

Scaling Networks v6

Companion Guide

Cisco Networking Academy

Cisco Press

800 East 96th Street

Indianapolis, Indiana 46240 USA

Scaling Networks v6 Companion Guide

Cisco Networking Academy

Copyright © 2018 Cisco Systems, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 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

First Printing August 2017

Library of Congress Control Number: 2017946462

ISBN-13: 978-1-58713-434-0

ISBN-10: 1-58713-434-9

Warning and Disclaimer

This book is designed to provide information about the Cisco Networking Academy Scaling Networks course. 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 authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

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

Editor-in-Chief

Mark Taub

Alliances Manager,

Cisco Press

Ron Fligge

Product Line Manager

Brett Bartow

Executive Editor

Mary Beth Ray

Managing Editor

Sandra Schroeder

Development Editor

Ellie C. Bru

Senior Project Editor

Tonya Simpson

Copy Editor

Kitty Wilson

Technical Editor

Rick McDonald

Editorial Assistant

Vanessa Evans

Cover Designer

Ockomon House

Composition

codeMantra

Indexer

Erika Millen

Proofreader

Abigail Manheim

This book is part of the Cisco Networking Academy® series from Cisco Press. The products in this series support and complement the Cisco Networking Academy curriculum. If you are using this book outside the Networking Academy, then you are not preparing with a Cisco trained and authorized Networking Academy provider.

For more information on the Cisco Networking Academy or to locate a Networking Academy, Please visit www.netacad.com



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.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

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 feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

About the Contributing Authors

Bob Vachon is a professor at Cambrian College in Sudbury, Ontario, Canada, where he teaches network infrastructure courses. He has worked and taught in the computer networking and information technology field since 1984. Since 2002, he has collaborated on various CCNA, CCNA Security, CCNP, Cybersecurity, and IoT projects for the Cisco Networking Academy as team lead, lead author, and subject matter expert. He enjoys playing guitar and being outdoors.

Allan Johnson entered the academic world in 1999, after 10 years as a business owner/operator, to dedicate his efforts to his passion for teaching. He holds both an MBA and an MEd in training and development. He taught CCNA courses at the high school level for seven years and has taught both CCNA and CCNP courses at Del Mar College in Corpus Christi, Texas. In 2003, Allan began to commit much of his time and energy to the CCNA Instructional Support Team, providing services to Networking Academy instructors worldwide and creating training materials. He now works full time for Cisco Networking Academy as curriculum lead.

Contents at a Glance

	Introduction	xx
Chapter 1	LAN Design	1
Chapter 2	Scaling VLANs	47
Chapter 3	STP	105
Chapter 4	EtherChannel and HSRP	179
Chapter 5	Dynamic Routing	219
Chapter 6	EIGRP	273
Chapter 7	EIGRP Tuning and Troubleshooting	365
Chapter 8	Single-Area OSPF	415
Chapter 9	Multiarea OSPF	493
Chapter 10	OSPF Tuning and Troubleshooting	527
Appendix A	Answers to the Review Questions	591
	Glossary	603
	Index	621

Contents

	Introduction	xx
Chapter 1	LAN Design	1
	Objectives	1
	Key Terms	1
	Introduction (1.0.1.1)	3
	Campus Wired LAN Designs (1.1)	4
	Cisco Validated Designs (1.1.1)	4
	<i>The Need to Scale the Network (1.1.1.1)</i>	4
	<i>Hierarchical Design Model (1.1.1.2)</i>	6
	Expanding the Network (1.1.2)	8
	<i>Design for Scalability (1.1.2.1)</i>	8
	<i>Planning for Redundancy (1.1.2.2)</i>	10
	<i>Failure Domains (1.1.2.3)</i>	11
	<i>Increasing Bandwidth (1.1.2.4)</i>	13
	<i>Expanding the Access Layer (1.1.2.5)</i>	14
	<i>Fine-tuning Routing Protocols (1.1.2.6)</i>	15
	Selecting Network Devices (1.2)	17
	Switch Hardware (1.2.1)	17
	<i>Switch Platforms (1.2.1.1)</i>	17
	<i>Port Density (1.2.1.2)</i>	21
	<i>Forwarding Rates (1.2.1.3)</i>	22
	<i>Power over Ethernet (1.2.1.4)</i>	23
	<i>Multilayer Switching (1.2.1.5)</i>	24
	Router Hardware (1.2.2)	26
	<i>Router Requirements (1.2.2.1)</i>	26
	<i>Cisco Routers (1.2.2.2)</i>	27
	<i>Router Hardware (1.2.2.3)</i>	28
	Managing Devices (1.2.3)	29
	<i>Managing IOS Files and Licensing (1.2.3.1)</i>	30
	<i>In-Band versus Out-of-Band Management (1.2.3.2)</i>	30
	<i>Basic Router CLI Commands (1.2.3.3)</i>	31
	<i>Basic Router Show Commands (1.2.3.4)</i>	34
	<i>Basic Switch CLI Commands (1.2.3.5)</i>	38
	<i>Basic Switch Show Commands (1.2.3.6)</i>	40
	Summary (1.3)	43
	Practice	44
	Check Your Understanding Questions	45

Chapter 2	Scaling VLANs	47
	Objectives	47
	Key Terms	47
	Introduction (2.0.1.1)	48
	VTP, Extended VLANs, and DTP (2.1)	48
	VTP Concepts and Operation (2.1.1)	49
	<i>VTP Overview (2.1.1.1)</i>	49
	<i>VTP Modes (2.1.1.2)</i>	50
	<i>VTP Advertisements (2.1.1.3)</i>	52
	<i>VTP Versions (2.1.1.4)</i>	53
	<i>Default VTP Configuration (2.1.1.5)</i>	53
	<i>VTP Caveats (2.1.1.6)</i>	55
	VTP Configuration (2.1.2)	57
	<i>VTP Configuration Overview (2.1.2.1)</i>	57
	<i>Step 1—Configure the VTP Server (2.1.2.2)</i>	58
	<i>Step 2—Configure the VTP Domain Name and Password (2.1.2.3)</i>	59
	<i>Step 3—Configure the VTP Clients (2.1.2.4)</i>	60
	<i>Step 4—Configure VLANs on the VTP Server (2.1.2.5)</i>	60
	<i>Step 5—Verify That the VTP Clients Have Received the New VLAN Information (2.1.2.6)</i>	62
	Extended VLANs (2.1.3)	63
	<i>VLAN Ranges on Catalyst Switches (2.1.3.1)</i>	63
	<i>Creating a VLAN (2.1.3.2)</i>	65
	<i>Assigning Ports to VLANs (2.1.3.3)</i>	66
	<i>Verifying VLAN Information (2.1.3.4)</i>	67
	<i>Configuring Extended VLANs (2.1.3.5)</i>	69
	Dynamic Trunking Protocol (2.1.4)	71
	<i>Introduction to DTP (2.1.4.1)</i>	71
	<i>Negotiated Interface Modes (2.1.4.2)</i>	72
	Troubleshoot Multi-VLAN Issues (2.2)	75
	Inter-VLAN Configuration Issues (2.2.1)	75
	<i>Deleting VLANs (2.2.1.1)</i>	75
	<i>Switch Port Issues (2.2.1.2)</i>	77
	<i>Verify Switch Configuration (2.2.1.3)</i>	79
	<i>Interface Issues (2.2.1.4)</i>	81
	<i>Verify Routing Configuration (2.2.1.5)</i>	82
	IP Addressing Issues (2.2.2)	83
	<i>Errors with IP Addresses and Subnet Masks (2.2.2.1)</i>	83
	<i>Verifying IP Address and Subnet Mask Configuration Issues (2.2.2.2)</i>	85
	VTP and DTP Issues (2.2.3)	88

Troubleshoot VTP Issues (2.2.3.1) 88

Troubleshoot DTP Issues (2.2.3.2) 89

Layer 3 Switching (2.3) 89

Layer 3 Switching Operation and Configuration (2.3.1) 90

Introduction to Layer 3 Switching (2.3.1.1) 90

Inter-VLAN Routing with Switch Virtual

Interfaces (2.3.1.2) 91

Inter-VLAN Routing with Switch Virtual

Interfaces (Con't.) (2.3.1.3) 92

Inter-VLAN Routing with Routed Ports (2.3.1.4) 94

Troubleshoot Layer 3 Switching (2.3.2) 95

Layer 3 Switch Configuration Issues (2.3.2.1) 95

Example: Troubleshooting Layer 3 Switching (2.3.2.2) 96

Summary (2.4) 99

Practice 99

Check Your Understanding Questions 100

Chapter 3

STP 105

Objectives 105

Key Terms 105

Introduction (3.0.1.1) 107

Spanning Tree Concepts (3.1) 108

Purpose of Spanning Tree (3.1.1) 108

Redundancy at OSI Layers 1 and 2 (3.1.1.1) 108

Issues with Layer 1 Redundancy: MAC Database

Instability (3.1.1.2) 109

Issues with Layer 1 Redundancy: Broadcast

Storms (3.1.1.3) 111

Issues with Layer 1 Redundancy: Duplicate Unicast

Frames (3.1.1.4) 113

STP Operation (3.1.2) 114

Spanning Tree Algorithm: Introduction (3.1.2.1) 114

Spanning Tree Algorithm: Port Roles (3.1.2.2) 117

Spanning Tree Algorithm: Root Bridge (3.1.2.3) 119

Spanning Tree Algorithm: Root Path Cost (3.1.2.4) 121

Port Role Decisions for RSTP (3.1.2.5) 124

Designated and Alternate Ports (3.1.2.6) 127

802.1D BPDU Frame Format (3.1.2.7) 128

802.1D BPDU Propagation and Process (3.1.2.8) 131

Extended System ID (3.1.2.9) 136

Varieties of Spanning Tree Protocols (3.2) 140

Overview (3.2.1) 140

	<i>Types of Spanning Tree Protocols (3.2.1.1)</i>	140
	<i>Characteristics of the Spanning Tree Protocols (3.2.1.2)</i>	141
	PVST+ (3.2.2)	143
	<i>Overview of PVST+ (3.2.2.1)</i>	143
	<i>Port States and PVST+ Operation (3.2.2.2)</i>	144
	<i>Extended System ID and PVST+ Operation (3.2.2.3)</i>	146
	Rapid PVST+ (3.2.3)	148
	<i>Overview of Rapid PVST+ (3.2.3.1)</i>	148
	<i>RSTP BPDUs (3.2.3.2)</i>	149
	<i>Edge Ports (3.2.3.3)</i>	150
	<i>Link Types (3.2.3.4)</i>	152
	Spanning Tree Configuration (3.3)	153
	PVST+ Configuration (3.3.1)	153
	<i>Catalyst 2960 Default Configuration (3.3.1.1)</i>	153
	<i>Configuring and Verifying the Bridge ID (3.3.1.2)</i>	154
	<i>PortFast and BPDU Guard (3.3.1.3)</i>	156
	<i>PVST+ Load Balancing (3.3.1.4)</i>	158
	Rapid PVST+ Configuration (3.3.2)	160
	<i>Spanning Tree Mode (3.3.2.1)</i>	161
	STP Configuration Issues (3.3.3)	163
	<i>Analyzing the STP Topology (3.3.3.1)</i>	164
	<i>Expected Topology versus Actual Topology (3.3.3.2)</i>	164
	<i>Overview of Spanning Tree Status (3.3.3.3)</i>	165
	<i>Spanning Tree Failure Consequences (3.3.3.4)</i>	166
	<i>Repairing a Spanning Tree Problem (3.3.3.5)</i>	169
	Switch Stacking and Chassis Aggregation (3.3.4)	169
	<i>Switch Stacking Concepts (3.3.4.1)</i>	169
	<i>Spanning Tree and Switch Stacks (3.3.4.2)</i>	171
	Summary (3.4)	173
	Practice	174
	Check Your Understanding Questions	174
Chapter 4	EtherChannel and HSRP	179
	Objectives	179
	Key Terms	179
	Introduction (4.0.1.1)	180
	Link Aggregation Concepts (4.1)	181
	Link Aggregation (4.1.1)	181
	<i>Introduction to Link Aggregation (4.1.1.1)</i>	181
	<i>Advantages of EtherChannel (4.1.1.2)</i>	182

- EtherChannel Operation (4.1.2) 183
 - Implementation Restrictions (4.1.2.1)* 183
 - Port Aggregation Protocol (4.1.2.2)* 185
 - Link Aggregation Control Protocol (4.1.2.3)* 186

Link Aggregation Configuration (4.2) 188

- Configuring EtherChannel (4.2.1) 188
 - Configuration Guidelines (4.2.1.1)* 188
 - Configuring Interfaces (4.2.1.2)* 189
- Verifying and Troubleshooting EtherChannel (4.2.2) 191
 - Verifying EtherChannel (4.2.2.1)* 191
 - Troubleshooting EtherChannel (4.2.2.2)* 194

First Hop Redundancy Protocols (4.3) 198

- Concept of First Hop Redundancy Protocols (4.3.1) 198
 - Default Gateway Limitations (4.3.1.1)* 198
 - Router Redundancy (4.3.1.2)* 199
 - Steps for Router Failover (4.3.1.3)* 200
 - First Hop Redundancy Protocols (4.3.1.5)* 201
- HSRP Operations (4.3.2) 202
 - HSRP Overview (4.3.2.1)* 203
 - HSRP Versions (4.3.2.2)* 204
 - HSRP Priority and Preemption (4.3.2.3)* 204
 - HSRP States and Timers (4.3.2.4)* 205
- HSRP Configuration (4.3.3) 206
 - HSRP Configuration Commands (4.3.3.1)* 206
 - HSRP Sample Configuration (4.3.3.2)* 207
 - HSRP Verification (4.3.3.3)* 208
- HSRP Troubleshooting (4.3.4) 209
 - HSRP Failure (4.3.4.1)* 209
 - HSRP Debug Commands (4.3.4.2)* 210
 - Common HSRP Configuration Issues (4.3.4.3)* 213

Summary (4.4) 214

Practice 215

Check Your Understanding Questions 216

Chapter 5 Dynamic Routing 219

Objectives 219

Key Terms 219

Introduction (5.0.1.1) 221

Dynamic Routing Protocols (5.1) 222

- Types of Routing Protocols (5.1.1) 222
 - Classifying Routing Protocols (5.1.1.1)* 222

<i>IGP and EGP Routing Protocols (5.1.1.2)</i>	224
<i>Distance Vector Routing Protocols (5.1.1.3)</i>	226
<i>Link-State Routing Protocols (5.1.1.4)</i>	226
<i>Classful Routing Protocols (5.1.1.5)</i>	228
<i>Classless Routing Protocols (5.1.1.6)</i>	231
<i>Routing Protocol Characteristics (5.1.1.7)</i>	233
<i>Routing Protocol Metrics (5.1.1.8)</i>	234
Distance Vector Dynamic Routing (5.2)	236
Distance Vector Fundamentals (5.2.1)	236
<i>Dynamic Routing Protocol Operation (5.2.1.1)</i>	236
<i>Cold Start (5.2.1.2)</i>	237
<i>Network Discovery (5.2.1.3)</i>	238
<i>Exchanging the Routing Information (5.2.1.4)</i>	239
<i>Achieving Convergence (5.2.1.5)</i>	241
Distance Vector Routing Protocol Operation (5.2.2)	242
<i>Distance Vector Technologies (5.2.2.1)</i>	242
<i>Distance Vector Algorithm (5.2.2.2)</i>	242
Types of Distance Vector Routing Protocols (5.2.3)	245
<i>Routing Information Protocol (5.2.3.1)</i>	245
<i>Enhanced Interior-Gateway Routing Protocol (5.2.3.2)</i>	246
Link-State Dynamic Routing (5.3)	248
Link-State Routing Protocol Operation (5.3.1)	248
<i>Shortest Path First Protocols (5.3.1.1)</i>	248
<i>Dijkstra's Algorithm (5.3.1.2)</i>	248
<i>SPF Example (5.3.1.3)</i>	249
Link-State Updates (5.3.2)	251
<i>Link-State Routing Process (5.3.2.1)</i>	251
<i>Link and Link-State (5.3.2.2)</i>	252
<i>Say Hello (5.3.2.3)</i>	256
<i>Building the Link-State Packet (5.3.2.4)</i>	257
<i>Flooding the LSP (5.3.2.5)</i>	258
<i>Building the Link-State Database (5.3.2.6)</i>	259
<i>Building the SPF Tree (5.3.2.7)</i>	260
<i>Adding OSPF Routes to the Routing Table (5.3.2.8)</i>	264
Link-State Routing Protocol Benefits (5.3.3)	264
<i>Why Use Link-State Protocols? (5.3.3.1)</i>	264
<i>Disadvantages of Link-State Protocols (5.3.3.2)</i>	265
<i>Protocols That Use Link-State (5.3.3.3)</i>	267
Summary (5.4)	268
Practice	269
Check Your Understanding Questions	269

Chapter 6 EIGRP 273

Objectives 273

Key Terms 273

Introduction (6.0.1.1) 274

EIGRP Characteristics (6.1) 274

EIGRP Basic Features (6.1.1) 274

Features of EIGRP (6.1.1.1) 274

Protocol Dependent Modules (6.1.1.2) 276

Reliable Transport Protocol (6.1.1.3) 278

Authentication (6.1.1.4) 279

EIGRP Packet Types (6.1.2) 279

EIGRP Packet Types (6.1.2.1) 279

EIGRP Hello Packets (6.1.2.2) 280

EIGRP Update and Acknowledgment Packets (6.1.2.3) 281

EIGRP Query and Reply Packets (6.1.2.4) 283

EIGRP Messages (6.1.3) 284

Encapsulating EIGRP Messages (6.1.3.1) 284

EIGRP Packet Header and TLV (6.1.3.2) 285

Implement EIGRP for IPv4 (6.2) 289

Configure EIGRP with IPv4 (6.2.1) 289

EIGRP Network Topology (6.2.1.1) 289

Autonomous System Numbers (6.2.1.2) 291

*The **router eigrp** Command (6.2.1.3) 292*

EIGRP Router ID (6.2.1.4) 293

Configuring the EIGRP Router ID (6.2.1.5) 295

*The **network** Command (6.2.1.6) 296*

*The **network** Command and Wildcard Mask (6.2.1.7) 298*

Passive Interface (6.2.1.8) 300

Verify EIGRP with IPv4 (6.2.2) 302

Verifying EIGRP: Examining Neighbors (6.2.2.1) 302

*Verifying EIGRP: **show ip protocols***

Command (6.2.2.2) 304

Verifying EIGRP: Examine the IPv4 Routing

Table (6.2.2.3) 306

EIGRP Operation (6.3) 309

EIGRP Initial Route Discovery (6.3.1) 309

EIGRP Neighbor Adjacency (6.3.1.1) 310

EIGRP Topology Table (6.3.1.2) 311

EIGRP Convergence (6.3.1.3) 312

EIGRP Metrics (6.3.2) 313

EIGRP Composite Metric (6.3.2.1) 313

Examining Interface Metric Values (6.3.2.2) 315

	<i>Bandwidth Metric (6.3.2.3)</i>	316
	<i>Delay Metric (6.3.2.4)</i>	319
	<i>How to Calculate the EIGRP Metric (6.3.2.5)</i>	320
	<i>Calculating the EIGRP Metric (6.3.2.6)</i>	321
	DUAL and the Topology Table (6.3.3)	323
	<i>DUAL Concepts (6.3.3.1)</i>	323
	<i>Introduction to DUAL (6.3.3.2)</i>	324
	<i>Successor and Feasible Distance (6.3.3.3)</i>	324
	<i>Feasible Successors, Feasibility Condition, and Reported Distance (6.3.3.4)</i>	326
	<i>Topology Table: show ip eigrp topology Command (6.3.3.5)</i>	328
	<i>Topology Table: show ip eigrp topology Command (Cont.) (6.3.3.6)</i>	329
	<i>Topology Table: No Feasible Successor (6.3.3.7)</i>	332
	DUAL and Convergence (6.3.4)	334
	<i>DUAL Finite State Machine (FSM) (6.3.4.1)</i>	334
	<i>DUAL: Feasible Successor (6.3.4.2)</i>	335
	<i>DUAL: No Feasible Successor (6.3.4.3)</i>	338
	Implement EIGRP for IPv6 (6.4)	341
	EIGRP for IPv6 (6.4.1)	341
	<i>EIGRP for IPv6 (6.4.1.1)</i>	341
	<i>Compare EIGRP for IPv4 and IPv6 (6.4.1.2)</i>	342
	<i>IPv6 Link-local Addresses (6.4.1.3)</i>	344
	Configure EIGRP for IPv6 (6.4.2)	345
	<i>EIGRP for IPv6 Network Topology (6.4.2.1)</i>	345
	<i>Configuring IPv6 Link-local Addresses (6.4.2.2)</i>	347
	<i>Configuring the EIGRP for IPv6 Routing Process (6.4.2.3)</i>	349
	<i>The ipv6 eigrp Interface Command (6.4.2.4)</i>	350
	Verifying EIGRP for IPv6 (6.4.3)	352
	<i>IPv6 Neighbor Table (6.4.3.1)</i>	352
	<i>The show ip protocols Command (6.4.3.2)</i>	354
	<i>The EIGRP for IPv6 Routing Table (6.4.3.3)</i>	355
	Summary (6.5)	358
	Practice	359
	Check Your Understanding Questions	360
Chapter 7	EIGRP Tuning and Troubleshooting	365
	Objectives	365
	Key Terms	365
	Introduction (7.0.1.1)	366

