within a certain timeframe, the port is pruned from the multicast group. This condition makes optimal use of bandwidth.

10 *Name one multicast routing protocol that is used by Cisco routers for router to router communication?*

PIM is the multicast routing protocol that is most used by Cisco routers for router-to-router communication. You could also use MBGP or DVMRP.

11 *What is the MBONE? What multicast routing protocol is used throughout the MBONE?*

The MBONE is the Internet Multicast backbone. DVMRP is used on the MBONE. Cisco routers perform PIM to DVMRP conversion.

12 *What are the two basic tasks associated with configuring IP multicasts on a Cisco router?*

The first task is to enable IP multicast routing on the router itself. The second task is to configure PIM on associated interfaces.

13 *What process is used within PIM router-to-router communication, which is also used in an OSPF network over a multi-access network?*

The process of election of a Designated Router (DR) is much the same as used in OSPF. The DR is responsible for polling the LAN for host group membership.

14 *Name one of the three reasons why an interface would be placed in the oilist for a multicast group.*

A PIM neighbor was heard on an interface, a host serviced by the interface has joined a group, or the interface was manually configured to join a group.

15 *Define Rendezvous Point and determine under what circumstances Rendezvous Point would be used.*

Rendezvous Point (RP) is used in PIM sparse mode. The RP keeps track of multicast groups. Hosts that send multicast packets are registered with the RP. The RP sends join messages toward the source. Packets are then forwarded on a shared distribution tree. When no RP is known, the packet is flooded in a dense mode fashion.

16 *What troubleshooting command can be used to determine which routers belong to an IGMP group?*

The **ping** command specifying the group multicast address will cause all routers in that group to respond.

- 17 *Assume that you have a router named Router1 connected to other networks from a wide area perspective and a Catalyst switch connected to Router1 on port Ethernet 0. What commands would be required to enable basic CGMP in this network?*

On Router1, you must use the **ip cgmp** command to enable CGMP. On the Catalyst, you must use the **set cgmp enable** command.

- 18 *What happens when a CGMP-enabled router receives an IGMP control packet?*

The router will take the IGMP packet and create a CGMP packet, which contains the request type, the multicast group address, and the MAC address of the host. The router will then send the CGMP packet to a well-known address to which all Catalyst series switches listen. Then, the Catalyst will be able to send that multicast traffic only to the specified MAC address

Answers to Chapter 12 “Do I Know This Already?” Quiz

- 1 *Define an access policy.*

A policy is a firm’s documented standard of network access for their users.

- 2 *What is the access layer defined as?*

The access layer is defined at the point at which a user enters the network.

- 3 *Is HTTP access normally enabled on a Cisco router? What is the main purpose of using HTTP?*

HTTP access is normally disabled on a Cisco router. The main purpose of using it is to make configuration easier.

- 4 *Name at least two components relating to controlling access to network devices.*

Physical security, passwords, virtual terminal access, and privilege levels.

- 5 *What way of accessing a network device requires a password?*

All methods of device access should have a password applied.

- 6 *What feature of the Cisco IOS protects a console connection left unattended?*

Session timeouts provide a level of security by timing out the connection after a predefined period.

- 7 *What does the **access-class** command do when applied to a virtual terminal configuration?*

The **access-class** command is used as a means of allowing only certain hosts access to the virtual terminal lines or outbound traffic.

- 8** *What VLAN is the default VLAN for a Catalyst switch and why is it a good idea to change this?*

The default VLAN is VLAN 1. The reason it's a good idea to change this is that anyone plugging into a port will automatically be in VLAN 1 without further configuration. Because VLAN 1 is commonly used as the management VLAN, this represents a risk of the wrong person getting access to this switch or others within the network.

- 9** *What does port security do on a Catalyst series switch?*

Port security is the feature that can be used to limit access to only specified MAC addresses. All others will be denied access.

- 10** *What is the range of numerical representation of a standard IP access list? An extended access list?*

A standard IP access list can range from 1 to 99. An extended access list can range from 100 to 199, and also 1300 to 1999 in version 12.0 IOS. .

- 11** *Should a standard or an extended access list be used when filtering a particular host?*

When filtering a single host, an extended access list must be used.

- 12** *When implementing route filtering, what type of access list is used—a standard or an extended access list?*

Only standard access lists are used when filtering routes or routing update traffic.

- 13** *In general, what type of policies should be implemented in the core layer?*

Only Quality of Service (QoS) policies should be implemented at the core layer.

Answers to Chapter 12 Q&A Section

- 1** *Define an access policy.*

A policy is a firm's documented standard of network access for their users.

- 2** *What is the access layer defined as?*

The access layer is defined at the point at which a user enters the network.

- 3** *Is HTTP access normally enabled on a Cisco router? What is the main purpose of using HTTP?*

HTTP access is normally disabled on a Cisco router. The main purpose of using it is to make configuration easier.

- 4** *Name at least two components relating to controlling access to network devices.*

Physical security, passwords, virtual terminal access, and privilege levels.

5 *What way of accessing a network device requires a password?*

All methods of device access should have a password applied.

6 *What feature of the Cisco IOS protects a console connection left unattended?*

Session timeouts provide a level of security by timing out the connection after a predefined period.

7 *What does the **access-class** command do when applied to a virtual terminal configuration?*

The **access-class** command is used as a means of allowing only certain hosts access to the virtual terminal lines.

8 *What VLAN is the default VLAN for a Catalyst switch and why is it a good idea to change this?*

The default VLAN is VLAN 1. The reason it's a good idea to change this is that anyone plugging into a port will automatically be in VLAN 1 without further configuration. Because VLAN 1 is commonly used as the management VLAN, this represents a risk of the wrong person getting access to this switch or others within the network.

9 *What does port security do on a Catalyst series switch?*

Port security is the feature that can be used to limit access to only specified MAC addresses. All others will be denied access.

10 *What is the range of numerical representation of a standard IP access list? An extended access list?*

A standard IP access list can range from 1 to 99. An extended access list can range from 100 to 199.

11 *Should a standard or an extended access list be used when filtering a particular host?*

When filtering a single host, an extended access list must be used.

12 *When implementing route filtering, what type of access list is used—a standard or an extended access list?*

Only standard access lists are used when filtering routes or routing update traffic.

13 *In general, what type of policies should be implemented in the core layer?*

Only Quality of Service (QoS) policies should be implemented at the core layer.

14 *Which physical access method of a Cisco router should be disabled if not used?*

The auxiliary port should be disabled if not being used.

15 *What is the virtual terminal connection commonly called?*

A virtual terminal connection is commonly known as a Telnet session.

16 *What does the banner do?*

A banner can be used as a warning message for the purposes of informing potential users that they are connecting to a corporate owned entity and must have specific permission to use it.

17 *Why is it important to have physical security for a network device?*

Because most devices have a back door that can be exploited by physical access.

18 *What does the Cisco command **login local** do on a router?*

The **login local** command requires that you have preconfigured a username/password pair using the command `username xxxx password xxxx`. This is done in lieu of using authentication servers such as TACACS or RADIUS.

Answers to Chapter 13 “Do I Know This Already?” Quiz

1 *What is the main method of out-of-band management for Cisco switches?*

The main method of out-of-band management is the console connection.

2 *What is an application that uses SNMP to perform in-band management?*

CiscoWorks 2000.

3 *CDP operates at what layer of the OSI model?*

CDP operates at the data link layer.

4 *What is the command to verify that RMON is enabled on the switch?*

`show snmp`

5 *Using a troubleshooting model, what step is generally taken after ascertaining all the facts?*

Developing an action plan.

6 *What is the default value for the read-write community string?*

`private`

7 *How many simultaneous Telnet sessions are supported on a Cisco switch?*

Eight.

8 *What command shows information on the modules installed in a particular switch?*

`show module`

- 9 *What command shows the contents of memory displaying MAC addresses and their associated ports?*

show cam dynamic

- 10 *What command would I use if I wanted to see information on Spanning Tree?*

show spantree

- 11 *What physical layer troubleshooting tool is similar to “sonar” in that it bounces a signal to determine length.*

Time domain reflector.

- 12 *What is the name of the tool that decodes the various protocols in captured packets?*

A network analyzer.

Answers to Chapter 13 Q&A Section

- 1 *What is the main method of out-of-band management for Cisco switches?*

The main method of out-of-band management is the console connection.

- 2 *What is an application that uses SNMP to perform in-band management?*

CiscoWorks 2000.

- 3 *CDP operates at what layer of the OSI model?*

CDP operates at the Data Link Layer.

- 4 *What is the command to verify that RMON is enabled on the switch?*

show snmp

- 5 *Using a troubleshooting model, what step is generally taken after ascertaining all the facts?*

Developing an action plan.

- 6 *What is the default value for the read-write community string?*

private

- 7 *How many simultaneous Telnet sessions are supported on a Cisco switch?*

Eight.

- 8 *What command shows information on the modules installed in a particular switch?*

show module

- 9 Explain the function of the **show flash** command.
Displays the Flash code names, version numbers, and sizes.
- 10 What command displays content addressable memory?
show cam dynamic
- 11 What is the command to display CDP information about neighboring systems?
show cdp neighbor
- 12 What command would I use if I wanted to see information on spanning tree for VLAN1?
show spantree
- 13 What physical layer troubleshooting tool is similar to “sonar” in that it bounces a signal to determine length.
Time domain reflector.
- 14 What is the name of the tool that decodes the various protocols in captured packets?
A network analyzer.
- 15 Explain the function of the **show mac** command.
Displays the MAC counters for all of the installed modules.
- 16 Explain the function of the **show config** command.
Displays the current system configuration.
- 17 What command is used to display errors?
show log
- 18 Explain the function of the **show port** command.
Displays the port status and counters for all installed modules.



Numerics

- 14-1 scenario, 464–470
- 14-2 scenario, 465–472
- 14-3 scenario, 467–474
- 80/20 rule, 29

A

- access
 - policies, 398–405
 - remote, 82
 - route processor, 250
 - Telnet client, 438
 - users, 81
 - virtual terminal, 404–405
- access control list (ACL), 284
- access layer, 31
 - policy, 406–407
 - switches, 33
- access-class command, 404, 409
- access-group command, 409
- ACL (access control list), 284
- active routers, 313–314
- Address Resolution Protocol (ARP), 313
- addresses
 - LANE resolution, 218
 - MAC, 313
- addressing
 - ATM, 211–212
 - multicast, 343–344
- administration, switches, 80–85
- advertisements
 - MLS-RP, 271
 - VTP, 115–116
- aging, configuring, 286
- AIPs (ATM Interface Processors), 219
- all routes explorer (ARE), 75
- analyzers
 - network, 453
 - SPAN, 441–442
- ANSI (American National Standard Institute), 73
- answers
 - 14-1 scenario, 469–470
 - 14-2 scenario, 471–472

- 14-3 scenario, 472–474
- Application Specific Integrate Circuits (ASIC), 21, 353
- ARE (all routes explorer), 75
- ARP (Address Resolution Protocol), 307, 313
- ASIC (Application Specific Integrated Circuits), 21, 353
- assigning
 - Ethernet port mode, 87
 - port descriptions, 86
 - port speed, 86
 - ports, 103
 - VLAN ID, 276
- ATM, 70
 - addressing, 211–212
 - LANE (LAN Emulation), 248–249
 - configuring, 220–227
 - membership, 216–220
 - overview, 208–212
 - protocols, 213
 - virtual circuits, 211
- ATM Interface Processors (AIPs), 219
- auto mode, 153
- auto-config-atm-address option, 224
- autonegotiation selection priorities, 72
- Auto-RP, 378
- autosensing ports, 88

B

- B connectors, 75
- bandwidth, 150–155
- BCMSN (Building Cisco Multilayer Switched Networks), 8
- BCRAN (Building Cisco Remote Access Networks), 8
- behavior, active/standby routers, 314
- between, 246
- blocking redundant paths, 160
- blocks
 - core, 43–48
 - switch, 40–42
- BPDU (Bridge Protocol Data Units), 160
- breakout boxes, 453
- Bridge Protocol Data Units (BPDUs), 160
- bridging
 - loops, 156–159

- Root Bridge, 174–179
- Token Rings' transparent bridges, 75–76
- Broadcast and Unknown Server (BUS), 215–217
- broadcast
 - domains, 24
 - traffic, 340
- BSCN (Building Scalable Networks), 8
- Building Cisco Multilayer Switched Networks (BCMSN), 8
- Building Cisco Remote Access Networks (BCRAN), 8
- Building Scalable Cisco Networks (BSCN), 8
- bundling ports, EtherChannel, 150
- BUS (Broadcast and Unknown Server), 215, 217, 223

C

- cables
 - B connectors, 75
 - console port, 77
 - Ethernet, 71
 - Fast Ethernet, 72
 - testers, 452
 - Type 2, 75
- caching
 - entries, 289–290
 - MLS, 272–274
- campus networks
 - configuration, 30–36
 - hierarchical designs, 32–37
 - models, 23–30
 - modular designs, 39–47
- Catalyst switches, 35–37
 - access layer, 33
 - Gigabit Ethernet port cables/connectors, 78
 - managing, 80–85
 - ports, 150
- CBT (Core Based Trees), 361, 373
- CCDP (Cisco Certified Design Professional), 3
- CCIE (Cisco Certified Interworking Expert), 4
- CCNP (Cisco Certified Network Professional), 3
- CCO (Cisco Connection Online), 407
- CDDI (Copper Distribution Data Interface), 70
- CDP (Cisco Discovery Protocol), 83–84, 439
- CEF (Cisco Express Forwarding), 35
- cells, 209
- certifications
 - exams, 5
 - overview, 4–6
 - recommended training paths, 8
- CGMP (Cisco Group Management Protocol), 353, 383–384
- characteristics of multicast traffic, 342
- circuits, virtual ATM, 211
- Cisco
 - exams required for certifications, 5
 - overview of certifications, 4, 6
 - recommended training paths, 8
- Cisco Certified Interworking Expert (CCIE), 4
- Cisco Certified Design Professional (CCDP), 3
- Cisco Certified Network Professional (CCNP), 3
- Cisco Connection Online (CCO), 407
- Cisco Discovery Protocol (CDP), 83–84
- Cisco Discovery Protocol (CDP), 439
- Cisco Express Forwarding (CEF), 35
- Cisco Group Management Protocol (CGMP), 353, 383–384
- Cisco Internetworking Troubleshooting (CIT), 8
- Cisco LAN Switching Configuration (CLSC), 11
- CiscoWorks 2000, 442–443
- CiscoWorks for Switched Networks (CWSI), 105
- CIT (Cisco Internetworking Troubleshooting), 8
- class D multicast addresses, 343
- clear command, 80
- clear mls include all command, 288
- clear mls include command, 288
- clear trunk command, 113
- CLI (command-line interface), 80
- CLI-based switches
 - EtherChannel configuration, 155
 - passwords, 81
 - Port ID, 181
 - Root Bridge configuration, 178–179
 - Root Path Cost, 180
- CLI-based user interfaces, 80
- clients