Tune EIGRP (7.1) 366

- Automatic Summarization (7.1.1) 366
 - Network Topology (7.1.1.1) 367*
 - EIGRP Automatic Summarization (7.1.1.2) 369*
 - Configuring EIGRP Automatic Summarization (7.1.1.3) 371*
 - Verifying Auto-Summary: show ip protocols (7.1.1.4) 372*
 - Verifying Auto-Summary: Topology Table (7.1.1.5) 375*
 - Verifying Auto-Summary: Routing Table (7.1.1.6) 376*
 - Summary Route (7.1.1.7) 378*
 - Summary Route (Cont.) (7.1.1.8) 379*
- Default Route Propagation (7.1.2) 380
 - Propagating a Default Static Route (7.1.2.1) 380*
 - Verifying the Propagated Default Route (7.1.2.2) 382*
 - EIGRP for IPv6: Default Route (7.1.2.3) 383*
- Fine-tuning EIGRP Interfaces (7.1.3) 384
 - EIGRP Bandwidth Utilization (7.1.3.1) 385*
 - Hello and Hold Timers (7.1.3.2) 386*
 - Load Balancing IPv4 (7.1.3.3) 388*
 - Load Balancing IPv6 (7.1.3.4) 390*

Troubleshoot EIGRP (7.2) 392

- Components of Troubleshooting EIGRP (7.2.1) 392
 - Basic EIGRP Troubleshooting Commands (7.2.1.1) 392*
 - Components (7.2.1.2) 394*
- Troubleshoot EIGRP Neighbor Issues (7.2.2) 397
 - Layer 3 Connectivity (7.2.2.1) 397*
 - EIGRP Parameters (7.2.2.2) 398*
 - EIGRP Interfaces (7.2.2.3) 399*
- Troubleshoot EIGRP Routing Table Issues (7.2.3) 401
 - Passive Interface (7.2.3.1) 401*
 - Missing Network Statement (7.2.3.2) 403*
 - Autosummarization (7.2.3.3) 405*

Summary (7.3) 410

Practice 411

Check Your Understanding Questions 412

Chapter 8 Single-Area OSPF 415

Objectives 415

Key Terms 415

Introduction (8.0.1.1) 416

OSPF Characteristics (8.1) 416

Open Shortest Path First (8.1.1)	416
<i>Evolution of OSPF (8.1.1.1)</i>	417
<i>Features of OSPF (8.1.1.2)</i>	418
<i>Components of OSPF (8.1.1.3)</i>	419
<i>Link-State Operation (8.1.1.4)</i>	420
<i>Single-Area and Multiarea OSPF (8.1.1.5)</i>	424
OSPF Messages (8.1.2)	426
<i>Encapsulating OSPF Messages (8.1.2.1)</i>	426
<i>Types of OSPF Packets (8.1.2.2)</i>	428
<i>Hello Packet (8.1.2.3)</i>	428
<i>Hello Packet Intervals (8.1.2.4)</i>	430
<i>Link-State Updates (8.1.2.5)</i>	430
OSPF Operation (8.1.3)	431
<i>OSPF Operational States (8.1.3.1)</i>	432
<i>Establish Neighbor Adjacencies (8.1.3.2)</i>	433
<i>OSPF DR and BDR (8.1.3.3)</i>	435
<i>Synchronizing OSPF Databases (8.1.3.4)</i>	438
Single-Area OSPFv2 (8.2)	440
OSPF Router ID (8.2.1)	441
<i>OSPF Network Topology (8.2.1.1)</i>	441
<i>Router OSPF Configuration Mode (8.2.1.2)</i>	442
<i>Router IDs (8.2.1.3)</i>	442
<i>Configuring an OSPF Router ID (8.2.1.4)</i>	444
<i>Modifying a Router ID (8.2.1.5)</i>	445
<i>Using a Loopback Interface as the Router ID (8.2.1.6)</i>	447
Configure Single-Area OSPFv2 (8.2.2)	448
<i>Enabling OSPF on Interfaces (8.2.2.1)</i>	448
<i>Wildcard Mask (8.2.2.2)</i>	448
<i>The network Command (8.2.2.3)</i>	449
<i>Passive Interface (8.2.2.4)</i>	450
<i>Configuring Passive Interfaces (8.2.2.5)</i>	451
OSPF Cost (8.2.3)	453
<i>OSPF Metric = Cost (8.2.3.1)</i>	454
<i>OSPF Accumulates Costs (8.2.3.2)</i>	455
<i>Adjusting the Reference Bandwidth (8.2.3.3)</i>	456
<i>Default Interface Bandwidths (8.2.3.4)</i>	460
<i>Adjusting the Interface Bandwidth (8.2.3.5)</i>	462
<i>Manually Setting the OSPF Cost (8.2.3.6)</i>	463
Verify OSPF (8.2.4)	464
<i>Verify OSPF Neighbors (8.2.4.1)</i>	465
<i>Verify OSPF Protocol Settings (8.2.4.2)</i>	466
<i>Verify OSPF Process Information (8.2.4.3)</i>	466
<i>Verify OSPF Interface Settings (8.2.4.4)</i>	468

Single-Area OSPFv3 (8.3) 469

- OSPFv2 vs. OSPFv3 (8.3.1) 469
 - OSPFv3 (8.3.1.1)* 469
 - Similarities Between OSPFv2 and OSPFv3 (8.3.1.2)* 471
 - Differences Between OSPFv2 and OSPFv3 (8.3.1.3)* 471
 - Link-Local Addresses (8.3.1.4)* 472
- Configuring OSPFv3 (8.3.2) 473
 - OSPFv3 Network Topology (8.3.2.1)* 473
 - Link-Local Addresses (8.3.2.2)* 475
 - Assigning Link-Local Addresses (8.3.2.3)* 476
 - Configuring the OSPFv3 Router ID (8.3.2.4)* 477
 - Modifying an OSPFv3 Router ID (8.3.2.5)* 479
 - Enabling OSPFv3 on Interfaces (8.3.2.6)* 481
- Verify OSPFv3 (8.3.3) 481
 - Verify OSPFv3 Neighbors (8.3.3.1)* 482
 - Verify OSPFv3 Protocol Settings (8.3.3.2)* 483
 - Verify OSPFv3 Interfaces (8.3.3.3)* 483
 - Verify the IPv6 Routing Table (8.3.3.4)* 484

Summary (8.4) 486**Practice 487****Check Your Understanding Questions 488****Chapter 9 Multiarea OSPF 493****Objectives 493****Key Terms 493****Introduction (9.0.1.1) 494****Multiarea OSPF Operation (9.1) 494**

- Why Multiarea OSPF? (9.1.1) 494
 - Single-Area OSPF (9.1.1.1)* 494
 - Multiarea OSPF (9.1.1.2)* 495
 - OSPF Two-Layer Area Hierarchy (9.1.1.3)* 498
 - Types of OSPF Routers (9.1.1.4)* 499
- Multiarea OSPF LSA Operation (9.1.2) 501
 - OSPF LSA Types (9.1.2.1)* 502
 - OSPF LSA Type 1 (9.1.2.2)* 502
 - OSPF LSA Type 2 (9.1.2.3)* 503
 - OSPF LSA Type 3 (9.1.2.4)* 504
 - OSPF LSA Type 4 (9.1.2.5)* 505
 - OSPF LSA Type 5 (9.1.2.6)* 506
- OSPF Routing Table and Types of Routes (9.1.3) 506
 - OSPF Routing Table Entries (9.1.3.1)* 507
 - OSPF Route Calculation (9.1.3.2)* 508

	Configuring Multiarea OSPF (9.2)	509
	Configuring Multiarea OSPF (9.2.1)	510
	<i>Implementing Multiarea OSPF (9.2.1.1)</i>	510
	<i>Configuring Multiarea OSPFv2 (9.2.1.2)</i>	511
	<i>Configuring Multiarea OSPFv3 (9.2.1.3)</i>	513
	Verifying Multiarea OSPF (9.2.2)	515
	<i>Verifying Multiarea OSPFv2 (9.2.2.1)</i>	515
	<i>Verify General Multiarea OSPFv2 Settings (9.2.2.2)</i>	515
	<i>Verify the OSPFv2 Routes (9.2.2.3)</i>	516
	<i>Verify the Multiarea OSPFv2 LSDB (9.2.2.4)</i>	517
	<i>Verify Multiarea OSPFv3 (9.2.2.5)</i>	518
	Summary (9.3)	522
	Practice	523
	Check Your Understanding Questions	524
Chapter 10	OSPF Tuning and Troubleshooting	527
	Objectives	527
	Key Terms	527
	Introduction (10.0.1.1)	528
	Advanced Single-Area OSPF Configurations (10.1)	528
	OSPF in Multiaccess Networks (10.1.1)	528
	<i>OSPF Network Types (10.1.1.1)</i>	528
	<i>Challenges in Multiaccess Networks (10.1.1.2)</i>	531
	<i>OSPF Designated Router (10.1.1.3)</i>	533
	<i>Verifying DR/BDR Roles (10.1.1.4)</i>	535
	<i>Verifying DR/BDR Adjacencies (10.1.1.5)</i>	538
	<i>Default DR/BDR Election Process (10.1.1.6)</i>	540
	<i>DR/BDR Election Process (10.1.1.7)</i>	541
	<i>The OSPF Priority (10.1.1.8)</i>	544
	<i>Changing the OSPF Priority (10.1.1.9)</i>	544
	Default Route Propagation (10.1.2)	547
	<i>Propagating a Default Static Route in OSPFv2 (10.1.2.1)</i>	547
	<i>Verifying the Propagated IPv4 Default Route (10.1.2.2)</i>	549
	<i>Propagating a Default Static Route in OSPFv3 (10.1.2.3)</i>	551
	<i>Verifying the Propagated IPv6 Default Route (10.1.2.4)</i>	552
	Fine-tuning OSPF Interfaces (10.1.3)	554
	<i>OSPF Hello and Dead Intervals (10.1.3.1)</i>	554
	<i>Modifying OSPFv2 Intervals (10.1.3.2)</i>	555
	<i>Modifying OSPFv3 Intervals (10.1.3.3)</i>	557

Troubleshooting Single-Area OSPF Implementations (10.2) 560

Components of Troubleshooting Single-Area OSPF (10.2.1) 560

Overview (10.2.1.1) 560

OSPF States (10.2.1.2) 560

OSPF Troubleshooting Commands (10.2.1.3) 562

Components of Troubleshooting OSPF (10.2.1.4) 566

Troubleshoot Single-Area OSPFv2 Routing Issues (10.2.2) 569

Troubleshooting Neighbor Issues (10.2.2.1) 569

Troubleshooting OSPFv2 Routing Table

Issues (10.2.2.2) 573

Troubleshoot Single-Area OSPFv3 Routing Issues (10.2.3) 576

OSPFv3 Troubleshooting Commands (10.2.3.1) 576

Troubleshooting OSPFv3 (10.2.3.2) 580

Troubleshooting Multiarea OSPFv2 and OSPFv3 (10.2.4) 582

Multiarea OSPF Troubleshooting Skills (10.2.4.1) 582

Multiarea OSPF Troubleshooting Data

Structures (10.2.4.2) 583

Summary (10.3) 585

Practice 587

Check Your Understanding Questions 587

Appendix A Answers to the Review Questions 591

Glossary 603

Index 621

Reader Services

Register your copy at www.ciscopress.com/title/9781587134340 for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.ciscopress.com/register and log in or create an account*. Enter the product ISBN 9781587134340 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

*Be sure to check the box that you would like to hear from us to receive exclusive discounts on future editions of this product.

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:

- **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).
- *Italic* indicates arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([]) indicate an optional element.
- Braces ({ }) indicate a required choice.
- Braces within brackets ([{ }]) indicate a required choice within an optional element.

Introduction

Scaling Networks v6 Companion Guide is the official supplemental textbook for the Cisco Network Academy CCNA Routing & Switching Scaling Networks course. Cisco Networking Academy is a comprehensive program that delivers information technology skills to students around the world. The curriculum emphasizes real-world practical application, while providing opportunities for you to gain the skills and hands-on experience needed to design, install, operate, and maintain networks in small to medium-sized businesses, as well as enterprise and service provider environments.

This book provides a ready reference that explains the same networking concepts, technologies, protocols, and devices as the online curriculum. This book emphasizes key topics, terms, and activities and provides some alternate explanations and examples than are available in the course. You can use the online curriculum as directed by your instructor and then use this book's study tools to help solidify your understanding of all the topics.

Who Should Read This Book

The book, as well as the course, is designed as an introduction to data network technology for those pursuing careers as network professionals as well as those who need only an introduction to network technology for professional growth. Topics are presented concisely, starting with the most fundamental concepts and progressing to a comprehensive understanding of network communication. The content of this text provides the foundation for additional Cisco Networking Academy courses and preparation for the CCNA Routing and Switching certification.

Book Features

The educational features of this book focus on supporting topic coverage, readability, and practice of the course material to facilitate your full understanding of the course material.

Topic Coverage

The following features give you a thorough overview of the topics covered in each chapter so that you can make constructive use of your study time:

- **Objectives:** Listed at the beginning of each chapter, the objectives reference the core concepts covered in the chapter. The objectives match the objectives stated

in the corresponding chapters of the online curriculum; however, the question format in this book encourages you to think about finding the answers as you read the chapter.

- **Notes:** These are short sidebars that point out interesting facts, timesaving methods, and important safety issues.
- **Chapter summaries:** At the end of each chapter is a summary of the chapter's key concepts. It provides a synopsis of the chapter and serves as a study aid.
- **Practice:** At the end of chapter is a full list of the labs, class activities, and Packet Tracer activities to refer to for study time.

Readability

The following features have been updated to assist your understanding of the networking vocabulary:

- **Key terms:** Each chapter begins with a list of key terms, along with a page-number reference from within the chapter. The terms are listed in the order in which they are explained in the chapter. This handy reference allows you to find a term, flip to the page where the term appears, and see the term used in context. The Glossary defines all the key terms.
- **Glossary:** This book contains an all-new Glossary with more than 250 terms.

Practice

Practice makes perfect. This new Companion Guide offers you ample opportunities to put what you learn into practice. You will find the following features valuable and effective in reinforcing the instruction that you receive:

- **“Check Your Understanding” questions and answer key:** Updated review questions are presented at the end of each chapter as a self-assessment tool. These questions match the style of questions that you see in the online course. Appendix A, “Answers to the ‘Check Your Understanding’ Questions,” provides an answer key to all the questions and includes an explanation of each answer.
- **Labs and activities:** Throughout each chapter, you will be directed back to the online course to take advantage of the activities created to reinforce concepts. In addition, at the end of each chapter, there is a practice section that collects a list of all the labs and activities to provide practice with the topics introduced in the chapter. The labs, class activities, and Packet Tracer instructions are available in the companion *Scaling Networks v6 Labs & Study Guide* (ISBN 9781587134333). The Packet Tracer PKA files are found in the online course.



Packet Tracer
Activity

Video

- **Page references to online course:** After headings, you will see, for example, (1.1.2.3). This number refers to the page number in the online course so that you can easily jump to that spot online to view a video, practice an activity, perform a lab, or review a topic.

Lab Study Guide

The supplementary book *Scaling Networks v6 Labs & Study Guide*, by Allan Johnson (ISBN 9781587134333), includes a Study Guide section and a Lab section for each chapter. The Study Guide section offers exercises that help you learn the concepts, configurations, and troubleshooting skill crucial to your success as a CCNA exam candidate. Some chapters include unique Packet Tracer activities available for download from the book's companion website. The Labs and Activities section contains all the labs, class activities, and Packet Tracer instructions from the course.



About Packet Tracer Software and Activities

Interspersed throughout the chapters you'll find many activities to work with the Cisco Packet Tracer tool. Packet Tracer allows you to create networks, visualize how packets flow in the network, and use basic testing tools to determine whether the network would work. When you see this icon, you can use Packet Tracer with the listed file to perform a task suggested in this book. The activity files are available in the course. Packet Tracer software is available only through the Cisco Networking Academy website. Ask your instructor for access to Packet Tracer.

How This Book Is Organized

This book corresponds closely to the Cisco Academy Scaling Networks course and is divided into 10 chapters, one appendix, and a Glossary of key terms:

- **Chapter 1, “LAN Design”:** This chapter discusses strategies that can be used to systematically design a highly functional network, such as the hierarchical network design model and appropriate device selections. The goals of network design are to limit the number of devices impacted by the failure of a single network device, provide a plan and path for growth, and create a reliable network.
- **Chapter 2, “Scaling VLANs”:** This chapter examines the implementation of inter-VLAN routing using a Layer 3 switch. It also describes issues encountered when implementing VTP, DTP and inter-VLAN routing.