- LEC, 214
 - Telnet access, 438
- CLSC (Cisco LAN Switching Configuration), 11
- Cluster Builder, 85
- collapsed core blocks, 44
- collisions, 24–26
- command-line interface (CLI), 80
- commands, 324, 433, 447
 - access-class, 404, 409
 - access-group, 409
 - clear, 80
 - clear mls include, 288
 - clear mls include all, 288
 - clear trunk, 113
 - configure terminal, 104
 - copy running-config startup-config, 80
 - debug, 451–453
 - debug standby, 324
 - distribute-list, 413
 - distribution-list, 409
 - enable, 433
 - enable secret level level password, 403
 - encapsulation, 252
 - end, 104
 - exit, 104, 402
 - hostname name, 251
 - interface, 104, 112
 - interface atm, 223
 - ip default-gateway ip-address, 256
 - ip igmp query-interval, 348
 - ip multicast ttl-threshold, 381
 - ip pim rp-address, 379
 - lane server-bus ethernet, 223
 - line vty-number vty-range, 404
 - mls rp ip input-acl, 284
 - mls rp vtp-domain, 276
 - name database, 223
 - network, 255
 - no description, 86
 - no enable password level password, 81
 - no ip multicast-routing, 374
 - no mls rp ip, 275, 278
 - no mls rp ip input-acl, 284
 - no mls rp management-interface, 280
 - no mls rp vlan-id, 276
 - no shutdown, 87, 252
 - no standby group preempt, 319
 - no standby group timers, 320
 - ping, 343, 449
 - port secure max-mac-count, 408
 - privilege level 3, 403
 - privilege levels, 402
 - quit, 438
 - router, 255
 - session, 222
 - session Catalyst switch, 251
 - session switch, 222
 - set, 80, 408
 - set cgmp disable, 384
 - set interface, 433
 - set interface s10, 433
 - set mls agingtime, 286
 - set mls agingtime fast, 286
 - set mls disable, 285
 - set mls enable, 285
 - set mls flow, 283
 - set mls include, 288
 - set port name, 86
 - set snap community, 434
 - set snmp rmon enable, 441
 - set snmp trap enable, 434
 - set spantree portcost, 180
 - set spantree portpri, 181
 - set spantree portvlancost, 180
 - set spantree portvlanpri, 181
 - set spantree priority, 178
 - set spantree root, 178
 - set spantree root secondary, 179
 - set tokenring distrib-crf enable, 129
 - set trunk, 113
 - set vlan 101 3/1,3-7, 105
 - set vlan name, 104
 - setsnmp trap, 434
 - show cdp neighbors detail, 84
 - show interface, 83, 87
 - show ip, 82
 - show ip arp, 313
 - show ip igmp, 351
 - show ip mroute, 381
 - show ip neighbor, 376
 - show lane bus, 225
 - show lane client, 226
 - show lane database, 226

- show lane default, 222, 224
- show lane server, 225
- show mac, 450
- show mls, 287
- show mls entry, 289
- show mls rp, 277, 288
- show mls rp interface interface
 - number, 282
- show mls rp ip interface VLAN41, 282
- show mls rp vtp-domain, 277
- show module, 250
- show port, 450
- show port capabilities, 154
- show port channel info, 155
- show run, 317
- show snmp, 436
- show span, 442
- show spantree, 173
- show spantree module/port, 170
- show standby, 323
- show standby brief, 323
- show system, 446–447
- show users, 402
- show vlan, 105
- show vtp domain, 126
- slip attach, 432
- slip deatch, 432
- span, 442
- standby ip, 317
- standby preempt, 319
- standby use-bia, 310
- switchport access vlan, 104
- switchport mode access, 104
- switchport mode trunk, 112
- switchport trunk allowed vlan add, 112
- switchport trunk encapsulation {is, 112
- switchport trunk vlan remove, 112
- traceroute, 450
- troubleshooting with show commands, 446–447
- undebug all, 324
- vlan database, 104, 120
- vlan name, 104
- vtp {client server transport}, 120
- vtp domain, 120
- vtp password, 120
- comments, assigning, 86
- Common Spanning Tree (CST), 172
- communicating between switches, 83–85
- components
 - IP multicast, 373
 - LANE, 213, 219–220
- configuration
 - access layer port security, 407
 - access lists, 410
 - BUS, 223
 - cache aging, 286
 - campus networks, 30–36
 - CGMP, 383–384
 - dual core blocks, 45–46
 - dynamic VLANs, 105
 - EtherChannel, 154–155
 - filtering routes, 413–415
 - Hello messages, 319
 - hierarchical network design, 32, 34–37
 - hostname/system name, 80
 - HSRP, 317–324
 - HTTP access, 406
 - IGMP, 382
 - InterVLAN routing, 245–256
 - LANE, 220–227
 - LECS, 223–224
 - LES, 223
 - MLS, 275–281, 287–288
 - MLS-SE, 285–290
 - modular network designs, 39, 41–47
 - PIM, 376
 - privilege levels, 402
 - Root Bridge, 178–179
 - route processor, 254–255
 - routes, 413–415
 - RP, 377–380
 - SLIP, 432–433
 - static VLANs, 104
 - STP, 172–187
 - switches, 86–88
 - TTL, 381
 - user access, 81
 - VLAN, 102, 111–113
 - VTP, 119–121
- configure terminal command, 104
- congestion, relieving, 24
- connectivity, 71–74
 - ATM LANE, 213–227

- core blocks, 43–47
 - desktop, 70–74
 - switch blocks, 40–42
 - switches, 77–79
 - VLAN, 251–254
- connectors, 77
- consoles
 - connections, 430
 - port cables, 77
- controlling routing update traffic, 413
- Copper Distribution Data Interface (CDDI), 70
- copy running-config startup-config
 - command, 80
- Core Based Trees (CBT), 361, 373
- core blocks, 43–47
- core layers, 32
 - policy, 415
 - switches, 36
- count parameter, 376
- CRC (cyclic redundancy check), 109, 440
- CST (Common Spanning Tree), 172
- customization, STP, 179–182
- CWSI (CiscoWorks for Switched Networks), 105
- cyclic redundancy check (CRC), 109, 440