- **Chapter 3, “STP”:** This chapter focuses on the protocols used to manage Layer 2 redundancy. It also covers some of the potential redundancy problems and their symptoms.
- **Chapter 4, “EtherChannel and HSRP”:** This chapter describes EtherChannel and the methods used to create an EtherChannel. It also focuses on the operations and configuration of Hot Standby Router Protocol (HSRP), a first-hop redundancy protocol. Finally, the chapter examines a few potential redundancy problems and their symptoms.
- **Chapter 5, “Dynamic Routing”:** This chapter introduces dynamic routing protocols. It explores the benefits of using dynamic routing protocols, how different routing protocols are classified, and the metrics routing protocols use to determine the best path for network traffic. In addition, the characteristics of dynamic routing protocols and the differences between the various routing protocols are examined.
- **Chapter 6, “EIGRP”:** This chapter introduces EIGRP and provides basic configuration commands to enable it on a Cisco IOS router. It also explores the operation of the routing protocol and provides more detail on how EIGRP determines the best path.
- **Chapter 7, “EIGRP Tuning and Troubleshooting”:** This chapter describes EIGRP tuning features, the configuration mode commands to implement these features for both IPv4 and IPv6, and the components and commands used to troubleshoot OSPFv2 and OSPFv3.
- **Chapter 8, “Single-Area OSPF”:** This chapter covers basic single-area OSPF implementations and configurations.
- **Chapter 9, “Multiarea OSPF”:** This chapter discusses basic multiarea OSPF implementations and configurations.
- **Chapter 10, “OSPF Tuning and Troubleshooting”:** This chapter describes OSPF tuning features, the configuration mode commands to implement these features for both IPv4 and IPv6, and the components and commands used to troubleshoot OSPFv2 and OSPFv3.
- **Appendix A, “Answers to the Review Questions”:** This appendix lists the answers to the “Check Your Understanding” review questions that are included at the end of each chapter.
- **Glossary:** The Glossary provides you with definitions for all the key terms identified in each chapter.

Objectives

Upon completion of this chapter, you will be able to answer the following questions:

- What are the appropriate hierarchical network designs for small businesses?
- What are the considerations for designing a scalable network?
- What switch hardware features are appropriate to support network requirements in small to medium-sized business networks?
- What types of routers are available for small to medium-sized business networks?
- What are the basic configuration settings for a Cisco IOS device?

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary.

mission-critical services Page 3

enterprise network Page 3

network operations center (NOC) Page 5

hierarchical design model Page 6

access layer Page 7

distribution layer Page 7

core layer Page 7

collapsed core design Page 7

multilayer switch Page 9

Redundant links Page 9

link aggregation Page 9

redundancy Page 10

Spanning Tree Protocol (STP) Page 10

failure domain Page 11

wireless access point (AP) Page 12

building switch block Page 13

departmental switch block Page 13

EtherChannel Page 13

port channel interface Page 14

load balancing Page 14

Open Shortest Path First (OSPF) Page 15

Enhanced Interior Gateway Routing Protocol (EIGRP) Page 15

link-state routing protocol Page 15

single-area OSPF Page 15

multiarea OSPF Page 15

distance vector routing protocol Page 16

form factor Page 17

Power over Ethernet (PoE) Page 17

campus LAN switch Page 17

cloud-managed switch Page 18

data center switch Page 18

service provider switch Page 18

virtual networking switch Page 18

fixed configuration Page 19

modular configuration Page 19

stackable configuration Page 19

rack unit Page 20

supervisor engine Page 21

port density Page 21

small form-factor pluggable (SFP) Page 22

forwarding rates Page 22

wire speed Page 22

application-specific integrated circuits (ASIC) Page 24

branch router Page 27

network edge router Page 28

service provider router Page 28

Cisco Internetwork Operating System (IOS) Page 29

IOS image Page 30

out-of-band management Page 30

in-band management Page 30

PuTTY Page 31

TeraTerm Page 31

Introduction (1.0.1.1)

There is a tendency to discount a network as just simple plumbing, to think that all you have to consider is the size and the length of the pipes or the speeds and feeds of the links, and to dismiss the rest as unimportant. Just as the plumbing in a large stadium or high rise has to be designed for scale, purpose, redundancy, protection from tampering or denial of operation, and the capacity to handle peak loads, a network requires similar consideration. As users depend on a network to access the majority of the information they need to do their jobs and to transport their voice or video with reliability, the network must be able to provide resilient, intelligent transport.

As a business grows, so does its networking requirements. Businesses rely on the network infrastructure to provide *mission-critical services*. Network outages can result in lost revenue and lost customers. Network designers must design and build an *enterprise network* that is scalable and highly available.

The campus local area network (LAN) is the network that supports devices people use within a location to connect to information. The campus LAN can be a single switch at a small remote site up to a large multi-building infrastructure, supporting classrooms, office space, and similar places where people use their devices. The campus design incorporates both wired and wireless connectivity for a complete network access solution.

This chapter discusses strategies that can be used to systematically design a highly functional network, such as the hierarchical network design model and appropriate device selections. The goals of network design are to limit the number of devices impacted by the failure of a single network device, provide a plan and path for growth, and create a reliable network.



Class Activity 1.0.1.2: Network by Design

Refer to *Scaling Networks v6 Labs & Study Guide* and the online course to complete this activity.

Your employer is opening a new branch office. You have been reassigned to the site as the network administrator, and your job will be to design and maintain the new branch network. The network administrators at the other branches used the Cisco three-layer hierarchical model when designing their networks. You decide to use the same approach. To get an idea of what using the hierarchical model can do to enhance the design process, you research the topic.

Campus Wired LAN Designs (1.1)

Enterprise networks come in all sizes. There are small networks consisting of a few hosts, medium-sized networks consisting of a few hundred hosts, and large networks consisting of thousands of hosts. Besides the number of hosts these networks must support, consideration must be given to the applications and services that must be supported to meet the organizational goals.

Fortunately, proven methods are available to design all types of networks. The Cisco Enterprise Architecture is an example of a proven campus network design.

In this section, you will learn why it is important to design a scalable hierarchical network.

Cisco Validated Designs (1.1.1)

Networks must be scalable, which means they must be able to accommodate an increase or a decrease in size. The focus of this topic is to discover how the hierarchical design model is used to help accomplish this task.

The Need to Scale the Network (1.1.1.1)

Businesses increasingly rely on their network infrastructure to provide mission-critical services. As businesses grow and evolve, they hire more employees, open branch offices, and expand into global markets. These changes directly affect the requirements of a network.

The LAN is the networking infrastructure that provides access to network communication services and resources for end users and devices spread over a single floor or building. A campus network is created by interconnecting a group of LANs that are spread over a small geographic area.

Campus network designs include small networks that use a single LAN switch, up to very large networks with thousands of connections. For example, in Figure 1-1, the company is located in a single location with one connection to the Internet.



Figure 1-1 A Small, Single-Location Company

In Figure 1-2, the company grows to multiple locations in the same city.

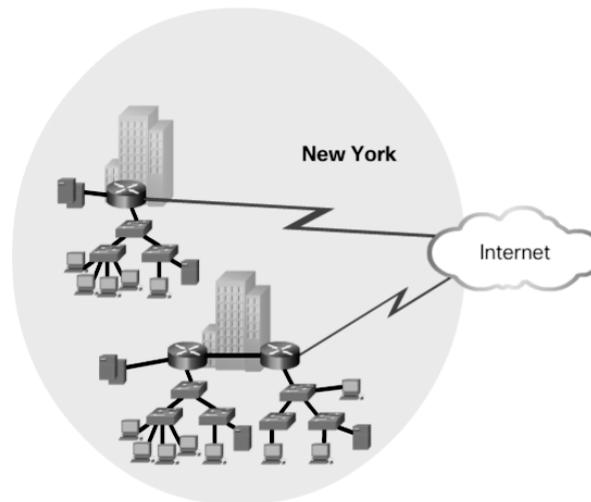


Figure 1-2 The Company Grows to Multiple Locations in the Same City

In Figure 1-3, the company continues to grow and expands to more cities. It also hires and connects teleworkers.

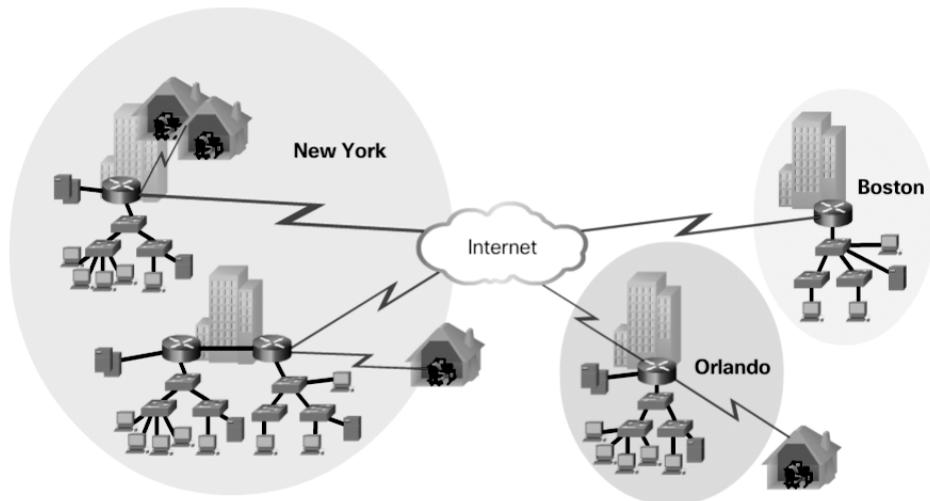


Figure 1-3 Enterprise Grows to Multiple Cities and Adds Teleworkers

In Figure 1-4, the company expands to other countries and centralizes management in a *network operations center (NOC)*.

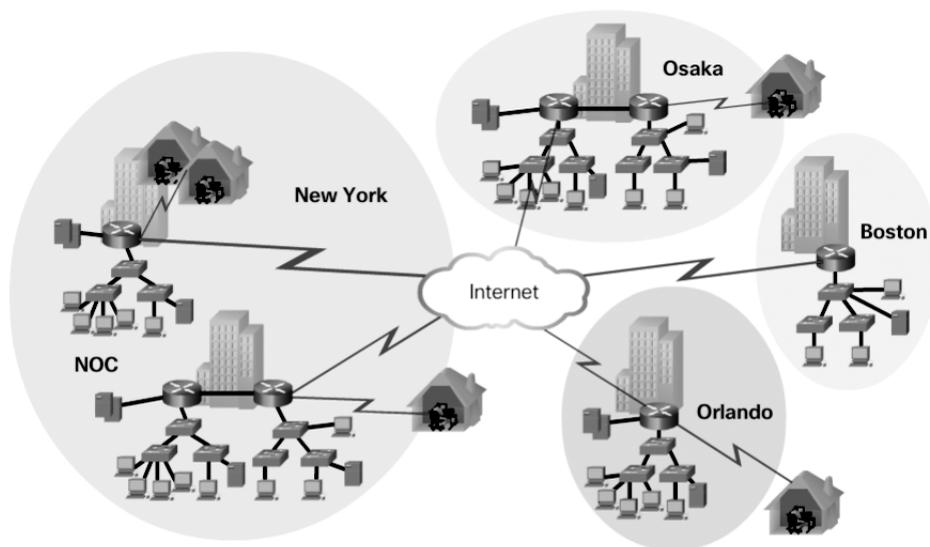


Figure 1-4 Enterprise Becomes Global and Centralizes Network Operations

In addition to supporting physical growth, a network must also support the exchange of all types of network traffic, including data files, email, IP telephony, and video applications for multiple business units.

Specifically, all enterprise networks must:

- Support mission-critical services and applications
- Support converged network traffic
- Support diverse business needs
- Provide centralized administrative control

To help campus LANs meet these requirements, a *hierarchical design model* is used.

Hierarchical Design Model (1.1.1.2)

The campus wired LAN enables communications between devices in a building or group of buildings, as well as interconnection to the WAN and Internet edge at the network core.

Early networks used a flat or meshed network design, in which large numbers of hosts were connected in the same network. Changes affected many hosts in this type of network architecture.

Campus wired LANs now use a hierarchical design model that divides network design into modular groups or layers. Dividing (or *breaking*) the network design into

layers enables each layer to implement specific functions. This simplifies the network design and the deployment and management of the network.

A hierarchical LAN design consists of the following three layers, as shown in Figure 1-5:

- Access layer
- Distribution layer
- Core layer

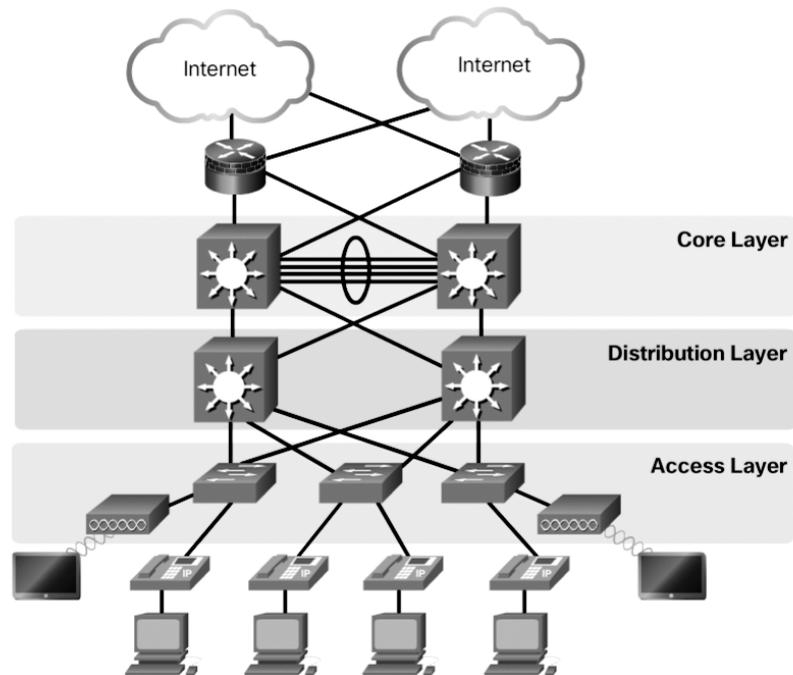


Figure 1-5 Hierarchical Design Model

Each layer is designed to meet specific functions.

The *access layer* provides endpoints and users direct access to the network. The *distribution layer* aggregates access layers and provides connectivity to services. Finally, the *core layer* provides connectivity between distribution layers for large LAN environments. User traffic is initiated at the access layer and passes through the other layers if the functionality of those layers is required.

Medium-sized to large enterprise networks commonly implement the three-layer hierarchical design model. However, some smaller enterprise networks may implement a two-tier hierarchical design, referred to as a *collapsed core design*. In a two-tier hierarchical design, the core and distribution layers are collapsed into one layer, reducing cost and complexity, as shown in Figure 1-6.

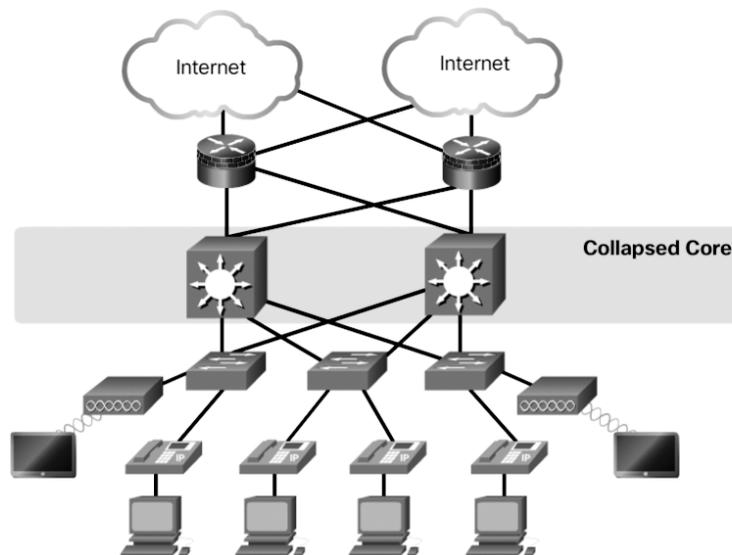


Figure 1-6 Collapsed Core

In flat or meshed network architectures, changes tend to affect a large number of systems. Hierarchical design helps constrain operational changes to a subset of the network, which makes it easy to manage and improves resiliency. Modular structuring of the network into small, easy-to-understand elements also facilitates resiliency via improved fault isolation.

Expanding the Network (1.1.2)

Networks must be scalable, which means they must be able to accommodate an increase or a decrease in size. The focus of this topic is to discover how the hierarchical design model is used to help accomplish this task.

Design for Scalability (1.1.2.1)

To support a large, medium, or small network, the network designer must develop a strategy to enable the network to be available and to scale effectively and easily. Included in a basic network design strategy are the following recommendations:

- Use expandable, modular equipment or clustered devices that can be easily upgraded to increase capabilities. Device modules can be added to the existing equipment to support new features and devices without requiring major equipment upgrades. Some devices can be integrated in a cluster to act as one device to simplify management and configuration.

- Design a hierarchical network to include modules that can be added, upgraded, and modified as necessary, without affecting the design of the other functional areas of the network. For example, you might create a separate access layer that can be expanded without affecting the distribution and core layers of the campus network.
- Create an IPv4 or IPv6 address strategy that is hierarchical. Careful address planning eliminates the need to re-address the network to support additional users and services.
- Use a router or *multilayer switch* to limit broadcasts and filter other undesirable traffic from the network. Use Layer 3 devices to filter and reduce traffic to the network core.

As shown in Figure 1-7, more advanced network design requirements include:

- A. Redundant links**—Implementing redundant links in the network between critical devices and between access layer and core layer devices.

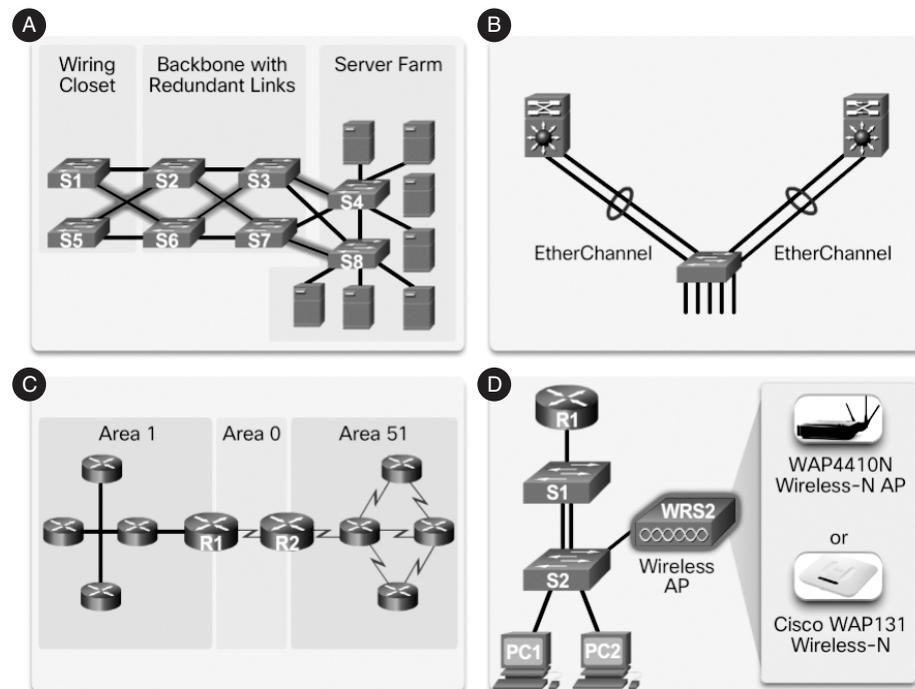


Figure 1-7 Design for Scalability

- B. Link aggregation**—Implementing multiple links between equipment, with either link aggregation (EtherChannel) or equal-cost load balancing, to increase

bandwidth. Combining multiple Ethernet links into a single, load-balanced EtherChannel configuration increases the available bandwidth. EtherChannel implementations can be used when budget restrictions prohibit purchasing high-speed interfaces and fiber runs.

- C. Scalable routing protocols**—Using a scalable routing protocol such as multiarea OSPF and implementing features within that routing protocol to isolate routing updates and minimize the size of the routing table.
- D. Wireless mobility**—Implementing wireless connectivity to allow for mobility and expansion.