D

- datagrams
 - IGMPv1, 347
 - IGMPv2, 350
- debug command, 324, 451–453, 381
- debugging
 - HRSP, 323–324
 - multicast, 381
- default gateways, HSRP, 306
- deleting passwords, 81
- dense mode
 - enabling PIM, 375
 - routing protocols, 358–359
- descriptions, assigning, 86
- designated ports, electing, 165–166
- designated router (DR), 376
- desirable mode, 153
- desktop Ethernet connectivity, 70

- devices
 - access layers, 31
 - core layers, 32
 - distribution layers, 31–32
 - hierarchical network design, 32–37
 - Layer 2 switching, 20
 - Layer 3 routing, 21
 - Layer 4 switching, 22
 - Layer 3 switching, 22
 - LEC, 224
 - MLS (multilayer switching), 23
 - networks
 - access layer policy, 406–407
 - core layer policy, 415
 - distribution layer policy, 408–415
 - managing, 400–405
 - switch blocks, 41–42
 - switches, 83–85
 - VLAN trunks, 106–110
- diagram, 463
- digital multimeters, 452
- disabling MLS, 274
- Distance Vector Multicast Routing Protocol (DVMRP), 359, 373
- distribute-list command, 413
- distribution, 31–35, 44
 - layer policy, 408–413, 415
 - traffic (EtherChannel), 151–153
 - trees, 355
- distribution-list command, 409
- DNS (Domain Name System), 286
- domains, VTP, 114–117, 227–279
- dot1q, 112
- double tagging frames, 109
- DR (designated router), 376
- DRiP (Duplicate Ring Protocol), 131
- DTP (Dynamic Trunking Protocol), 111
- dual core blocks, 45–46
- Duplicate Ring Protocol (DRiP), 131
- DVMRP (Distance Vector Multicast Routing Protocol), 359, 373
- Dynamic Trunking Protocol (DTP), 111
- dynamic VLANs, 103, 105

E

- early token releases, 75
- EBC (Ethernet Bundling Controller), 154
- EIGRP (Enhanced Interior Gateway Routing Protocol), 46, 415
- elan-name parameter, 223
- election
 - designated ports, 165–166
 - queriers using IGMPv2, 350
 - Root Bridge, 163
 - Root Ports, 163–164
- Embedded Remote Monitoring (RMON), 440–441
- enable command, 433
- enable login, 81
- enable secret level level password command, 403
- enabling
 - CDP, 83
 - debug, 323–324
 - IP multicast routing, 374
 - MLS, 278
 - PIM, 374–376
 - port security, 408
 - remote access, 82
 - VTP pruning, 125
- encapsulation command, 252
- end command, 104
- end-system identifier (ESI), 212
- end-to-end VLANs (virtual LANs), 106
- Enhanced Interior Gateway Routing Protocol (EIGRP), 46, 415
- entries, cache, 289–290
- equipment, network test, 451–453
- ESI (end-system identifier), 212
- establishing VLAN connectivity, 251–254
- EtherChannel, 150–155
- Ethernet Bundling Controller (EBC), 154
- Ethernets, 78
 - addresses for multicast, 344
 - desktop connectivity, 70
 - Fast Ethernet, 71–72
 - Gigabit, 73
 - port cables/connectors, 77
 - port mode, 87
 - troubleshooting, 448–449

- exams, 5–12
- exclusive-OR (XOR), 151
- EXEC mode, 81
- exec-timeout, 402
- exit command, 104, 402
- Expires parameter, 377
- extended access list, IP, 411–413
- external router support, 288

F

- Fast EtherChannel (FEC), 73, 150–155
- Fast Ethernet, 71–72
- Fast Simple Server Redundancy Protocol (FSSRP), 220
- FDDI (Fiber Distributed Data Interface), 21, 441
- FDDI (Fiber Distribution Data Interface), 70
- FEC (Fast EtherChannel), 73, 150–155
- Fiber Distributed Data Interface (FDDI), 21, 441
- fiber-optic cable, testing, 452
- fields, TTL scope of delivery, 357
- filtering
 - routes, 413–415
 - traffic, 409
- flooding reverse paths, 359
- Flow Masks, 282, 284
- Forward Delay, 168–170
- forwarding
 - frames, 20–22
 - packets, 21–22
- fox boxes, 453
- frames
 - BPDUs, 160
 - double tagging, 109
 - forwarding, 160
 - Layer 2 switching, 20
 - Layer 3 switching, 22
 - multicast, 24
 - two-link EtherChannel distribution, 151
 - unknown unicast, 157
 - VLANs, 108–110
- FSSRP (Fast Simple Server Redundancy Protocol), 220

full-duplex mode, 71–72
 functionality, 20–23
 functions, TrBRF (Token Ring Bridge Relay Function), 127–128

G

gateways
 default HSRP, 306
 InterVLAN, 256
 GBICs (Gigabit Interface Converters), 78
 GEC (Gigabit EtherChannel), 74, 150–155
 general queries, IGMPv1, 348
 Get Nearest Server (GNS), 24
 get-request format, get-next-request format, 434
 Gigabit EtherChannel (GEC), 74, 150
 Gigabit Ethernets, 73, 78
 Gigabit Interface Converters (GBICs), 78
 GNS (Get Nearest Server), 24
 groups
 HRSP, 314–318
 IGMPv1, 348
 IGMPv2, 352
 members (HSRP), 310–312
 multicast, 341–342
 addressing, 343–344
 maintaining with IGMPv2, 352
 subscribing, 346
 shared distribution trees, 356–357
 source-specific distribution trees, 355

H

hardware-based bridging, 21
 have, 254
 Hello messages, 271, 319
 Hello Time, 170
 hellotime values, 319
 hierarchical network design, 30–37
 holdtime values, 319
 hostname configuration, 80
 hostname name command, 251

Hot Standby Router Protocol (HSRP),
 306–313
 configuring, 317–324
 interfaces, 320–322
 messages, 315
 operations, 313–316
 states, 316
 troubleshooting, 323–324
 HTTP (Hypertext Transfer Protocol) configuring, 406

I–J

IANA (Internet Assigned Numbers Authority), 343
 ICMP Router Discovery Protocol (IRDP), 308
 identifying
 frames, 108–110
 ports, 86
 switches, 80
 IEEE (Institute of Electrical and Electronic Engineers), 70
 802.1Q trunks, 247
 802.1Q protocol, 109–110
 IGMP (Internet Group Management Protocol), 346, 382
 IGMPv1, 347
 IGMPv2, 349–350
 ILMI (Integrated Local Management Interface), 213
 implementation of VLANs (virtual LANs), 105–106
 in-band management, 433–439
 inclusion lists, switches, 289
 input, access lists, 284
 Institute of Electrical and Electronics Engineers (IEEE), 70
 Integrated Local Management Interface (ILMI), 213
 integrated routers, InterVLAN, 249
 interface atm command, 223
 interface command, 104, 112
 interfaces
 HRSP, 317–322
 MLS management, 279–280
 olist (outgoing interface lists), 375
 PIM, 374–376
 VLAN connectivity, 251
 internal tagging, 110