Planning for Redundancy (1.1.2.2)

For many organizations, the availability of the network is essential to supporting business needs. *Redundancy* is an important part of network design for preventing disruption of network services by minimizing the possibility of a single point of failure. One method of implementing redundancy is to install duplicate equipment and provide failover services for critical devices.

Another method of implementing redundancy is using redundant paths, as shown in Figure 1-8. Redundant paths offer alternate physical paths for data to traverse the network. Redundant paths in a switched network support high availability. However, due to the operation of switches, redundant paths in a switched Ethernet network may cause logical Layer 2 loops. For this reason, *Spanning Tree Protocol (STP)* is required.

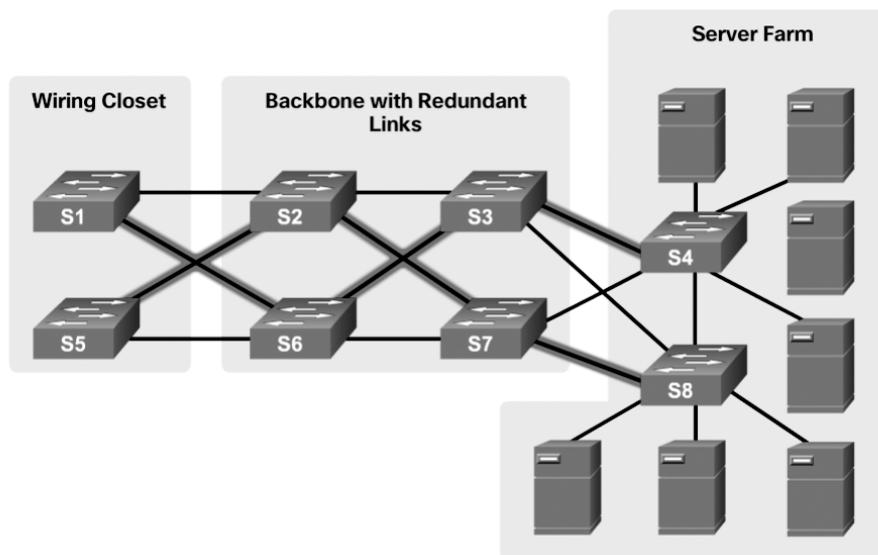


Figure 1-8 LAN Redundancy

STP eliminates Layer 2 loops when redundant links are used between switches. It does this by providing a mechanism for disabling redundant paths in a switched network until the path is necessary, such as when failures occur. STP is an open standard protocol used in a switched environment to create a loop-free logical topology.

Chapter 3, “STP,” provides more details about LAN redundancy and the operation of STP.

Failure Domains (1.1.2.3)

A well-designed network not only controls traffic but also limits the size of failure domains. A *failure domain* is the area of a network that is impacted when a critical device or network service experiences problems.

The function of the device that initially fails determines the impact of a failure domain. For example, a malfunctioning switch on a network segment normally affects only the hosts on that segment. However, if the router that connects this segment to others fails, the impact is much greater.

The use of redundant links and reliable enterprise-class equipment minimizes the chance of disruption in a network. Smaller failure domains reduce the impact of a failure on company productivity. They also simplify the troubleshooting process, thereby shortening the downtime for all users.

Figure 1-9 shows an example of the failure domain for a router.

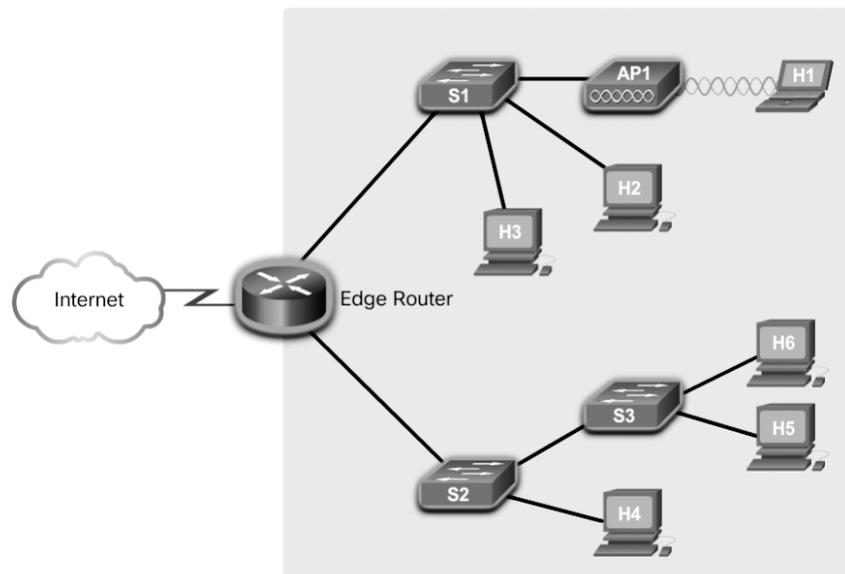


Figure 1-9 Failure Domain—Router

Figure 1-10 shows an example of the failure domain for a switch.

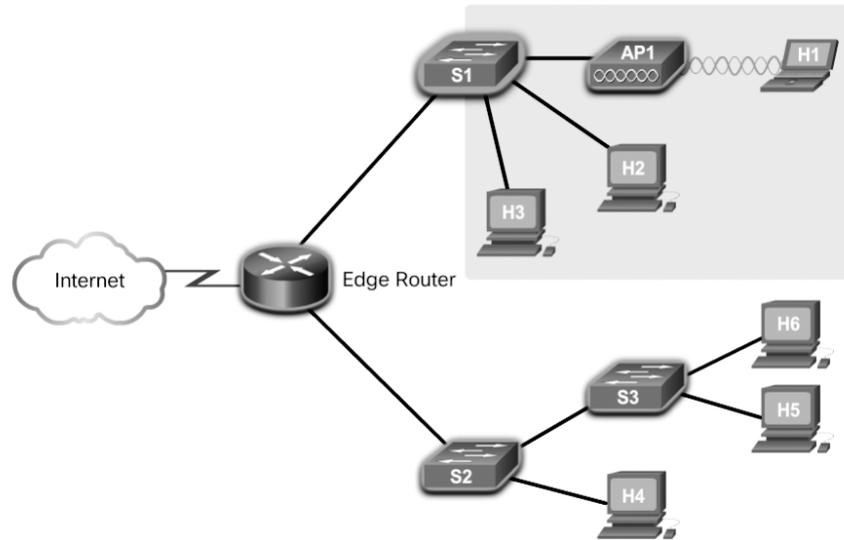


Figure 1-10 Failure Domain—Switch

Figure 1-11 shows an example of the failure domain for a *wireless access point (AP)*.

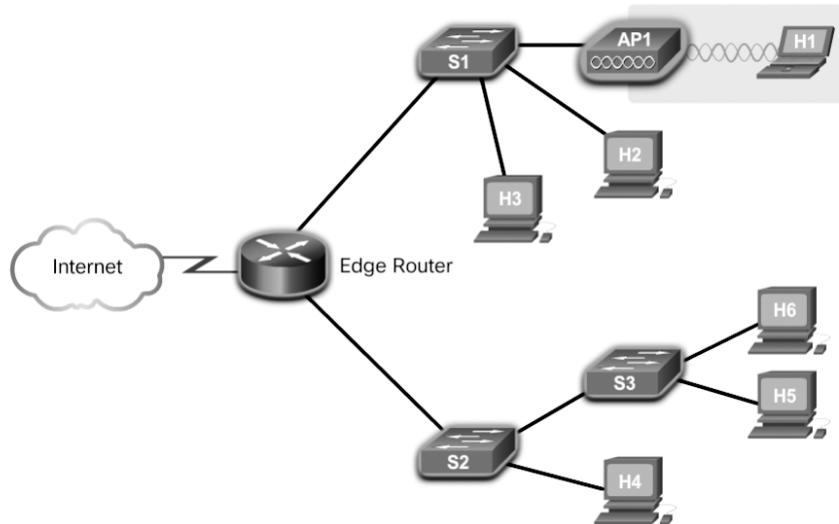


Figure 1-11 Failure Domain—Wireless Access Point

Because a failure at the core layer of a network can have a potentially large impact, the network designer often concentrates on efforts to prevent failures. These efforts can greatly increase the cost of implementing the network.

In the hierarchical design model, it is easiest and usually least expensive to control the size of a failure domain in the distribution layer. Limiting the size of failure domains in the distribution layer confines network errors to a smaller area and thereby affects fewer users. When using Layer 3 devices at the distribution layer, every router functions as a gateway for a limited number of access layer users.

Routers or multilayer switches are usually deployed in pairs, with access layer switches evenly divided between them. This configuration is referred to as a *building switch block* or a *departmental switch block*. Each switch block acts independently of the others. As a result, the failure of a single device does not cause the network to go down. Even the failure of an entire switch block does not affect a significant number of end users.

Increasing Bandwidth (1.1.2.4)

In hierarchical network design, some links between access and distribution switches may need to process a greater amount of traffic than other links. As traffic from multiple links converges onto a single, outgoing link, it is possible for that link to become a bottleneck.

Link aggregation allows an administrator to increase the amount of bandwidth between devices by creating one logical link by grouping several physical links together. *EtherChannel* is a form of link aggregation used in switched networks, as shown in Figure 1-12.

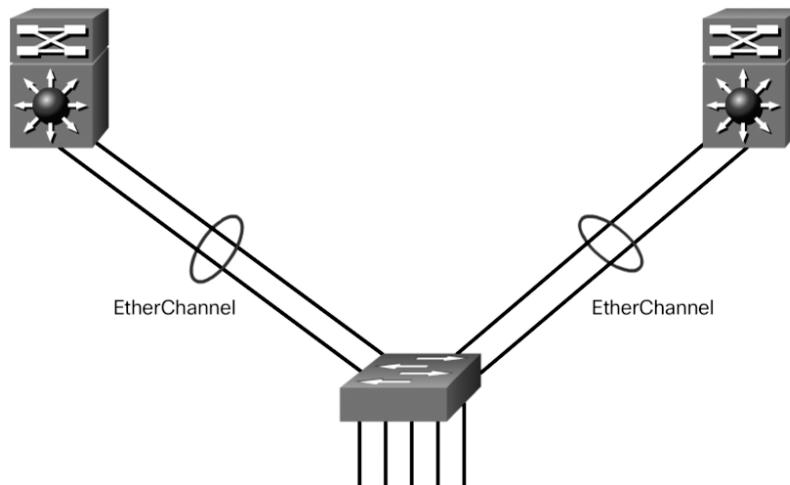


Figure 1-12 Advantages of EtherChannel

EtherChannel uses the existing switch ports. Therefore, additional costs to upgrade the link to a faster and more expensive connection are not necessary. The EtherChannel is seen as one logical link, using an EtherChannel interface.

On a Cisco Catalyst switch, an EtherChannel is configured as a *port channel interface*. Most configuration tasks are done on the port channel interface instead of on each individual port to ensure configuration consistency throughout the links.

Finally, the EtherChannel configuration takes advantage of *load balancing* between links that are part of the same EtherChannel, and depending on the hardware platform, one or more load balancing methods can be implemented.

EtherChannel operation and configuration are covered in more detail Chapter 4, “EtherChannel and HSRP.”

Expanding the Access Layer (1.1.2.5)

A network must be designed to be able to expand network access to individuals and devices as needed. An increasingly important aspect of extending access layer connectivity is wireless connectivity. Providing wireless connectivity offers many advantages, such as increased flexibility, reduced costs, and the ability to grow and adapt to changing network and business requirements.

To communicate wirelessly, end devices require a wireless network interface card (NIC) that incorporates a radio transmitter/receiver and the required software driver to make it operational. In addition, a wireless router or a wireless access point (AP) is required for users to connect, as shown in Figure 1-13.

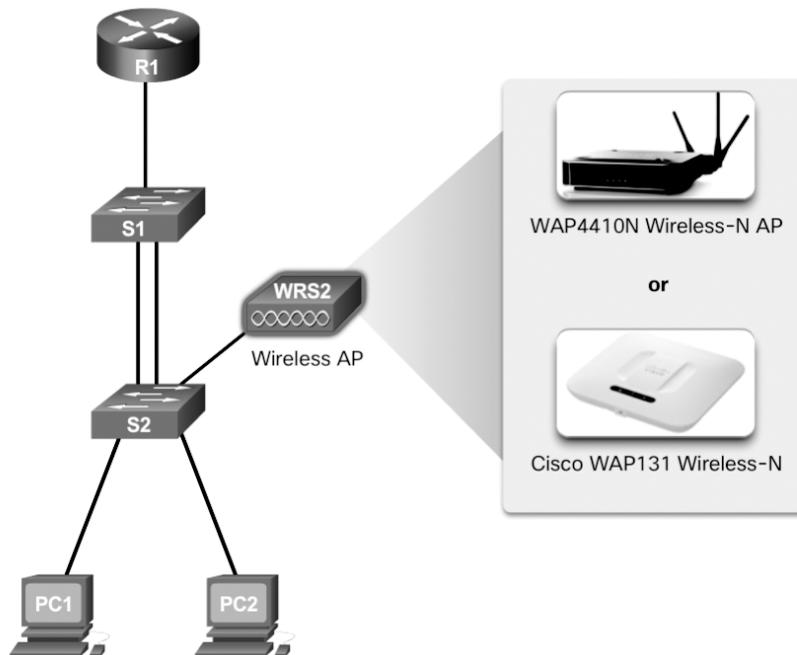


Figure 1-13 Wireless LANs

Implementing a wireless network involves many considerations, such as the types of wireless devices to use, wireless coverage requirements, interference considerations, and security considerations.

Fine-tuning Routing Protocols (1.1.2.6)

Advanced routing protocols, such as *Open Shortest Path First (OSPF)* and *Enhanced Interior Gateway Routing Protocol (EIGRP)*, are used in large networks.

A *link-state routing protocol* such as OSPF, as shown in Figure 1-14, works well for larger hierarchical networks where fast convergence is important.

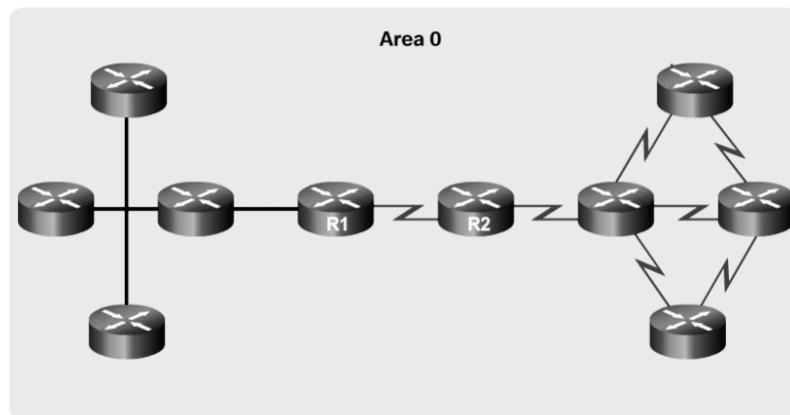


Figure 1-14 Single-Area OSPF

OSPF routers establish and maintain neighbor adjacency or adjacencies with other connected OSPF routers. When routers initiate an adjacency with neighbors, an exchange of link-state updates begins. Routers reach a FULL state of adjacency when they have synchronized views on their link-state database. With OSPF, link-state updates are sent when network changes occur. *Single-area OSPF* configuration and concepts are covered in Chapter 8, “Single-Area OSPF.”

In addition, OSPF supports a two-layer hierarchical design, referred to as *multiarea OSPF*, as shown in Figure 1-15.

All multiarea OSPF networks must have an Area 0, also called the backbone area. Non-backbone areas must be directly connected to area 0. Chapter 9, “Multiarea OSPF,” introduces the benefits, operation, and configuration of multiarea OSPF. Chapter 10, “OSPF Tuning and Troubleshooting,” covers more advanced features of OSPF.

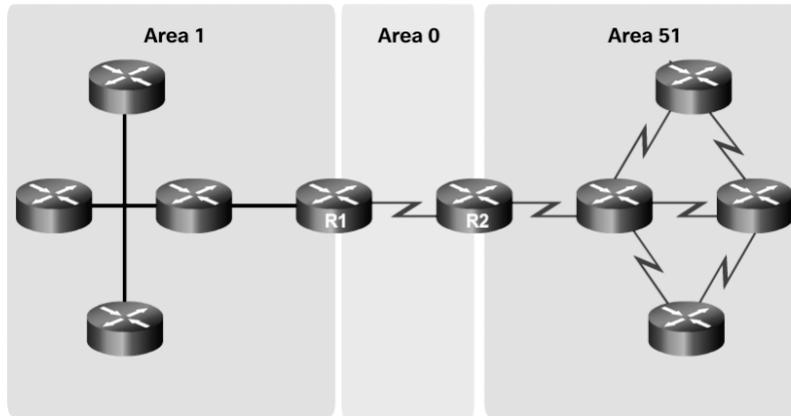


Figure 1-15 Multiarea OSPF

Another popular routing protocol for larger networks is EIGRP. Cisco developed EIGRP as a proprietary *distance vector routing protocol* with enhanced capabilities. Although configuring EIGRP is relatively simple, the underlying features and options of EIGRP are extensive and robust. For example, EIGRP uses protocol-dependent modules (PDM), which enable support for IPv4 and IPv6 routing tables, as shown in Figure 1-16.

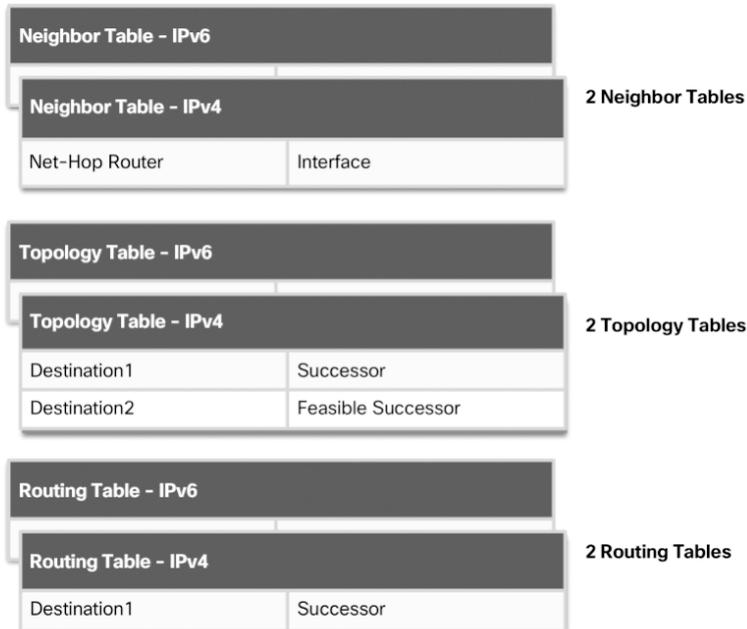


Figure 1-16 EIGRP Protocol-Dependent Modules (PDM)

EIGRP contains many features that are not found in any other routing protocols. It is an excellent choice for large multiprotocol networks that use primarily Cisco devices.

Chapter 6, “EIGRP,” introduces the operation and configuration of the EIGRP routing protocol, and Chapter 7, “EIGRP Tuning and Troubleshooting,” covers some of the more advanced configuration options of EIGRP.

**Interactive
Graphic****Activity 1.1.2.7: Identify Scalability Terminology**

Refer to the online course to complete this activity.

Selecting Network Devices (1.2)

Switches and routers are core network infrastructure devices. Therefore, selecting them appears to be a fairly simple task. However, many different models of switches and routers are available. Different models provide various numbers of ports, different forwarding rates, and unique feature support.

In this section, you will learn how to select network devices based on feature compatibility and network requirements.