Internet Assigned Numbers Authority (IANA), 343
 Internet Group Management Protocol (IGMP), 346, 382
 Inter-Switch Link (ISL), 109
 interVLAN
 route processor, 250, 254–255
 routing, 245–256
 IOS-based switches, 155, 181
 IP (Internet Protocol)
 addressing schemes (VLANs), 105–106
 extended access list, 411–413
 multicast services, 373–381
 standard access list, 410
 ip default-gateway ip-address command, 256
 ip http access-class 1 statement, 406
 ip igmp query-interval command, 348
 ip multicast ttl-threshold command, 381
 ip pim rp-address command, 379
 ip routing command, 255
 IRDP (ICMP Router Discovery Protocol), 308
 ISL (Inter-Switch Link), 109, 247
 joining groups, 347, 350

K–L

LAN (local access network) segmentation model, 25–30
 LAN Emulation (LANE), 111
 LAN Emulation Client (LEC), 214
 LAN Emulation Configuration Server (LECS), 215–217
 LAN Emulation Server (LES), 215–216
 LANE (LAN Emulation), 111
 address resolution, 218
 ATM, 213–214, 248–249
 configuring, 220–227
 membership, 216–220
 components, 219–220
 viewing, 224–227
 VLAN connectivity, 253
 lane server-bus ethernet command, 223
 Layer 2 switching, 20
 Layer 3 switching, 22, 48
 Layer 4 switching, 22
 Layer 3 routing, 21

layers, 20–23
 access, 31
 policy, 406–407
 switches, 33
 ATM models, 210
 core, 32
 policy, 415
 switches, 36
 distribution, 31–32
 collapsed core blocks, 44
 policy, 408–413, 415
 switches, 34–35
 MLS, 272–275
 physical, 447–453
 leave, CGMP, 384
 leaving groups, 348, 352
 LEC (LAN Emulation Client), 214, 223–224
 LECS (LAN Emulation Configuration Server), 215–217
 LES (LAN Emulation Server), 215–216, 223
 line vty-number vty-range command, 404
 lines, virtual terminal access, 404–405
 links
 bandwidth, 150–155
 multiple physical, 246
 trunks, 247–252
 link-state advertisement (LSA), 360
 lists
 ACL, 284
 input access, 284
 IP, 410–413
 output, 283–284
 switch inclusion, 289
 loads, 151–153
 local VLANs (virtual LANs), 106
 locating virtual router MAC addresses, 313
 login, 81
 login, passwords, 401
 loop-free topology, 161–162
 loops, 156–160
 low latency, high throughput, 209
 LSA (link-state advertisement), 360

M

- maintenance. *See also* troubleshooting
 - IGMPv2, 352
 - multicast groups, 346
- management
 - domains, 119
 - in-band, 433–439
 - MLS, 279–280
 - multicast traffic, 345–354
 - network devices, 400–405
 - out-of-band, 430, 432–433
 - switches, 80–85
 - VTP, 119
- Management Information Base (MIB), 434
- mapping multicast addresses to Ethernets, 344
- masks, Flow Masks, 282–284
- Max Age, 170
- maximum transmission unit (MTU), 104
- members, 310–312
- membership
 - LANE, 216–220
 - queries, IGMPv1, 348
 - VLAN, 103–105
- messages
 - Hello, 271, 319
 - HSRP, 315
- MIB (Management Information Base), 434
- MLS (Multilayer Switching), 23, 34
 - caching, 272–274, 289–290
 - configuring, 287–288
 - disabling, 274
 - enabling, 278
 - Flow Masks, 282, 284
 - management interface, 279–280
- mls rp ip input-acl command, 284
- mls rp vtp-domain command, 276
- MLSP (Multilayer Switching Protocol), 271
- MLS-RP (Multilayer Switching Route Processor), 270
 - advertisements, 271
 - external router support, 288
 - verifying, 280
- MLS-SE (Multilayer Switching Switch Engine), 270, 282–290
- models
 - ATM, 210
 - campus networks, 23–30
 - LAN segmentation, 25–26
 - networks, 28–30
 - shared networks, 24–25
 - troubleshooting, 444–445
- modes
 - auto, 153
 - dense mode routing protocols, 358–359
 - desirable, 153
 - Ethernet port, 87
 - EXEC, 81
 - full-duplex, 71
 - privileged, 81
 - sparse mode routing protocols, 360–361
 - Token Ring port, 88
 - VTP (VLAN Trunking Protocol), 115–120
- modification of TCN, 171
- modular network designs, 39–47
- monitoring
 - RMON, 440–441
 - switches, 430–442
- monitors, network, 453
- MOSPF (Multicast Open Shortest Path First), 359
- MSAUs (multistation access units), 75
- MSFC (Multilayer Switch Feature Card), 35, 250
- MTU (maximum transmission unit), 104
- multicast
 - addresses, 343–344
 - debugging, 381
 - frames, 24
 - overview, 338–342
 - routing protocols, 358–361
 - services, 373–381
 - traffic, 24
 - managing, 345–354
 - routing, 354–357
- Multicast Open Shortest Path First (MOSPF), 359
- Multilayer Switch Feature Card (MSFC), 35, 250
- Multilayer Switching. *See* MLS
- Multilayer Switching Protocol (MLSP), 271
- Multilayer Switching Route Processor (MLS-RP), 270
- Multilayer Switching Switch Engine (MLS-SE), 270

multimedia, 341–342
 multiple groups, 312
 multiple physical links, 246
 multiservice traffic, 209
 multistation access units (MSAUs), 75

N

name database command, 223
 native VLAN, 110
 near-end crosstalk (NEXT), 452
 neighbors, 376
 NetFlow feature Card (NFFC), 35
 NetFlow logic, 34
 network command, 255
 network management system (NMS), 434
 network service access point (NSAP), 211–212
 networks
 access layer, 31
 access policies, 398–405
 analyzers, 453
 campus
 configuration, 30–36
 Ethernet connectivity, 70–74
 models, 23–30
 modular designs, 39–47
 Token Ring connectivity, 74–76
 core layer, 32
 devices
 access layer policy, 406–407
 core layer policy, 415
 distribution layer policy, 408–415
 managing, 400–405
 distribution layer, 31–32, 408–415
 Ethernet, 448–449
 Gigabit Ethernets, 73
 interVLAN routing, 245–256
 IP multicast services, 373–381
 monitors, 453
 multicast traffic
 managing, 345–354
 routing, 354–357
 predictable models, 30
 queriers, 346
 shared models, 24–25
 SPAN, 441–442
 STP, 156–172
 switches, 77–85
 switching functionality, 20–23
 test equipment, 451–453
 testing, 449
 Token Ring, 74–76, 126–131
 traffic models, 28–29
 troubleshooting, 444–445
 unicast traffic, 338
 VLAN, 102
 Network-to-Network INterface (NNI), 208
 NEXT (near-end crosstalk), 452
 NFFC (NetFlow Feature Card), 35
 NMS (network management system), 434
 NNI (Network-to-Network Interface), 208
 no description command, 86
 no enable password level password
 command, 81
 no ip multicast-routing command, 374
 no mls rp ip command, 275, 278
 no mls rp ip input-acl command, 284
 no mls rp management-interface
 command, 280
 no mls rp vlan-id command, 276
 no shutdown command, 87, 252
 no standby group preempt command, 319
 no standby group timers command, 320
 non-root switches, 163–164
 nonvolatile random-access memory (NVRAM), 118
 NSAP (network service access point), 211–212
 NVRAM (nonvolatile random-access memory), 118

O

oilist (outgoing interface lists), 375
 one-armed routers, 247, 249
 one-to-one correspondence, 105–106
 operations, HSRP, 313–316
 optical time domain reflectometer
 (OTDR), 452
 optimization of networks, full-duplex mode, 71
 options
 auto-config-atm-address, 224
 interVLAN routing, 245–256
 OSPF (Open Shortest Path First), 46

OTDR (optical time domain reflectometer), 452
 outgoing interface lists (oilist), 375
 out-of-band management, 430–433
 output
 lists, 283–284
 VTP status, 122
 overview
 certification, 4–6
 ATM, 208–212
 exams, 6
 HSRP, 306
 multicast, 338–342
 topics of exams, 7–8

P

packets
 broadcast traffic, 340
 Flow Masks, 282, 284
 Layer 3 routing, 21
 Layer 4 switching, 22
 MLS, 272–275
 multicast traffic
 managing, 345–354
 routing, 354–357
 PAGP, 153
 unicast traffic, 338
 PAGP (Port Aggregation Protocol), 153
 parameters
 count, 376
 elan-name, 223
 Expires, 377
 passwords, 400–402
 paths, blocking, 160
 PDU (protocol data unit), 20
 permanent host groups, 343
 permanent virtual circuits (PVCs), 211
 Per-VLAN Spanning Tree (PVST), 172
 physical access, network devices, 400
 physical interfaces, establishing, 251
 physical layer, troubleshooting, 447, 449–453
 physical links, routing multiple, 246
 PIM (Protocol Independent Multicast), 373–376
 PIM SM (Protocol Independent Multicast Sparse Mode), 361
 PIMDM (Protocol Independent Multicast Mode), 360
 ping, 82
 ping command, 343, 449
 placement
 of LANE components, 219
 of Root Bridge, 174–179
 PNNI (Private Network-to-Network Interface), 213
 policies
 access, 398–405
 access layer, 406–407
 core layer, 415
 distribution layer, 408–415
 Port Aggregation Protocol (PAGP), 153
 Port ID, tuning, 181
 port secure max-mac-count command, 408
 port-based membership, 103
 ports
 access layer security, 407
 console cables, 77
 designated, 165–166
 EtherChannel, 150
 Ethernet mode, 87
 Ethernets, 77
 Gigabit Ethernets, 78
 identifying, 86
 multicast traffic, 353
 Root, 163–164
 speed, 86
 STP states, 168
 switches, 86–88
 Token Rings, 88
 predictable network models, 30
 preempting HRSP standby, 319
 prefixes, 212
 preserving VLAN identification, 109
 prevention. *See also* security
 collisions, 25–26
 loops, 159–160
 priority HRSP standby, configuring, 318
 Private Network-to-Network Interface (PNNI), 213
 privilege levels, 402–403
 privileged mode, 81
 privilege level 3 command, 403

protocol data unit (PDU), 20
 Protocol Independent Multicast (PIM), 373
 Protocol Independent Multicast Dense Mode (PIMDM), 360
 Protocol Independent Multicast Sparse Mode (PIM SM), 361
 protocols
 analyzers, 453
 ARP, 313
 ATM, 213
 CDP, 439
 CGMP, 353, 383–384
 dense mode routing, 358–359
 DRiP, 131
 DTP, 111
 DVRMP, 359
 EIGRP, 415
 FSSRP, 220
 HSRP, 306–313
 IEEE 802.1Q, 109–110
 IGMP, 346, 382
 IGMPv1, 347
 IGMPv2, 349
 IRDP, 308
 ISL, 109
 MLSP, 271
 multicast routing, 358–361
 PAgP, 153
 RIP, 308
 routing, 22
 SLIP, 432–433
 SNMP, 121, 434–439
 sparse mode routing, 360–361
 SSRP, 220
 STP, 156–172
 TFTP, 286
 VTP, 111
 configuring, 119–121
 pruning, 123–125
 proxy ARP, 307
 pruning VTP (VLAN Trunking Protocol), 123
 PVCs (permanent virtual circuits), 211
 PVST+ (Per-VLAN Spanning Tree Plus), 173

Q–R

queriers, 346–352
 queries, IGMPv1, 348
 quit command, 438
 ranges, 344
 recommendations for studying, 9–12
 redundant bridging, 158
 redundant switches, preventing loops, 159–160
 relationships, VLANs (virtual LANs), 105–106
 relieving network congestion, 24
 remote access, 82
 Rendezvous Point (RP), configuring, 377–380
 replicated unicast traffic, 339
 report suppression, 348
 resolution, addresses
 restricting access, 410. *See also* security
 reverse path flooding, 359
 Reverse Path Forwarding (RPF), 355
 RIF (Routing Information Field), 75
 Ring-In connectors, 75
 Ring-Out connectors, 75
 RIP (Routing Information Protocol), 24, 308
 RJ-45 connectors, 75
 RMON (Embedded Remote Monitoring), 440–441
 rollover cables, 77
 Root Bridge, 161–162, 174–179
 Root Path Cost, 180–181
 Root Ports, electing, 163–164

 route processor
 accessing, 250
 configuring, 254–255
 Route Switch Feature Card (RSFC), 35, 250
 Route Switch Module (RSM), 35, 250
 router command, 255
 routers, 377
 access lists, 410
 active, 314
 distribution trees, 355–357
 external support, 288

- HSRP, 306–313
 - active routers, 313
 - operations, 313–316
 - standby interfaces, 317–318
- integrated routers, 249
- MLS, 272–275, 284
- multicast, 343
- proxy ARP, 307
- TTL field scope of delivery, 357
- virtual, 313
- routes, filtering, 413–415
- routing
 - InterVLAN, 245–256
 - IP multicast, 374
 - Layer3, 21
 - multiple physical links, 246
 - multicast, 373
 - protocols, 358–361
 - traffic, 354–357
 - protocols, 22
 - STP, 156–172
 - trunk links, 247, 249
 - update traffic, 413
- Routing Information Field (RIF), 75
- Routing Information Protocol (RIP), 24, 308
- RP (Rendezvous Point), 377–380
- RPF (Reverse Path Forwarding), 355
- RSFC (Route Switch Feature Card), 35, 250
- RSM (Route Switch Module), 35, 250
- rules, 80/20, 29