Switch Hardware (1.2.1)

Various types of switch platforms are available. Each platform differs in terms of physical configuration and *form factor*, the number of ports, and the features supported, including *Power over Ethernet (PoE)* and routing protocols.

The focus of this topic is on how to select the appropriate switch hardware features to support network requirements in small to medium-sized business networks.

Switch Platforms (1.2.1.1)

When designing a network, it is important to select the proper hardware to meet current network requirements, as well as allow for network growth. Within an enterprise network, both switches and routers play a critical role in network communication.

There are five categories of switches for enterprise networks, as shown in Figure 1-17:

- *Campus LAN switch*—To scale network performance in an enterprise LAN, there are core, distribution, access, and compact switches. These switch platforms vary from fanless switches with eight fixed ports to 13-blade switches supporting hundreds of ports. Campus LAN switch platforms include the Cisco 2960, 3560, 3650, 3850, 4500, 6500, and 6800 Series.

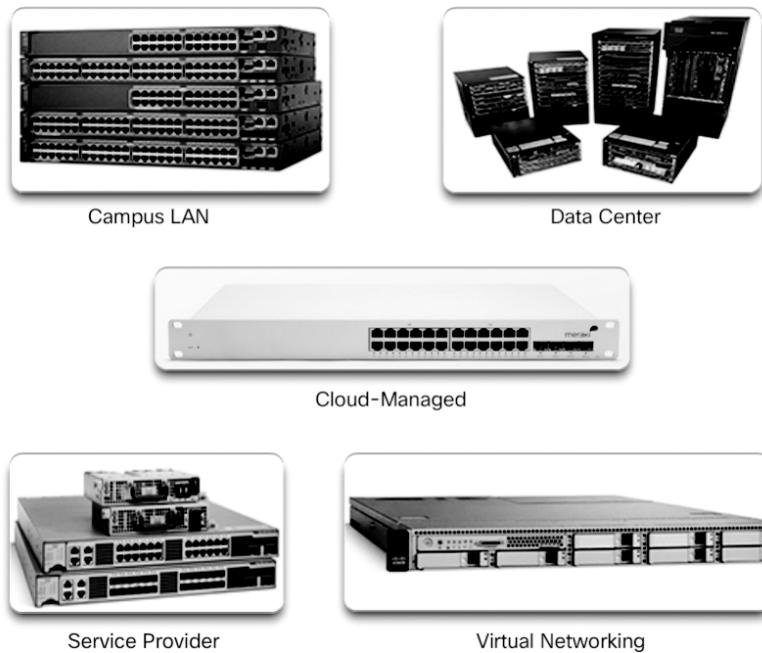
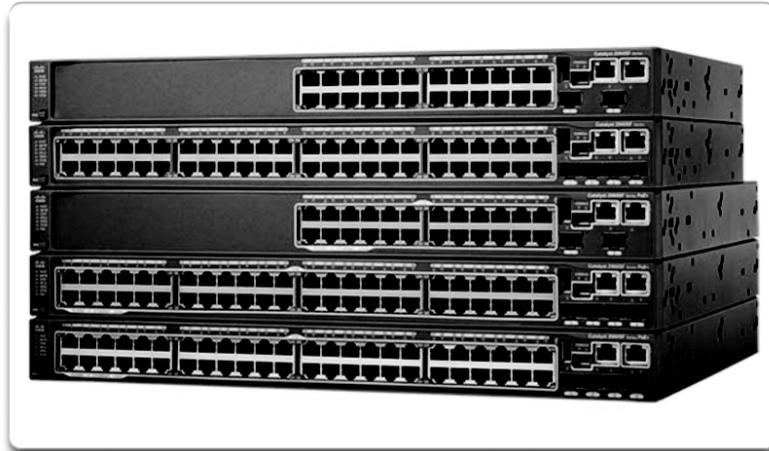


Figure 1-17 Switch Platforms

- **Cloud-managed switch**—The Cisco Meraki cloud-managed access switches enable virtual stacking of switches. They monitor and configure thousands of switch ports over the web, without the intervention of onsite IT staff.
- **Data center switch**—A data center should be built based on switches that promote infrastructure scalability, operational continuity, and transport flexibility. The data center switch platforms include the Cisco Nexus Series switches and the Cisco Catalyst 6500 Series switches.
- **Service provider switch**—Service provider switches fall under two categories: aggregation switches and Ethernet access switches. Aggregation switches are carrier-grade Ethernet switches that aggregate traffic at the edge of a network. Service provider Ethernet access switches feature application intelligence, unified services, virtualization, integrated security, and simplified management.
- **Virtual networking switch**—Networks are becoming increasingly virtualized. Cisco Nexus virtual networking switch platforms provide secure multitenant services by adding virtualization intelligence technology to the data center network.

When selecting switches, network administrators must determine the switch form factors. These include *fixed configuration* (Figure 1-18), *modular configuration* (Figure 1-19), or *stackable configuration* (Figure 1-20).



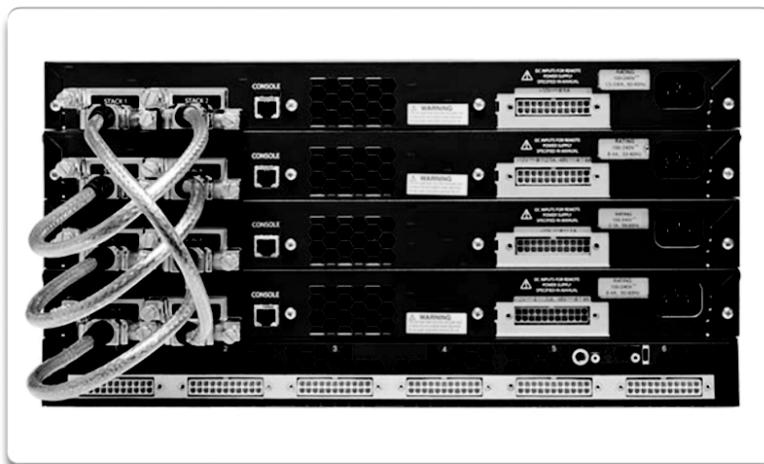
Features and options are limited to those that originally come with the switch.

Figure 1-18 Fixed Configuration Switches



The chassis accepts line cards that contain the ports.

Figure 1-19 Modular Configuration Switches



Stackable switches, connected by a special cable, effectively operate as one large switch.

Figure 1-20 Stackable Configuration Switches

The amount of space that a device occupies in a network rack is also an important consideration. *Rack unit* is a term used to describe the thickness of a rack-mountable network device. Defined in EIA-310, a unit (U) describes a device with a standard height of 4.45 centimeters (1 3/4 inches) and width of 48.26 centimeters (19 inches). For example, the fixed configuration switches shown in Figure 1-18 are all one rack unit (1U).

Besides the device form factor, other device selection considerations must be made. Table 1-1 describes some of these considerations.

Table 1-1 Considerations When Selecting Network Devices

Consideration	Description
Cost	The cost of a switch depends on the number and speed of the interfaces, supported features, and expansion capability.
Port density	The port density describes how many ports are available on the switch. Network switches must support the appropriate number of devices on the network.
Port speed	The speed of the network connection is of primary concern to end users.
Forwarding rate	This rate defines the processing capabilities of a switch by rating how much data the switch can process per second. For instance, distribution layer switches should provide higher forwarding rates than access layer switches.

Consideration	Description
Size of frame buffers	Switches with large frame buffers are better able to store frames when there are congested ports to servers or other areas of the network.
PoE support	Power over Ethernet (PoE) is used to power access points, IP phones, security cameras, and even compact switches. Demand for PoE is increasing.
Redundant power	Some stackable and modular chassis-based switches support redundant power supplies.
Reliability	Switches should provide continuous access to the network. Therefore, select switches with reliable redundant features including redundant power supplies, fans, and <i>supervisor engines</i> .
Scalability	The number of users on a network typically grows over time. Therefore, select switches that provide the opportunity for growth.

Some of these considerations are now described in more detail.

Port Density (1.2.1.2)

The *port density* of a switch refers to the number of ports available on a single switch. Figure 1-21 shows the port densities of three different switches.

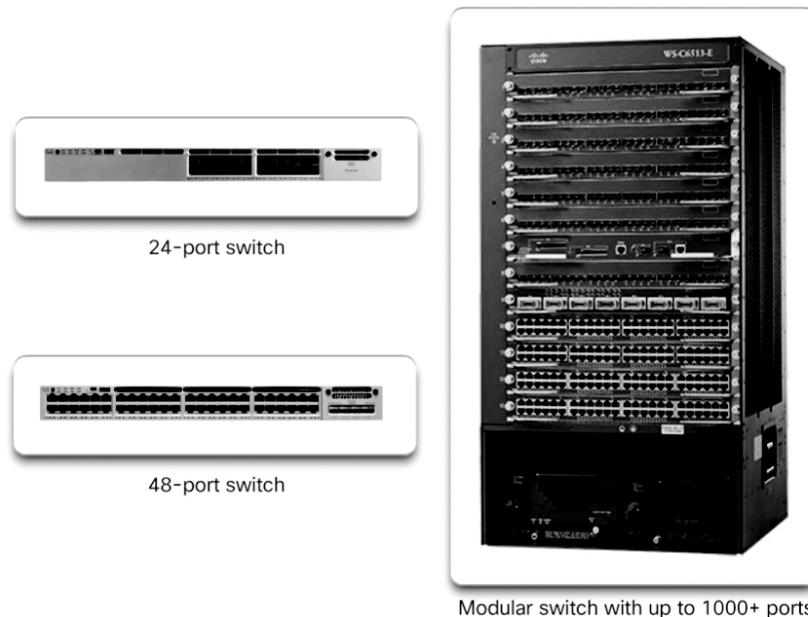


Figure 1-21 Port Densities

Fixed configuration switches support a variety of port density configurations. The Cisco Catalyst 3850 24-port and 48-port switches are shown on the left in the figure. The 48-port switch has an option for 4 additional ports for *small form-factor pluggable (SFP)* devices. SFPs are small compact, hot-pluggable transceivers used on some switches to provide flexibility when choosing network media. SFP transceivers are available for copper and fiber Ethernet, Fibre Channel networks, and more.

Modular switches can support very high port densities through the addition of multiple switch port line cards. The modular Catalyst 6500 switch shown on the right in the figure can support in excess of 1000 switch ports.

Large networks that support many thousands of network devices require high-density modular switches to make the best use of space and power. Without high-density modular switches, a network would need many fixed configuration switches to accommodate the number of devices that need network access—and this approach can consume many power outlets and a lot of closet space.

A network designer must also consider the issue of uplink bottlenecks: A series of fixed configuration switches may consume many additional ports for bandwidth aggregation between switches, for the purpose of achieving target performance. With a single modular switch, bandwidth aggregation is less problematic because the backplane of the chassis can provide the necessary bandwidth to accommodate the devices connected to the switch port line cards.

Forwarding Rates (1.2.1.3)

Forwarding rates define the processing capabilities of a switch by rating how much data the switch can process per second. Switch product lines are classified by forwarding rates, as shown in Figure 1-22.

Forwarding rates are an important consideration when selecting a switch. If its forwarding rate is too low, a switch cannot accommodate full wire-speed communication across all of its switch ports. *Wire speed* is a term used to describe the data rate that each Ethernet port on the switch is capable of attaining. Data rates can be 100 Mb/s, 1 Gb/s, 10 Gb/s, or 100 Gb/s.

For example, a typical 48-port gigabit switch operating at full wire speed generates 48 Gb/s of traffic. If the switch supports a forwarding rate of only 32 Gb/s, it cannot run at full wire speed across all ports simultaneously.

Access layer switches are usually physically limited by their uplinks to the distribution layer. However, they typically do not need to operate at full wire speed. Therefore, less expensive, lower-performing switches can be used at the access layer. The more expensive, higher-performing switches can be used at the distribution and core layers, where the forwarding rate has a greater impact on network performance.

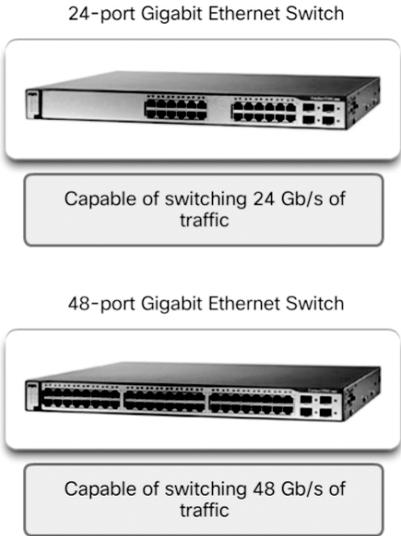


Figure 1-22 Forwarding Rate

Power over Ethernet (1.2.1.4)

PoE allows a switch to deliver power to a device over the existing Ethernet cabling. This feature can be used by IP phones and some wireless access points. Figure 1-23 shows PoE ports on various devices.

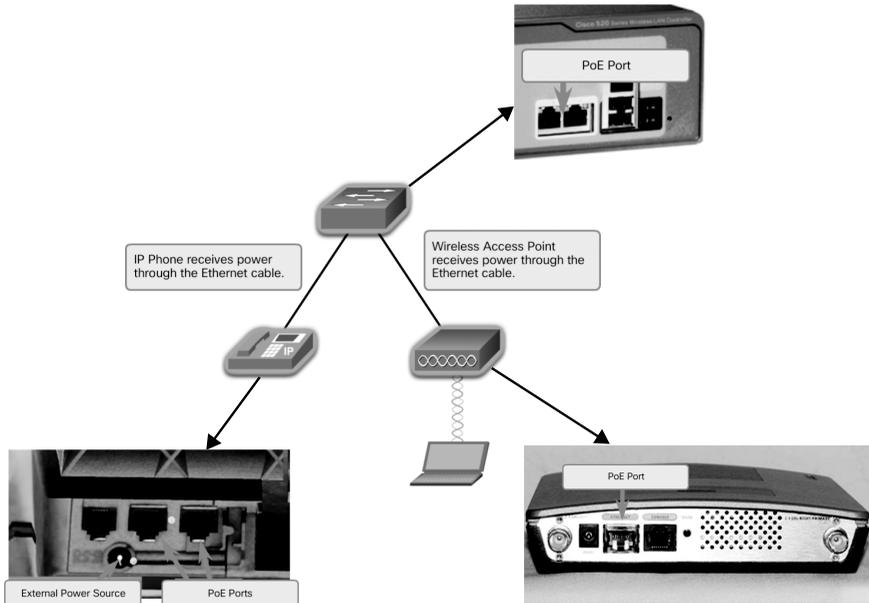


Figure 1-23 Power over Ethernet

PoE increases flexibility when installing wireless access points and IP phones because these devices can be installed anywhere that there is an Ethernet cable. Therefore, a network administrator should ensure that the PoE features are required because switches that support PoE are expensive.

The Cisco Catalyst 2960-C and 3560-C Series compact switches support PoE pass-through. PoE pass-through allows a network administrator to power PoE devices connected to the switch, as well as the switch itself, by drawing power from certain upstream switches. Figure 1-24 shows the PoE ports on a Cisco Catalyst 2960-C.

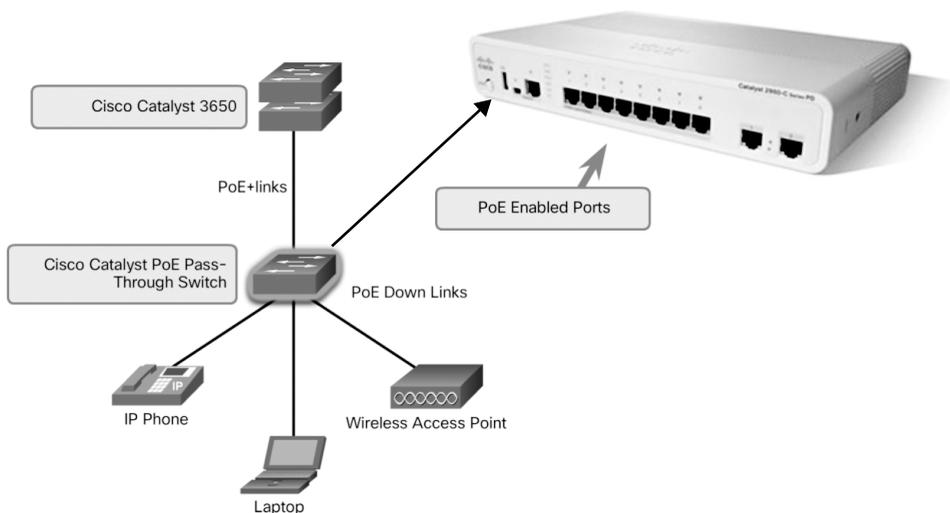


Figure 1-24 PoE Pass-through

Multilayer Switching (1.2.1.5)

Multilayer switches are typically deployed in the core and distribution layers of an organization's switched network. Multilayer switches are characterized by their capability to build a routing table, support a few routing protocols, and forward IP packets at a rate close to that of Layer 2 forwarding. Multilayer switches often support specialized hardware, such as *application-specific integrated circuits (ASIC)*. ASICs along with dedicated software data structures can streamline the forwarding of IP packets independently of the CPU.

There is a trend in networking toward a pure Layer 3 switched environment. When switches were first used in networks, none of them supported routing; now, almost all switches support routing. It is likely that soon all switches will incorporate a route processor because the cost is decreasing relative to other constraints.

As shown in Figure 1-25, the Catalyst 2960 switches illustrate the migration to a pure Layer 3 environment. With IOS versions prior to 15.x, these switches supported only one active switched virtual interface (SVI). With IOS 15.x, these switches now support multiple active SVIs. This means that a Catalyst 2960 switch can be remotely accessed via multiple IP addresses on distinct networks.

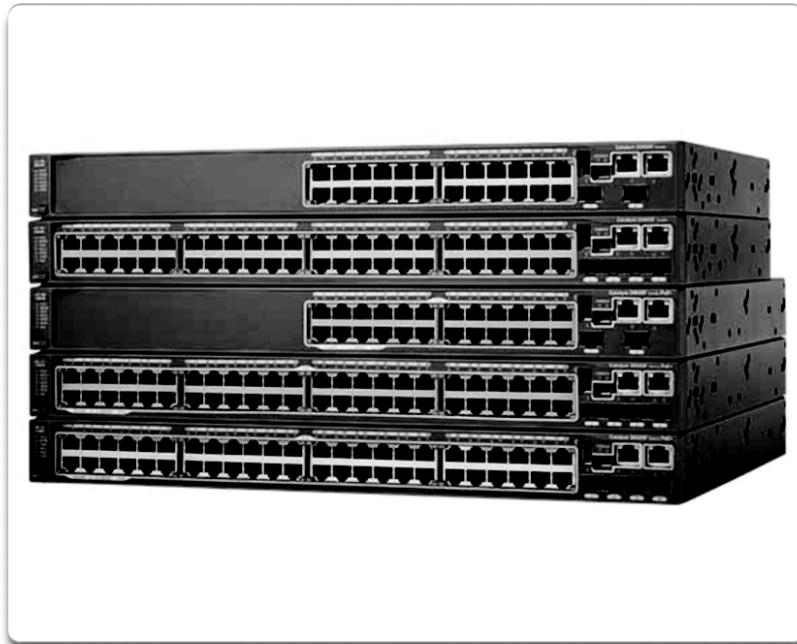


Figure 1-25 Cisco Catalyst 2960 Series Switches

**Interactive
Graphic**

Activity 1.2.1.6: Selecting Switch Hardware

Refer to the online course to complete this activity.