S

- SAID (Security Association Identifier), 104
- SAP (Service Advertisement Protocol), 24
- SAR (Segmentation and Reassembly), 209
- scaling
 - core blocks, 47
 - Layer 2 switching, 21
 - link bandwidth, 150–155
- scenarios
 - 14-1, 464, 469–470
 - 14-2, 465–466, 471–472
 - 14-3, 467–468, 472–474
- scope of delivery, TTL fields, 357

- security
 - access policies, 398–407
 - core layer policy, 415
 - distribution layer policy, 408–415
 - passwords, 81
- Security Association Identifier (SAID), 104
- segmentation, 22–26
- Segmentation and Reassembly (SAR), 209
- selecting DR (designated router), 376
- selectors, 212
- Serial Line Internet Protocol (SLIP), 432–433
- servers
 - BUS, 215–217
 - LECS, 216–217
 - LES, 215–216
 - multicast, 341–342
 - unicast traffic, 338
- Service Advertisement Protocol (SAP), 24
- services, multicast, 373–381
- session Catalyst switch command, 251
- session command, 222
- session switch command, 222
- set cgm disable command, 384
- set command, 80, 408
- set interface s10 command, 433
- set mls agingtime command, 286
- set mls agingtime fast command, 286
- set mls disable command, 285
- set mls enable command, 285
- set mls flow command, 283
- set mls include command, 288
- set port name command, 86
- set snap community command, 434
- set snmp rmon enable command, 441
- set snmp trap command, 434
- set snmp trap enable command, 434
- set spantree portcost command, 180
- set spantree portpri command, 181
- set spantree portvlancost commands, 180
- set spantree portvlanpri commands, 181
- set spantree priority command, 178
- set spantree root command, 178
- set spantree root secondary command, 179
- set tokenring distrib-crf enable command, 129
- set trunk command, 113
- set vlan 101 3/1,3-7 command, 105
- set vlan name command, 104

- set-request format, 434
- shared distribution trees, 356–357
- shared network models, 24–25
- shielded twisted-pair, 452
- show cdp neighbors detail command, 84
- show commands, troubleshooting with, 446
- show interface command, 83, 87, 433
- show ip arp command, 313
- show ip command, 82
- show ip igmp command, 351
- show ip mroute command, 381
- show ip pim neighbor command, 376
- show lane bus command, 225
- show lane client command, 226
- show lane database command, 226
- show lane default command, 222, 224
- show lane server command, 225
- show mac command, 450
- show mls command, 287
- show mls entry command, 289
- show mls rp command, 277, 288
- show mls rp interface interface number command, 282
- show mls rp ip interface VLAN41 command, 282
- show mls rp vtp-domain command, 277
- show module command, 250
- show port capabilities command, 154
- show port channel info command, 155
- show port command, 450
- show run command, 317
- show snmp command, 436
- show span command, 442
- show spantree command, 173
- show spantree module/port command, 170
- show standby brief command, 323
- show standby command, 323
- show system command, 446–447
- show users command, 402
- show vlan command, 105
- show vtp domain command, 126
- Simple Network Management Protocol (SNMP), 121, 434–439
- Simple Server Redundancy Protocol (SSRP), 220
- single trunk links, 247, 249
- single-tagging, 110
- sizing
 - dual core blocks, 46
 - switch blocks, 41–42
- SLIP (Serial Line Internet Protocol), 432–433
- slip attach command, 432
- slip detach command, 432
- SMDS (Switched Multimegabit Data Service), 210
- SNMP (Simple Network Management Protocol), 121, 434–439
- source-route bridges (SRBs), 75–76
- source-route transparent bridging (SRT), 76
- source-specific distribution, 355
- SPAN (Switched Port Analyzer), 441–442
- Spanning-Tree Algorithm (STA), 160
- Spanning-Tree Protocol. *See* STP
- sparse mode, enabling, 375
- sparse mode routing protocols, 360–361
- sparse-dense mode, enabling, 376
- speed ports, assigning, 86
- SRBs (source-route bridges), 75–76
- SRT (source-route transparent bridging), 76
- SSRP (Simple Server Redundancy Protocol), 220
- STA (Spanning-Tree Algorithm), 160
- stacking switches, 85
- standard access list, IP, 410
- standby, configuring, 317–318
- standby ip command, 317
- standby preempt command, 319
- standby routers, behavior, 314
- standby states, redundant paths, 160
- standby use-bia command, 310
- statements, 401, 406
- states
 - HSRP, 316
 - STP, 168, 170
- static VLANs, 103–104
- status
 - HSRP, 323
 - VTP, 122
- STP (shielded twisted-pair), 452
- STP (Spanning-Tree Protocol), 156–172
- structures, multicast addresses, 343
- study recommendations, 9, 11–12
- subscribing to multicast groups, 346

support, external routers, 288. *See also* troubleshooting

switch blocks, 40–42

switch port aggregation, EtherChannel, 150–155

Switched Multimegabit Data Service (SMDS), 210

Switched Port Analyzer (SPAN), 441–442

switches

- access layer, 33
- ATM, 208
- BPDUs, 160
- Catalyst, 37
- CLI-based switches
 - EtherChannel configuration, 155
 - Port ID tuning, 181
 - Root Bridge configuration, 178–179
 - Root Path Cost tuning, 180
- clustering, 85
- communicating between, 83–85
- connecting, 77–79
- core layer, 36
- designated ports, 165–166
- distribution layer, 34–35
- identifying, 80
- in-band management, 433–439
- inclusion lists, 289
- InterVLAN routing, 245–256
- IOS-based
 - EtherChannel configuration, 155
 - Root Path Cost tuning, 181
- loops, 159–160
- managing, 80–85
- monitoring, 430–442
- OSI-based switches, 179
- out-of-band management, 430–433
- password configuration, 81
- ports
 - bundling, 150
 - configuration, 86–88
- Root Bridge, 161–162
- Root Port, 163–164
- SPAN, 441–442
- stacking, 85
- STP states, 168, 170
- TCN, 171

- transparent bridging, 158
- VLAN, 104–110
- VTP domains, 279

switching

- functionality, 20–23
- Layer 2, 20
- Layer 3, 22, 48
- Layer 4, 22
- MLS, 272–281
- multicast traffic using CGMP, 353
- overview of exam, 6
- switchport access vlan command, 104
- switchport mode access command, 104
- switchport mode trunk command, 112
- switchport trunk allowed vlan add command, 112
- switchport trunk allowed vlan remove command, 112
- switchport trunk encapsulation {isl dot1q} command, 112
- Sylvan Prometric testing centers, 8
- Sylvan Prometric testing centers, 8
- Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH), 211
- systems, name configuration, 80

T

- Tag Protocol Identifier (TPID), 110
- TCI (Tag Control Information), 110
- TCN (Topology Change Notification), 160, 171
- TDR (time domain reflectometers), 451
- Telnet
 - command, 433
 - client access, 438
- test equipment, networks, 451–453
- testing networks, 449. *See also* troubleshooting
- TFTP (Trivial File Transfer Protocol), 286
- the no debug all command, 324
- time domain refelectometer (TDR), 451
- time domain reflectometer (TDR), 452
- Time-To-Live. *See* TTL
- timers, 170–171, 319
- TLV (Type-Length-Value), 121
- Token Ring Bridge Relay Function (TrBRF), 127–128

- Token Rings
 - bridging, 75–76
 - desktop connectivity, 74–76
 - port cables/connectors, 79
 - port mode, 88
 - VLANs (virtual LANs), 126–131
- topics of exams, 7–8
- topologies
 - loop-free Root Bridges, 161–162
 - TCN, 171
- Topology Change Notification (TCN), 160, 171
- TPID (Tag Protocol Identifier), 110
- traceroute command, 450
- tracking HRSP interfaces, 320
- traffic
 - campus network configuration, 30–36
 - core blocks, 43–47
 - distribution layer, 409
 - EtherChannel, 151–153
 - multicast, 24
 - multiservice, 209
 - multicast, 338–342
 - managing, 345–354
 - routing, 354–357
 - network models, 28–29
 - routing, 413
 - SPAN, 441–442
 - unicast, 338
 - VLAN, 107
- training paths, 8
- transmissions
 - multicast, 343–344
 - unicast, 338
- transparent bridging, 158
- TrBRF (Token Ring Bridge Relay Function), 127–128
- TrCRF, 128–129
- trees, distribution, 355–357
- Trivial File Transfer Protocol (TFTP), 286
- troubleshooting
 - Ethernet, 448–449
 - HRSP, 323
 - networks, 444–449
 - physical layer, 447, 449–453
 - predictable network models, 30
 - show commands, 446
 - traceroute command, 450

- trunks, 28
 - DTP, 111
 - IEEE 802.1Q, 247
 - ISL, 247
 - links
 - enabling VLAN connectivity, 252
 - routing, 247
 - VLANs, 106–113
- TTL (Time-To-Live)
 - configuring, 381
 - HSRP messages, 315
 - multicast fields, 357
- tuning
 - Port ID, 181
 - Root Path Cost, 180–181
- two-link EtherChannel, distributing frames, 151
- Type 2 cables, 75
- Type-Length-Value (TLV), 121
- types of STP, 172

U

- UDP (User Datagram Protocol)
 - HSRP messages, 315
 - multicast traffic, 342
- undebg all command, 324
- undebg standby, 324
- undebg standby command, 324
- UNI User-Network Interface (UNI), 208
- unicast traffic, 338
- unknown unicast, 123
- unknown unicast frames, 157
- unshielded twisted-pair (UTP), 71, 452
- updating traffic, routing, 413
- use of book, 9, 11–12
- username student cisco, 401
- User-Network Interface (UNI), 208
- users
 - access, 81
 - privilege levels, 402–403
- UTP (unshielded twisted-pair), 71, 78, 452

V


- values
 - hellotime, 319
 - holdtime, 319
 - TLV (Type-Length-Value), 121
- VBR (variable bit rate), 210
- verifying
 - MLS-RP, 280, 282
 - PIM configuration, 376
 - port security, 408
 - SNMP, 437
- versions, VTP (VLAN Trunking Protocol), 120–122
- VID (VLAN Identifier), 110
- viewing
 - cache entries, 289–290
 - CDP (Cisco Discovery Protocol), 84
 - EtherChannel configuration, 155
 - LANE configuration, 224–227
 - PIM neighbors, 376
 - VTP domain information, 277
- virtual circuits, ATM, 211
- virtual LANs. *See* VLANs
- virtual path (VP), 211
- virtual router MAC addresses, locating, 313
- virtual terminal access, 404–405
- Virtual Trunk Protocol (VTP), 111
- VLAN (virtual LAN), 28
 - connectivity, 251–254
 - DTP (Dynamic Trunking Protocol), 111
 - IEEE 802.1Q, 109–110
 - ID, 276
 - implementing, 105–106
 - InterVLAN routing, 245–256
 - membership, 103–105
 - Token Rings, 126, 128, 130–131
 - trunks, 106, 108, 110
- vlan database command, 104, 120
- VLAN Identifier (VID), 110
- VLAN Membership Policy Server (VMPS), 105
- vlan name command, 104
- VLAN Trunking Protocol. *See* VTP
- VMPS (VLAN Membership Policy Server), 105
- volt-ohm meters, 452
- VP (virtual path), 211

- VPI/VCI addresses, 212
- VTP (VLAN Trunking Protocol), 111
 - configuring, 119–121
 - domains, 279
 - pruning, 123, 125
 - Token Ring VLANs, 130–131
- vtp {server client transport} command, 120
- vtp domain command, 120
- vtp password command, 120

W–Z

- Web sites, 8
- workgroups, 29
- XOR (exclusive-OR), 151
- XTAGs, 271–272

CCIE Professional Development




Cisco LAN Switching

Kennedy Clark, CCIE; Kevin Hamilton, CCIE

1-57870-094-9 • **AVAILABLE NOW**

This volume provides an in-depth analysis of Cisco LAN switching technologies, architectures, and deployments, including unique coverage of Catalyst network design essentials. Network designs and configuration examples are incorporated throughout to demonstrate the principles and enable easy translation of the material into practice in production networks.




Advanced IP Network Design

Alvaro Retana, CCIE; Don Slice, CCIE; and Russ White, CCIE

1-57870-097-3 • **AVAILABLE NOW**

Network engineers and managers can use these case studies, which highlight various network design goals, to explore issues including protocol choice, network stability, and growth. This book also includes theoretical discussion on advanced design topics.




Large-Scale IP Network Solutions

Khalid Raza, CCIE; and Mark Turner

1-57870-084-1 • **AVAILABLE NOW**

Network engineers can find solutions as their IP networks grow in size and complexity. Examine all the major IP protocols in-depth and learn about scalability, migration planning, network management, and security for large-scale networks.




Routing TCP/IP, Volume I

Jeff Doyle, CCIE

1-57870-041-8 • **AVAILABLE NOW**

This book takes the reader from a basic understanding of routers and routing protocols through a detailed examination of each of the IP interior routing protocols. Learn techniques for designing networks that maximize the efficiency of the protocol being used. Exercises and review questions provide core study for the CCIE Routing and Switching exam.

Cisco Press Fundamentals




Internet Routing Architectures, Second Edition

Sam Halabi with Danny McPherson

1-57870-233-x • **AVAILABLE NOW**

This book explores the ins and outs of interdomain routing network design with emphasis on BGP-4 (Border Gateway Protocol Version 4)--the de facto interdomain routing protocol. You will have all the information you need to make knowledgeable routing decisions for Internet connectivity in your environment.



Voice over IP Fundamentals

Jonathan Davidson and James Peters

1-57870-168-6 • **AVAILABLE NOW**

Voice over IP (VoIP), which integrates voice and data transmission, is quickly becoming an important factor in network communications. It promises lower operational costs, greater flexibility, and a variety of enhanced applications. This book provides a thorough introduction to this new technology to help experts in both the data and telephone industries plan for the new networks.

For the latest on Cisco Press resources and Certification and Training guides, or for information on publishing opportunities, visit www.ciscopress.com