**Packet Tracer
Activity**

Packet Tracer 1.2.1.7: Comparing 2960 and 3560 Switches

In this activity, you will use various commands to examine three different switching topologies and compare the similarities and differences between the 2960 and 3560 switches. You will also compare the routing table of a 1941 router with a 3560 switch.

Router Hardware (1.2.2)

Various types of router platforms are available. Like switches, routers differ in physical configuration and form factor, the number and types of interfaces supported, and the features supported.

The focus of this topic is on how to describe the types of routers available to support network requirements in small to medium-sized business networks.

Router Requirements (1.2.2.1)

In the distribution layer of an enterprise network, routing is required. Without the routing process, packets cannot leave the local network.

Routers play a critical role in networking by determining the best path for sending packets. They connect multiple IP networks by connecting homes and businesses to the Internet. They are also used to interconnect multiple sites within an enterprise network, providing redundant paths to destinations. A router can also act as a translator between different media types and protocols. For example, a router can accept packets from an Ethernet network and re-encapsulate them for transport over a serial network.

Routers use the network portion of the destination IP address to route packets to the proper destination. They select an alternate path if a link or path goes down. All hosts on a local network specify the IP address of the local router interface in their IP configuration. This router interface is the default gateway. The ability to route efficiently and recover from network link failures is critical to delivering packets to their destination.

Routers also serve other beneficial functions, as shown in Figure 1-26:

- Provide broadcast containment
- Provide enhanced security
- Connect remote locations
- Group users logically by application or department

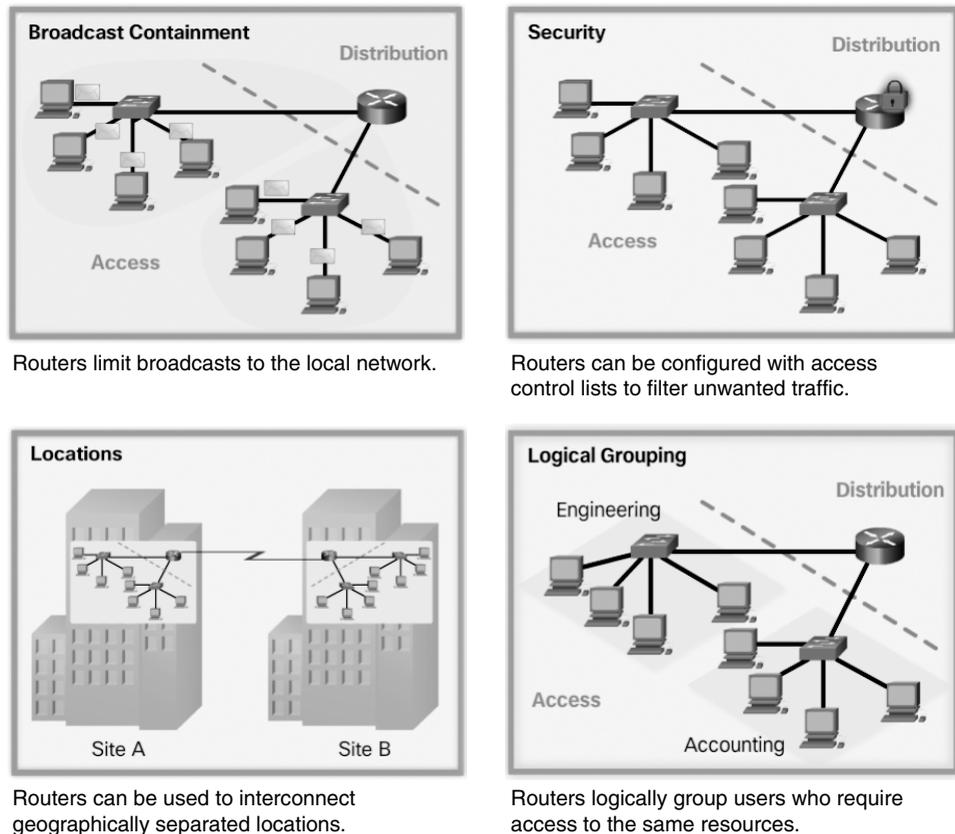


Figure 1-26 Router Functions

Cisco Routers (1.2.2.2)

As a network grows, it is important to select the proper routers to meet its requirements. As shown Figure 1-27, there are three categories of routers:

- **Branch router**—Branch routers optimize branch services on a single platform while delivering an optimal application experience across branch and WAN infrastructures. Maximizing service availability at the branch requires networks designed for 24x7x365 uptime. Highly available branch networks must ensure fast recovery from typical faults while minimizing or eliminating the impact on service, and they must provide simple network configuration and management.

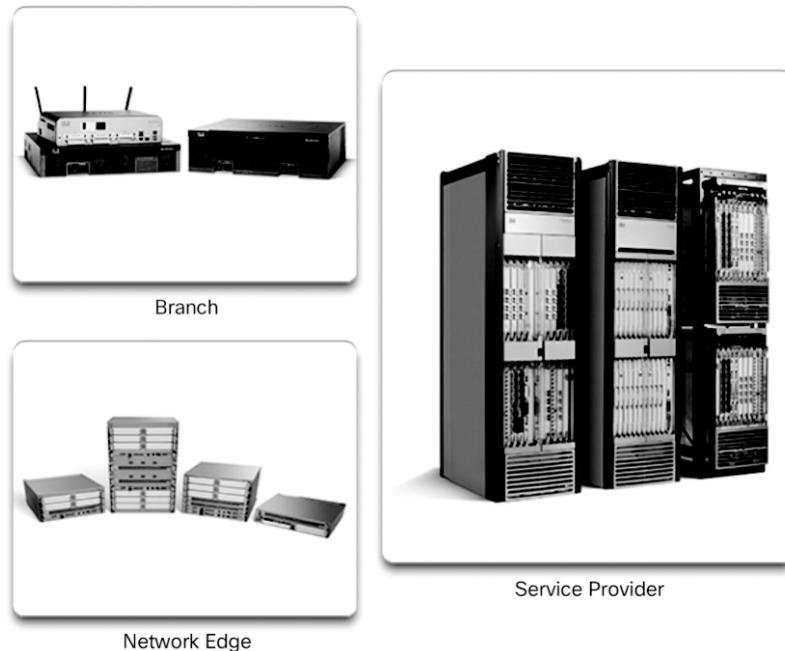


Figure 1-27 Router Platforms

- **Network edge router**—Network edge routers enable the network edge to deliver high-performance, highly secure, and reliable services that unite campus, data center, and branch networks. Customers expect a high-quality media experience and more types of content than ever before. Customers want interactivity, personalization, mobility, and control for all content. Customers also want to access content anytime and anyplace they choose, over any device—whether at home, at work, or on the go. Network edge routers must deliver enhanced quality of service and nonstop video and mobile capabilities.
- **Service provider router**—Service provider routers differentiate the service portfolio and increase revenues by delivering end-to-end scalable solutions and subscriber-aware services. Operators must optimize operations, reduce expenses, and improve scalability and flexibility to deliver next-generation Internet experiences across all devices and locations. These systems are designed to simplify and enhance the operation and deployment of service-delivery networks.

Router Hardware (1.2.2.3)

Routers are available in many form factors, as shown in Figure 1-28. Network administrators in an enterprise environment should be able to support a variety of routers, from a small desktop router to a rack-mounted or blade model.

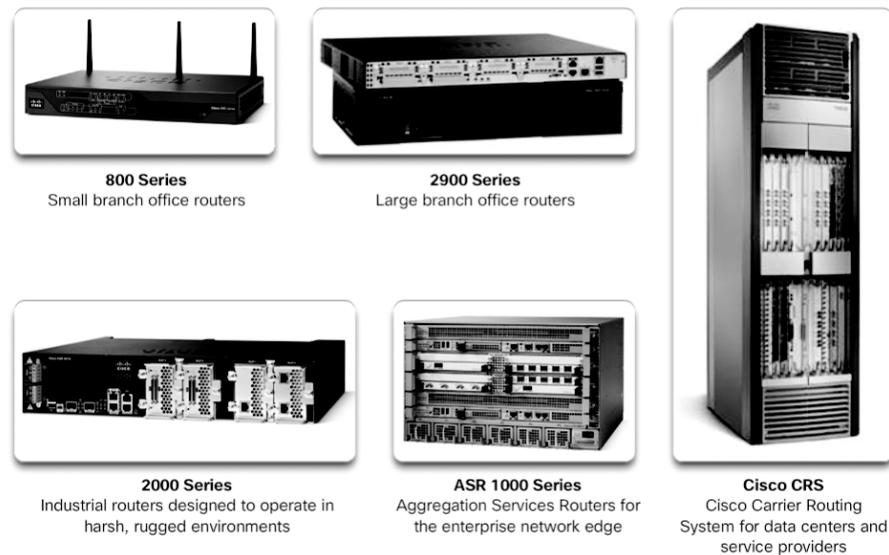


Figure 1-28 A Sampling of Cisco Routers

Routers can also be categorized as fixed configuration or modular. With the fixed configuration, the desired router interfaces are built in. Modular routers come with multiple slots that allow a network administrator to change the interfaces on the router. For example, a Cisco 1941 router is a small modular router. It comes with two built-in Gigabit Ethernet RJ-45 interfaces, and it also has two slots that can accommodate many different network interface modules. Routers come with a variety of different interfaces, such as Fast Ethernet, Gigabit Ethernet, serial, and fiber-optic.

Visit www.cisco.com/c/en/us/products/routers/product-listing.html for a comprehensive list of Cisco routers.

**Interactive
Graphic**

Activity 1.2.2.4: Identify the Router Category

Refer to the online course to complete this activity.

Managing Devices (1.2.3)

Regardless of the form factor and the features each IOS device supports, it requires the *Cisco Internetwork Operating System (IOS)* to be operational.

The focus of this topic is on the Cisco IOS, how to manage it, and how to configure basic settings on Cisco IOS routers and switches.

Managing IOS Files and Licensing (1.2.3.1)

With such a wide selection of network devices to choose from in the Cisco product line, an organization can carefully determine the ideal combination to meet the needs of employees and customers.

When selecting or upgrading a Cisco IOS device, it is important to choose the proper *IOS image* with the correct feature set and version. The IOS image refers to the package of routing, switching, security, and other internetworking technologies integrated into a single multitasking operating system. When a new device is shipped, it comes preinstalled with the software image and the corresponding permanent licenses for the customer-specified packages and features.

For routers, beginning with Cisco IOS Software Release 15.0, Cisco modified the process to enable new technologies within the IOS feature sets, as shown in Figure 1-29.

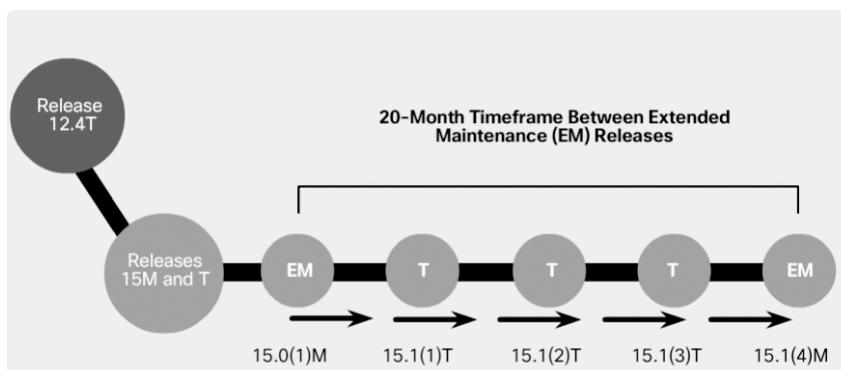


Figure 1-29 Cisco IOS Software Release 15 Family

In this figure, EM (or Extended Maintenance) releases are released approximately every 16 to 20 months. The T releases are between EM releases and are ideal for the very latest features and hardware support before the next EM release becomes available.

In-Band versus Out-of-Band Management (1.2.3.2)

Regardless of the Cisco IOS network device being implemented, there are two methods for connecting a PC to that network device for configuration and monitoring tasks: *out-of-band management* and *in-band management* (see Figure 1-30).

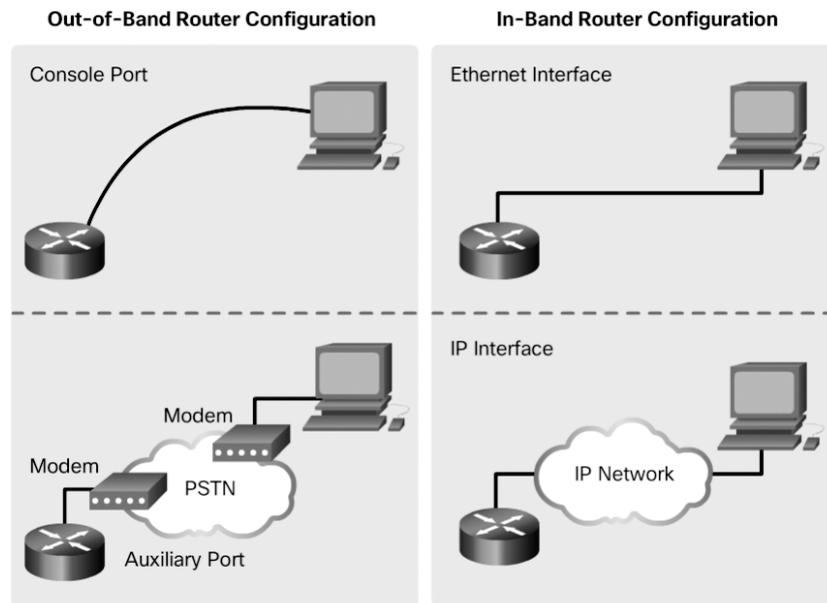


Figure 1-30 In-Band versus Out-of-Band Configuration Options

Out-of-band management is used for initial configuration or when a network connection is unavailable. Configuration using out-of-band management requires:

- A direct connection to a console or an AUX port
- A terminal emulation client (such as *PuTTY* or *TeraTerm*)

In-band management is used to monitor and make configuration changes to a network device over a network connection. Configuration using in-band management requires:

- At least one network interface on the device to be connected and operational
- Telnet, SSH, HTTP, or HTTPS to access a Cisco device

Note

Telnet and HTTP are less secure than the others listed here and are not recommended.

Basic Router CLI Commands (1.2.3.3)

A basic router configuration includes the host name for identification, passwords for security, assignment of IP addresses to interfaces for connectivity, and basic routing.

Example 1-1 shows the commands entered to enable a router with RIPv2. Verify and save configuration changes by using the **copy running-config startup-config** command.

Example 1-1 Enabling a Router with RIPv2

```
Router# configure terminal
Router(config)# hostname R1
R1(config)# enable secret class
R1(config)# line con 0
R1(config-line)# password cisco
R1(config-line)# login
R1(config-line)# exec-timeout 0 0
R1(config-line)# line vty 0 4
R1(config-line)# password cisco
R1(config-line)# login
R1(config-line)# exit
R1(config)# service password-encryption
R1(config)# banner motd $ Authorized Access Only! $
R1(config)#
R1(config)# interface GigabitEthernet0/0
R1(config-if)# description Link to LAN 1
R1(config-if)# ip address 172.16.1.1 255.255.255.0
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)#
R1(config)# interface Serial0/0/0
R1(config-if)# description Link to R2
R1(config-if)# ip address 172.16.3.1 255.255.255.252
R1(config-if)# clock rate 128000
R1(config-if)# no shutdown
R1(config-if)# interface Serial0/0/1
R1(config-if)# description Link to R3
R1(config-if)# ip address 192.168.10.5 255.255.255.252
R1(config-if)# no shutdown
R1(config-if)# exit
R1(config)#
R1(config)# router rip
R1(config-router)# version 2
R1(config-router)# network 172.16.0.0
R1(config-router)# network 192.168.10.0
R1(config-router)# end
R1#
R1# copy running-config startup-config
```

Example 1-2 shows the results of the configuration commands entered in Example 1-1. To clear the router configuration, use the **erase startup-config** command and then the **reload** command.

Example 1-2 Router Running Configuration

```
R1# show running-config
Building configuration...

Current configuration : 1242 bytes
!
Version 15.1
Service timestamps debug datetime msec
Service timestamps log datetime msec
Service password-encryption
!
hostname R1
!
enable secret class
!
<output omitted>
!
interface GigabitEthernet0/0
  description Link to LAN 1
  ip address 172.16.1.1 255.255.255.0
  no shutdown
!
interface Serial0/0/0
  description Link to R2
  ip address 172.16.3.1 255.255.255.252
  clock rate 128000
  no shutdown
!
interface Serial0/0/1
  description Link to R3
  ip address 192.168.10.5 255.255.255.252
  no shutdown
!
router rip
  version 2
  network 172.16.1.0
  network 192.168.10.0
!
banner motd ^C Authorized Access Only! ^C
!
line console 0
  password cisco
  login
```

```
exec-timeout 0 0
line aux 0
line vty 0 4
 password cisco
 login
```

Basic Router Show Commands (1.2.3.4)

A variety of IOS commands are commonly used to display and verify the operational status of the router and related IPv4 network functionality. Similar commands are available for IPv6; they replace **ip** with **ipv6**.

The following list describes routing-related and interface-related IOS router commands:

- **show ip protocols**—Displays information about the routing protocols configured. If RIP is configured, this includes the version of RIP, networks the router is advertising, whether automatic summarization is in effect, the neighbors the router is receiving updates from, and the default administrative distance, which is 120 for RIP (see Example 1-3).

Example 1-3 The **show ip protocols** Command

```
R1# show ip protocols

Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 26 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Default version control: send version 2, receive version 2
    Interface          Send  Recv  Triggered RIP  Key-chain
  GigabitEthernet0/0    2    2
  Serial0/0/0          2    2
  Serial0/0/1          2    2
  Interface          Send  Recv  Triggered RIP  Key-chain
Automatic network summarization is in effect
Maximum path: 4
Routing for Networks:
  172.16.0.0
  192.168.10.0
Routing Information Sources:
  Gateway          Distance    Last Update
  172.16.3.2        120        00:00:25
Distance: (default is 120)
```

- **show ip route**—Displays routing table information, including routing codes, known networks, administrative distance and metrics, how routes were learned, next hop, static routes, and default routes (see Example 1-4).

Example 1-4 The show ip route Command

```
R1# show ip route | begin Gateway
Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 5 subnets, 3 masks
C       172.16.1.0/24 is directly connected, GigabitEthernet0/0
L       172.16.1.1/32 is directly connected, GigabitEthernet0/0
C       172.16.3.0/30 is directly connected, Serial0/0/0
L       172.16.3.1/32 is directly connected, Serial0/0/0
R       172.16.5.0/24 [120/1] via 172.16.3.2, 00:00:25, Serial0/0/0
      192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.10.4/30 is directly connected, Serial0/0/1
L       192.168.10.5/32 is directly connected, Serial0/0/1
```

- **show interfaces**—Displays interface information and status, including the line (protocol) status, bandwidth, delay, reliability, encapsulation, duplex, and I/O statistics. If specified without a specific interface designation, all interfaces are displayed. If a specific interface is specified after the command, information about that interface only is displayed (see Example 1-5).

Example 1-5 The show interfaces Command

```
R1# show interfaces gigabitethernet 0/0
GigabitEthernet0/0 is up, line protocol is up (connected)
  Hardware is CN Gigabit Ethernet, address is 00e0.8fb2.de01 (bia 00e0.8fb2.de01)
  Description: Link to LAN 1
  Internet address is 172.16.1.1/24
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full Duplex, 100Mbps, media type is RJ45
<output omitted>
Serial0/0/0 is up, line protocol is up (connected)
  Hardware is HD64570
  Description: Link to R2
```

```

Internet address is 172.16.3.1/30
MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
Encapsulation HDLC, loopback not set, keepalive set (10 sec)
Last input never, output never, output hang never
Last clearing of "show interface" counters never
<output omitted>
Serial0/0/1 is up, line protocol is up (connected)
  Hardware is HD64570
  Description: Link to R3
  Internet address is 192.168.10.5/30
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation HDLC, loopback not set, keepalive set (10 sec)
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never

```

- **show ip interfaces**—Displays IP-related interface information, including protocol status, the IPv4 address, whether a helper address is configured, and whether an ACL is enabled on the interface. If specified without a specific interface designation, all interfaces are displayed. If a specific interface is specified after the command, information about that interface only is displayed (see Example 1-6).

Example 1-6 The show ip interface Command

```

R1# show ip interface gigabitEthernet 0/0
GigabitEthernet0/0 is up, line protocol is up
  Internet address is 172.16.1.1/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.5 224.0.0.6
  Outgoing access list is not set
  Inbound access list is not set
  Proxy ARP is enabled
  Local Proxy ARP is disabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachable are always sent

```

```

ICMP mask replies are never sent
IP fast switching is enabled
IP fast switching on the same interface is disabled
IP Flow switching is disabled
IP CEF switching is enabled
IP CEF switching turbo vector
IP multicast fast switching is enabled
IP multicast distributed fast switching is disabled
IP route-cache flags are Fast, CEF
Router Discovery is disabled
IP output packet accounting is disabled
IP access violation accounting is disabled
TCP/IP header compression is disabled
RTP/IP header compression is disabled
Policy routing is disabled
Network address translation is disabled
BGP Policy Mapping is disabled
Input features: MCI Check
IPv4 WCCP Redirect outbound is disabled
IPv4 WCCP Redirect inbound is disabled
IPv4 WCCP Redirect exclude is disabled

```

- **show ip interface brief**—Displays a summary status of all interfaces, including IPv4 addressing information and interface and line protocols status (see Example 1-7).

Example 1-7 The **show ip interface brief** Command

```

R1# show ip interface brief

```

Interface	IP-Address	OK?	Method	Status	Protocol
GigabitEthernet0/0	172.16.1.1	YES	manual	up	up
GigabitEthernet0/1	unassigned	YES	unset	administratively down	down
Serial10/0/0	172.16.3.1	YES	manual	up	up
Serial10/0/1	192.168.10.5	YES	manual	up	up
Vlan1	unassigned	YES	unset	administratively down	down

- **show protocols**—Displays information about the routed protocol that is enabled and the protocol status of interfaces (see Example 1-8).

Example 1-8 The `show protocols` Command

```
R1# show protocols
Global values:
  Internet Protocol routing is enabled
GigabitEthernet0/0 is up, line protocol is up
  Internet address is 172.16.1.1/24
GigabitEthernet0/1 is administratively down, line protocol is down
Serial0/0/0 is up, line protocol is up
  Internet address is 172.16.3.1/30
Serial0/0/1 is up, line protocol is up
  Internet address is 192.168.10.5/30
Vlan1 is administratively down, line protocol is down
```

- **show cdp neighbors**—Tests the Layer 2 connection and provides information about directly connected CDP enabled Cisco devices (see Example 1-9).

Example 1-9 The `show cdp neighbors` Command

```
R1# show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  D - Remote, C - CVTA, M - Two-port MAC Relay
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone
Device ID        Local Intrfce  Holdtme    Capability  Platform  Port ID
R2                Ser 0/0/0      136        R           C1900     Ser 0/0/0
R3                Ser 0/0/1      133        R           C1900     Ser 0/0/0
```

This command tests the Layer 2 connection and displays information on directly connected Cisco devices. The information it provides includes the device ID, the local interface the device is connected to, capability (R = router, S = switch), the platform, and the port ID of the remote device. The `details` option includes IP addressing information and the IOS version.

Basic Switch CLI Commands (1.2.3.5)

Basic switch configuration includes the host name for identification, passwords for security, and assignment of IP addresses for connectivity. In-band access requires the switch to have an IP address. Example 1-10 shows the commands entered to enable a switch.

Example 1-11 shows the results of the configuration commands that were entered in Example 1-10. Verify and save the switch configuration by using the `copy running-config startup-config` command. To clear the switch configuration, use the `erase startup-config` command and then the `reload` command. It may also be necessary to erase any VLAN information by using the command `delete flash:vlan.dat`. When switch configurations are in place, view the configurations by using the `show running-config` command.

Example 1-10 Enabling a Switch with a Basic Configuration

```
Switch# enable
Switch# configure terminal
Switch(config)# hostname S1
S1(config)# enable secret class
S1(config)# line con 0
S1(config-line)# password cisco
S1(config-line)# login
S1(config-line)# line vty 0 4
S1(config-line)# password cisco
S1(config-line)# login
S1(config-line)# service password-encryption
S1(config-line)# exit
S1(config)#
S1(config)# service password-encryption
S1(config)# banner motd $ Authorized Access Only! $
S1(config)#
S1(config)# interface vlan 1
S1(config-if)# ip address 192.168.1.5 255.255.255.0
S1(config-if)# no shutdown
S1(config-if)# exit
S1(config)# ip default-gateway 192.168.1.1
S1(config)#
S1(config)# interface fa0/2
S1(config-if)# switchport mode access
S1(config-if)# switchport port-security
S1(config-if)# end
S1#
S1# copy running-config startup-config
```

Example 1-11 Switch Running Configuration

```
S1# show running-config
<some output omitted>
version 15.0
service password-encryption
!
hostname S1
!
enable secret 4 06YFDUHH61wAE/kLkDq9BGho1QM5EnRtoyr8cHAUg.2
!
interface FastEthernet0/2
  switchport mode access
  switchport port-security
!
```

```
interface Vlan1
  ip address 192.168.1.5 255.255.255.0
  !
ip default-gateway 192.168.1.1
  !
banner motd ^C Authorized Access Only ^C
  !
line con 0
  exec-timeout 0 0
  password 7 1511021F0725
  login
line vty 0 4
  password 7 1511021F0725
  login
line vty 5 15
  login
  !
end
S1#
```

Basic Switch Show Commands (1.2.3.6)

Switches make use of the following common IOS commands for configuration, to check for connectivity, and to display current switch status:

- **show port-security interface**—Displays any ports that have security activated. To examine a specific interface, include the interface ID. Information included in the output includes the maximum addresses allowed, the current count, the security violation count, and action to be taken (see Example 1-12).

Example 1-12 The show port-security interface Command

```
S1# show port-security interface fa0/2
Port Security           : Enabled
Port Status             : Secure-up
Violation Mode          : Shutdown
Aging Time              : 0 mins
Aging Type              : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses   : 1
Total MAC Addresses     : 1
Configured MAC Addresses : 0
Sticky MAC Addresses    : 0
Last Source Address:Vlan : 0024.50d1.9902:1
Security Violation Count : 0
```

- **show port-security address**—Displays all secure MAC addresses configured on all switch interfaces (see Example 1-13).

Example 1-13 The show port-security address Command

```
S1# show port-security address
Secure Mac Address Table
-----
Vlan      Mac Address      Type                               Ports    Remaining Age
                                                (mins)
-----
1         0024.50d1.9902   SecureDynamic                      Fa0/2    -
-----
Total Addresses in System (excluding one mac per port)  : 0
Max Addresses limit in System (excluding one mac per port) : 1536
```

- **show interfaces**—Displays one or all interfaces with line (protocol) status, bandwidth, delay, reliability, encapsulation, duplex, and I/O statistics (see Example 1-14).

Example 1-14 The show interfaces Command

```
S1# show interfaces fa0/2
FastEthernet0/2 is up, line protocol is up (connected)
Hardware is Fast Ethernet, address is 001e.14cf.eb04 (bia 001e.14cf.eb04)
MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
    reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full-duplex, 100Mb/s, media type is 10/100BaseTX
input flow-control is off, output flow-control is unsupported
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:00:08, output 00:00:00, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 2000 bits/sec, 3 packets/sec
 59 packets input, 11108 bytes, 0 no buffer
  Received 59 broadcasts (59 multicasts)
  0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
  0 watchdog, 59 multicast, 0 pause input
```

```

0 input packets with dribble condition detected
886 packets output, 162982 bytes, 0 underruns
0 output errors, 0 collisions, 1 interface resets
0 unknown protocol drops
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier, 0 pause output
0 output buffer failures, 0 output buffers swapped out

```

- **show mac-address-table**—Displays all MAC addresses that the switch has learned, how those addresses were learned (dynamic/static), the port number, and the VLAN assigned to the port (see Example 1-15).

Example 1-15 The show mac address-table Command

```

S1# show mac address-table
      Mac Address Table
-----
Vlan    Mac Address      Type        Ports
----    -
All     0100.0ccc.cccc   STATIC      CPU
All     0100.0ccc.cccd   STATIC      CPU
All     0180.c200.0000   STATIC      CPU
All     0180.c200.0001   STATIC      CPU
  1     001e.4915.5405   DYNAMIC     Fa0/3
  1     001e.4915.5406   DYNAMIC     Fa0/4
  1     0024.50d1.9901   DYNAMIC     Fa0/1
  1     0024.50d1.9902   STATIC      Fa0/2
  1     0050.56be.0e67   DYNAMIC     Fa0/1
  1     0050.56be.c23d   DYNAMIC     Fa0/6
  1     0050.56be.df70   DYNAMIC     Fa0/
Total Mac Addresses for this criterion: 11
S1#

```

Like routers, switches also support the **show cdp neighbors** command.

The same in-band and out-of-band management techniques that apply to routers also apply to switch configuration.

Summary (1.3)



Class Activity 1.3.1.1: Layered Network Design Simulation

Refer to *Scaling Networks v6 Labs & Study Guide* and the online course to complete this activity.

As the network administrator for a very small network, you want to prepare a simulated-network presentation for your branch manager to explain how the network currently operates.

The small network includes the following equipment:

- One 2911 Series router
- One 3560 switch
- One 2960 switch
- Four user workstations (PCs or laptops)
- One printer

Interactive Graphic

Activity 1.3.1.2: Basic Switch Configurations

Refer to the online course to complete this activity.

Packet Tracer Activity

Packet Tracer 1.3.1.3: Skills Integration Challenge

Background/Scenario

You are a recently hired LAN technician, and your network manager has asked you to demonstrate your ability to configure a small LAN. Your tasks include configuring initial settings on two switches using the Cisco IOS and configuring IP address parameters on host devices to provide end-to-end connectivity. You are to use two switches and two hosts/PCs on a cabled and powered network.

The hierarchical network design model divides network functionality into the access layer, the distribution layer, and the core layer. A campus wired LAN enables communications between devices in a building or group of buildings, as well as interconnection to the WAN and Internet edge at the network core.

A well-designed network controls traffic and limits the size of failure domains. Routers and switches can be deployed in pairs so that the failure of a single device does not cause service disruptions.

A network design should include an IP addressing strategy, scalable and fast-converging routing protocols, appropriate Layer 2 protocols, and modular or clustered devices that can be easily upgraded to increase capacity.

A mission-critical server should have connections to two different access layer switches. It should have redundant modules when possible, as well as a power backup source. It may be appropriate to provide multiple connections to one or more ISPs.

Security monitoring systems and IP telephony systems must have high availability and often require special design considerations.

It is important to deploy the appropriate type of routers and switches for a given set of requirements, features and specifications, and expected traffic flow.

Practice

The following activities provide practice with the topics introduced in this chapter. The Labs and Class Activities are available in the companion *Scaling Networks v6 Labs & Study Guide* (ISBN 978-1-58713-433-3). The Packet Tracer activity instructions are also in the *Labs & Study Guide*. The PKA files are found in the online course.



Class Activities

Class Activity 1.0.1.2: Network by Design

Class Activity 1.3.1.1: Layered Network Design Simulation



Packet Tracer Activities

Packet Tracer 1.2.1.7: Comparing 2960 and 3560 Switches

Packet Tracer 1.3.1.3: Skills Integration Challenge

Check Your Understanding Questions

Complete all the review questions listed here to test your understanding of the topics and concepts in this chapter. The appendix “Answers to ‘Check Your Understanding’ Questions” lists the answers.

1. In the Cisco Enterprise Architecture, which two functional parts of the network are combined to form a collapsed core design? (Choose two.)
 - A. Access layer
 - B. Core layer
 - C. Distribution layer
 - D. Enterprise edge
 - E. Provider edge
2. Which design feature limits the impact of a distribution switch failure in an enterprise network?
 - A. The installation of redundant power supplies
 - B. The purchase of enterprise equipment that is designed for large traffic volume
 - C. The use of a collapsed core design
 - D. The use of the building switch block approach
3. What are two benefits of extending access layer connectivity to users through a wireless medium? (Choose two.)
 - A. Decreased number of critical points of failure
 - B. Increased bandwidth availability
 - C. Increased flexibility
 - D. Increased network management options
 - E. Reduced costs
4. As the network administrator, you have been asked to implement EtherChannel on the corporate network. What does this configuration consist of?
 - A. Grouping multiple physical ports to increase bandwidth between two switches
 - B. Grouping two devices to share a virtual IP address
 - C. Providing redundant devices to allow traffic to flow in the event of device failure
 - D. Providing redundant links that dynamically block or forward traffic

5. Which statement describes a characteristic of Cisco Meraki switches?
 - A. They are campus LAN switches that perform the same functions as Cisco 2960 switches.
 - B. They are cloud-managed access switches that enable virtual stacking of switches.
 - C. They are service provider switches that aggregate traffic at the edge of the network.
 - D. They promote infrastructure scalability, operational continuity, and transport flexibility.

6. What term is used to express the thickness or height of a switch?
 - A. Domain size
 - B. Module size
 - C. Port density
 - D. Rack unit

7. What are two functions of a router? (Choose two.)
 - A. It connects multiple IP networks.
 - B. It controls the flow of data through the use of Layer 2 addresses.
 - C. It determines the best path for sending packets.
 - D. It increases the size of the broadcast domain.
 - E. It manages the VLAN database.

8. Which two requirements must always be met to use in-band management to configure a network device? (Choose two.)
 - A. A direct connection to the console port
 - B. A direct connection to the auxiliary port
 - C. A terminal emulation client
 - D. At least one network interface that is connected and operational
 - E. Telnet, SSH, or HTTP access to the device

9. What are two ways to access a Cisco switch for out-of-band management? (Choose two.)
 - A. A connection that uses HTTP
 - B. A connection that uses the AUX port
 - C. A connection that uses the console port
 - D. A connection that uses SSH
 - E. A connection that uses Telnet

Scaling VLANs

Objectives

Upon completion of this chapter, you will be able to answer the following questions:

- How does VLAN Trunking Protocol (VTP) Version 1 compare with Version 2?
- How do you configure VTP Versions 1 and 2?
- How do you configure extended VLANs?
- How do you configure Dynamic Trunking Protocol (DTP)?
- How do you troubleshoot common inter-VLAN configuration issues?
- How do you troubleshoot common IP addressing issues in an inter-VLAN routed environment?
- How do you configure inter-VLAN routing using Layer 3 switching?

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary.

virtual local area network (VLANs) Page 48

trunk Page 48

VLAN Trunking Protocol (VTP) Page 48

Dynamic Trunking Protocol (DTP) Page 48

extended-range VLAN Page 48

vlan.dat Page 50

VTP domain Page 50

VTP advertisement Page 50

VTP mode Page 50

VTP server Page 50

VTP client Page 51

VTP transparent Page 51

summary advertisement Page 52

advertisement request Page 52

subset advertisement Page 52

normal-range VLAN Page 53

inter-VLAN routing Page 75

legacy inter-VLAN routing Page 77

router-on-a-stick inter-VLAN routing Page 78

Layer 3 inter-VLAN routing Page 90

routed port Page 90

switch virtual interface (SVI) Page 91

Cisco Express Forwarding Page 91

Introduction (2.0.1.1)

As the number of switches increases on a small or medium-sized business network, the overall administration required to manage *virtual local area networks (VLANs)* and *trunks* in the network becomes challenging. This chapter examines some of the strategies and protocols that can be used to manage VLANs and trunks.

VLAN Trunking Protocol (VTP) reduces administration in a switched network. A switch in VTP server mode can manage additions, deletions, and renaming of VLANs across the domain. For example, when a new VLAN is added on the VTP server, the VLAN information is distributed to all switches in the domain. This eliminates the need to configure the new VLAN on every switch. VTP is a Cisco proprietary protocol that is available on most of the Cisco Catalyst Series products.

Using VLANs to segment a switched network provides improved performance, manageability, and security. Trunks are used to carry information from multiple VLANs between devices. *Dynamic Trunking Protocol (DTP)* provides the ability for ports to automatically negotiate trunking between switches.

Because VLANs segment a network, and each is on its own network or subnet, a Layer 3 process is required to allow traffic to move from one VLAN to another.

This chapter examines the implementation of inter-VLAN routing using a Layer 3 switch. It also describes issues encountered when implementing VTP, DTP, and inter-VLAN routing.

VTP, Extended VLANs, and DTP (2.1)

Several technologies help simplify interswitch connectivity. VTP simplifies VLAN management in a switched network. VLANs are created and managed on VTP servers. Layer 2 access switches are typically configured as VTP clients that automatically update their VLAN database from VTP servers. Some Catalyst switches support the creation of *extended-range VLANs*. Extended-range VLANs, which are popular with service providers to segment their many clients, are numbered 1006 to 4094. Only transparent VTP mode switches can create extended VLANs. Finally, trunking must be enabled to transport VLAN frames between switches. DTP provides the ability for ports to automatically negotiate trunking between switches.

In this section, you will learn how to configure all of the enhanced interswitch connectivity technologies.

VTP Concepts and Operation (2.1.1)

VTP propagates and synchronizes VLAN information to other switches in the VTP domain. There are currently three versions of VTP: VTP Version 1, VTP Version 2, and VTP Version 3. The focus of this topic is to compare VTP Versions 1 and 2.

VTP Overview (2.1.1.1)

As the number of switches increases on a small or medium-sized business network, the overall administration required to manage VLANs and trunks in the network becomes challenging. In larger networks, VLAN management can become daunting. In Figure 2-1, assume that VLANs 10, 20, and 99 have already been implemented, and you must now add VLAN 30 to all switches. Manually adding the VLAN in this network would involve individually configuring 12 switches.

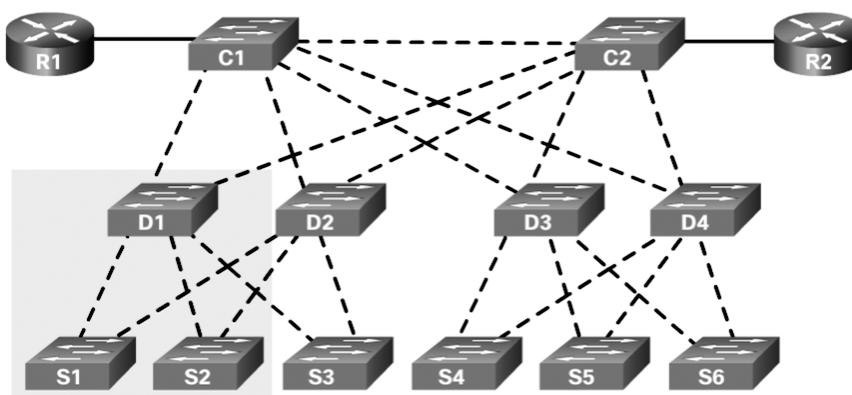


Figure 2-1 The VLAN Management Challenge

VTP allows a network administrator to manage VLANs on a master switch configured as a VTP server. The VTP server distributes and synchronizes VLAN information over trunk links to VTP-enabled client switches throughout the switched network. This minimizes problems caused by incorrect configurations and configuration inconsistencies.

Note

VTP only learns about normal-range VLANs (VLAN IDs 1 to 1005). Extended-range VLANs (IDs greater than 1005) are not supported by VTP Version 1 or Version 2. VTP Version 3 does support extended VLANs but is beyond the scope of this course.

Note

VTP stores VLAN configurations in a database called *vlan.dat*.

Table 2-1 provides a brief description of important components of VTP.

Table 2-1 VTP Components

VTP Components	Definition
<i>VTP domain</i>	<p>A VTP domain consists of one or more interconnected switches.</p> <p>All switches in a domain share VLAN configuration details by using VTP advertisements.</p> <p>Switches that are in different VTP domains do not exchange VTP messages.</p> <p>A router or Layer 3 switch defines the boundary of a domain.</p>
<i>VTP advertisements</i>	<p>Each switch in a VTP domain sends periodic VTP advertisements from each trunk port to a reserved Layer 2 multicast address.</p> <p>Neighboring switches receive these advertisements and update their VTP and VLAN configurations as necessary.</p>
<i>VTP modes</i>	<p>A switch can be configured as a VTP server, client, or transparent.</p>
VTP password	<p>Switches in the VTP domain can also be configured with a password.</p>

Note

VTP advertisements are not exchanged if the trunk between switches is inactive.

VTP Modes (2.1.1.2)

A switch can be configured in one of three VTP modes, as described in Table 2-2.

Table 2-2 VTP Modes

VTP Mode	Definition
<i>VTP server</i>	<p>VTP servers advertise the VTP domain VLAN information to other VTP-enabled switches in the same VTP domain.</p> <p>VTP servers store the VLAN information for the entire domain in NVRAM.</p> <p>The VTP server is where VLANs can be created, deleted, or renamed for the domain.</p>

VTP Mode	Definition
<i>VTP client</i>	<p>VTP clients function the same way as VTP servers, but you cannot create, change, or delete VLANs on a VTP client.</p> <p>A VTP client stores the VLAN information for the entire domain only while the switch is on.</p> <p>A switch reset deletes the VLAN information.</p> <p>You must configure VTP client mode on a switch.</p>
<i>VTP transparent</i>	<p>Transparent switches do not participate in VTP except to forward VTP advertisements to VTP clients and VTP servers.</p> <p>A VLAN that is created, renamed, or deleted on a transparent switch is local to that switch only.</p> <p>To create an extended VLAN, a switch must be configured as a VTP transparent switch when using VTP Version 1 or Version 2.</p>

Table 2-3 summarizes the operation of the three VTP modes.

Table 2-3 Comparing VTP Modes

VTP Question	VTP Server	VTP Client	VTP Transparent
What are the differences?	Manages domain and VLAN configuration. Multiple VTP servers can be configured.	Updates local VTP configurations. VTP client switches cannot change VLAN configurations.	Manages local VLAN configurations. VLAN configurations are not shared with the VTP network.
Does it respond to VTP advertisements?	Participates fully	Participates fully	Forwards only VTP advertisements
Is the global VLAN configuration preserved on restart?	Yes, global configurations are stored in NVRAM.	No, global configurations are stored in RAM only.	No, the local VLAN configuration is stored only in NVRAM.
Does it update other VTP-enabled switches?	Yes	Yes	No

Note

A switch that is in server or client mode with a higher configuration revision number than the existing VTP server updates all VLAN information in the VTP domain. (Configuration revision numbers are discussed later in this chapter.) As a best practice, Cisco recommends deploying a new switch in VTP transparent mode and then configuring the VTP domain specifics.

VTP Advertisements (2.1.1.3)

VTP includes three types of advertisements:

- *Summary advertisements*—These inform adjacent switches of the VTP domain name and configuration revision number.
- *Advertisement requests*—These are in response to a summary advertisement message when the summary advertisement contains a higher configuration revision number than the current value.
- *Subset advertisements*—These contain VLAN information, including any changes.

By default, Cisco switches issue summary advertisements every five minutes. Summary advertisements inform adjacent VTP switches of the current VTP domain name and the configuration revision number.

The configuration revision number is a 32-bit number that indicates the level of revision for a VTP packet. Each VTP device tracks the VTP configuration revision number that is assigned to it.

This information is used to determine whether the received information is more recent than the current version. The revision number increases by 1 each time you add a VLAN, delete a VLAN, or change a VLAN name. If the VTP domain name is changed or the switch is set to transparent mode, the revision number is reset to 0.

Note

To reset a configuration revision on a switch, change the VTP domain name and then change the name back to the original name.

When the switch receives a summary advertisement packet, the switch compares the VTP domain name to its own VTP domain name. If the name is different, the switch simply ignores the packet. If the name is the same, the switch then compares the configuration revision to its own revision. If its own configuration revision number is higher or equal to the packet's configuration revision number, the packet is ignored. If its own configuration revision number is lower, an advertisement request is sent, asking for the subset advertisement message.

The subset advertisement message contains the VLAN information with any changes. When you add, delete, or change a VLAN on the VTP server, the VTP server increments the configuration revision and issues a summary advertisement. One or several subset advertisements follow the summary advertisement containing the VLAN information, including any changes. This process is shown in Figure 2-2.

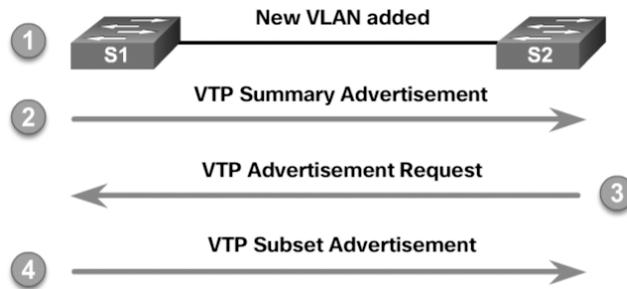


Figure 2-2 VTP Advertisements

VTP Versions (2.1.1.4)

VTP Version 1 (VTPv1) and Version 2 (VTPv2) are described in Table 2-4. Switches in the same VTP domain must use the same VTP version.

Table 2-4 VTP Versions

VTP Version	Definition
VTP Version 1	Default VTP mode on all switches. Supports <i>normal-range VLANs</i> only.
VTP Version 2	Supports normal-range VLANs only. Supports legacy Token Ring networks. Supports advanced features, including unrecognized Type-Length-Value (TLV), version-dependent transparent mode, and consistency checks.

Note

VTPv2 is not much different from VTPv1 and is generally configured only if legacy Token Ring support is required. The newest version is VTP Version 3 (VTPv3). VTPv3 is beyond the scope of this course.

Default VTP Configuration (2.1.1.5)

The `show vtp status` privileged EXEC command displays the VTP status. Executing the command on a Cisco 2960 Plus Series switch generates the output shown in Example 2-1.

Example 2-1 Verifying Default VTP Status

```

S1# show vtp status
VTP Version capable           : 1 to 3
VTP version running           : 1
VTP Domain Name               :
VTP Pruning Mode              : Disabled
VTP Traps Generation          : Disabled
Device ID                     : f078.167c.9900
Configuration last modified by 0.0.0.0 at 3-1-93 00:02:11

Feature VLAN:
-----
VTP Operating Mode            : Transparent
Maximum VLANs supported locally : 255
Number of existing VLANs      : 12
Configuration Revision         : 0
MD5 digest                    : 0x57 0xCD 0x40 0x65 0x63 0x59 0x47 0xBD
                               0x56 0x9D 0x4A 0x3E 0xA5 0x69 0x35 0xBC

S1#

```

Table 2-5 briefly describes the command output for the `show vtp status` parameters.

Table 2-5 Command Output Description

Command Output	Description
VTP Version capable and VTP version running	Display the VTP version that the switch is capable of running and the version that it is currently running. Switches implement VTPv1 by default. Newer switches may support VTPv3.
VTP Domain Name	Name that identifies the administrative domain for the switch. VTP domain name is case sensitive. The VTP domain name is NULL by default.
VTP Pruning Mode	Displays whether pruning is enabled or disabled (default). VTP pruning prevents flooded traffic from propagating to switches that do not have members in specific VLANs.
VTP Traps Generation	Displays whether VTP traps are sent to a network management station. VTP traps are disabled by default.
Device ID	The switch MAC address.

Command Output	Description
Configuration last modified	Date and time of the last configuration modification and IP address of the switch that caused the configuration change to the database.
VTP Operating Mode	Can be server (default), client, or transparent.
Maximum VLANs supported locally	The number of VLANs supported varies across switch platforms.
Number of existing VLANs	Includes the number of default and configured VLANs. The default number of existing VLANs varies across switch platforms.
Configuration Revision	The current configuration revision number on this switch. The revision number is a 32-bit number that indicates the level of revision for a VTP frame. The default configuration number for a switch is 0. Each time a VLAN is added or removed, the configuration revision number is incremented. Each VTP device tracks the VTP configuration revision number that is assigned to it.
MD5 digest	A 16-byte checksum of the VTP configuration.

VTP Caveats (2.1.1.6)

Some network administrators avoid VTP because it could potentially introduce false VLAN information into the existing VTP domain. The configuration revision number is used when determining whether a switch should keep its existing VLAN database or overwrite it with the VTP update sent by another switch in the same domain with the same password.

Adding a VTP-enabled switch to an existing VTP domain wipes out the existing VLAN configurations in the domain if the new switch is configured with different VLANs and has a higher configuration revision number than the existing VTP server. The new switch can be either a VTP server or a client switch. This propagation can be difficult to correct.

To illustrate this problem, refer to the example in Figure 2-3. The S1 switch is the VTP server, and the S2 and S3 switches are VTP clients. All switches are in the `cisco1` domain, and the current VTP revision is 17. In addition to the default VLAN 1, the VTP server (S1) has VLANs 10 and 20 configured. These VLANs have been propagated by VTP to the other two switches.

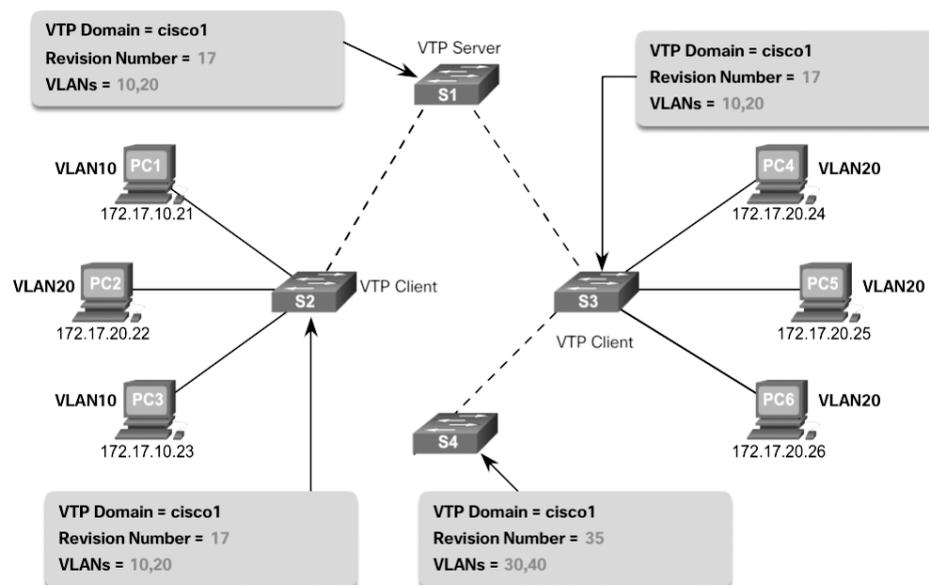


Figure 2-3 Incorrect VTP Configuration Revision Number Scenario

A network technician adds S4 to the network to address the need for additional capacity. However, the technician has not erased the startup configuration or deleted the VLAN.DAT file on S4. S4 has the same VTP domain name configured as the other two switches, but its revision number is 35, which is higher than 17, the current revision number on the other two switches.

S4 has VLAN 1 and is configured with VLANs 30 and 40. But it does not have VLANs 10 and 20 in its database. Unfortunately, because S4 has a higher revision number, all the other switches in the domain will sync to S4's revision. The result is that VLANs 10 and 20 will no longer exist on the switches, leaving clients that are connected to ports belonging to those nonexistent VLANs without connectivity.

Therefore, when a switch is added to a network, ensure that it has a default VTP configuration. The VTP configuration revision number is stored in NVRAM (or flash memory, on some platforms) and is not reset if you erase the switch configuration and reload it. To reset the VTP configuration revision number to 0, you have two options:

- Change the switch's VTP domain to a nonexistent VTP domain and then change the domain back to the original name.
- Change the switch's VTP mode to transparent and then back to the previous VTP mode.

Note

The commands to reset the VTP configuration revision number are discussed in the next topic.

Interactive Graphic**Activity 2.1.1.7: Identify VTP Concepts and Operations**

Refer to the online course to complete this activity.

VTP Configuration (2.1.2)

The focus of this topic is on how to configure VTP Versions 1 and 2.

VTP Configuration Overview (2.1.2.1)

Complete the following steps to configure VTP:

- Step 1.** Configure the VTP server.
- Step 2.** Configure the VTP domain name and password.
- Step 3.** Configure the VTP clients.
- Step 4.** Configure VLANs on the VTP server.
- Step 5.** Verify that the VTP clients have received the new VLAN information.

Figure 2-4 shows the reference topology used in this section for configuring and verifying a VTP implementation. Switch S1 is the VTP server, and S2 and S3 are VTP clients.

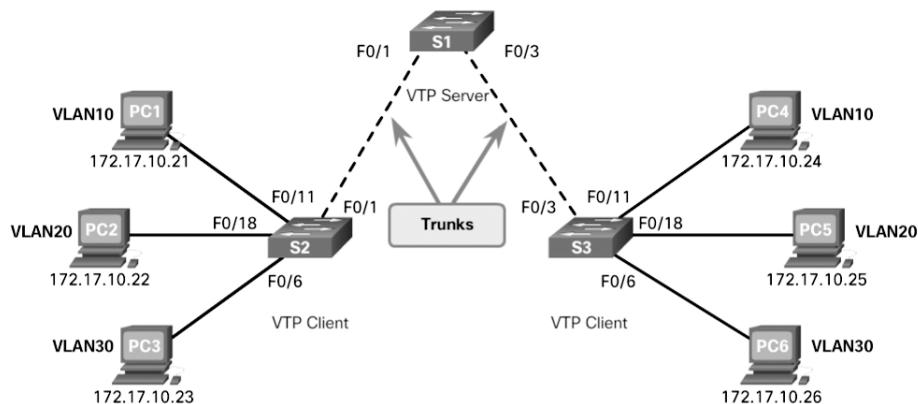


Figure 2-4 VTP Configuration Topology

Step 1—Configure the VTP Server (2.1.2.2)

Confirm that all switches are configured with default settings to avoid any issues with configuration revision numbers. Configure S1 as the VTP server by using the `vtp mode server` global configuration command, as shown in Example 2-2.

Example 2-2 Configuring VTP Server Mode

```
S1# conf t
Enter configuration commands, one per line. End with CNTL/Z.
S1(config)# vtp mode ?
    client      Set the device to client mode.
    off         Set the device to off mode.
    server      Set the device to server mode.
    transparent Set the device to transparent mode.

S1(config)# vtp mode server
Setting device to VTP Server mode for VLANs.
S1(config)# end
S1#
```

Issue the `show vtp status` command to confirm that S1 is a VTP server, as shown in Example 2-3.

Example 2-3 Verifying VTP Mode

```
S1# show vtp status
VTP Version capable      : 1 to 3
VTP version running      : 1
VTP Domain Name          :
VTP Pruning Mode         : Disabled
VTP Traps Generation     : Disabled
Device ID                 : f078.167c.9900
Configuration last modified by 0.0.0.0 at 3-1-93 00:02:11
Local updater ID is 0.0.0.0 (no valid interface found)

Feature VLAN:
-----
VTP Operating Mode       : Server
Maximum VLANs supported locally : 255
Number of existing VLANs : 5
Configuration Revision   : 0
MD5 digest               : 0x57 0xCD 0x40 0x65 0x63 0x59 0x47 0xBD
                          0x56 0x9D 0x4A 0x3E 0xA5 0x69 0x35 0xBC

S1#
```

Notice that the configuration revision number is still set to 0, and the number of existing VLANs is five. This is because no VLANs have yet been configured, and the switch does not belong to a VTP domain. The five VLANs are the default VLAN 1 and VLANs 1002 through 1005.

Step 2—Configure the VTP Domain Name and Password (2.1.2.3)

The domain name is configured by using the `vtp domain domain-name` global configuration command. In Example 2-4, the domain name is configured as `CCNA` on S1. S1 then sends out a VTP advertisement to S2 and S3. If S2 and S3 have the default configuration with the NULL domain name, both switches accept `CCNA` as the new VTP domain name. A VTP client must have the same domain name as the VTP server before it will accept VTP advertisements.

Example 2-4 Configuring the VTP Domain Name

```
S1(config)# vtp domain ?
WORD The ascii name for the VTP administrative domain.

S1(config)# vtp domain CCNA
Changing VTP domain name from NULL to CCNA
*Mar  1 02:55:42.768: %SW_VLAN-6-VTP_DOMAIN_NAME_CHG: VTP domain name changed to
CCNA.
S1(config)#
```

For security reasons, a password should be configured using the `vtp password password` global configuration command. In Example 2-5, the VTP domain password is set to `cisco12345`. All switches in the VTP domain must use the same VTP domain password to successfully exchange VTP messages.

Example 2-5 Configuring and Verifying the VTP Domain Password

```
S1(config)# vtp password cisco12345
Setting device VTP password to cisco12345
S1(config)# end

S1# show vtp password
VTP Password: cisco12345
S1#
```

Use the `show vtp password` command to verify the password entered, as shown in Example 2-5.


```

                                                    Fa0/19, Fa0/20, Fa0/21, Fa0/22
                                                    Fa0/23, Fa0/24, Gi0/1, Gi0/2
10 SALES active
20 MARKETING active
30 ACCOUNTING active
1002 fddi-default act/unsup
1003 token-ring-default act/unsup
1004 fddinet-default act/unsup
1005 trnet-default act/unsup
S1#

```

Notice that the three VLANs are now in the VLAN database. Verify the VTP status, as shown in Example 2-9.

Example 2-9 Verifying the VTP Status

```

S1# show vtp status
VTP Version capable      : 1 to 3
VTP version running     : 1
VTP Domain Name         : CCNA
VTP Pruning Mode        : Disabled
VTP Traps Generation    : Disabled
Device ID                : f078.167c.9900
Configuration last modified by 0.0.0.0 at 3-1-93 02:02:45
Local updater ID is 0.0.0.0 (no valid interface found)

Feature VLAN:
-----
VTP Operating Mode      : Server
Maximum VLANs supported locally : 255
Number of existing VLANs : 8
Configuration Revision  : 6
MD5 digest              : 0xFE 0x8D 0x2D 0x21 0x3A 0x30 0x99 0xC8
                        0xDB 0x29 0xBD 0xE9 0x48 0x70 0xD6 0xB6
*** MD5 digest checksum mismatch on trunk: Fa0/2 ***
S1#

```

Notice that the configuration revision number incremented six times, from the default 0 to 6. This is because three new named VLANs were added. Each time the administrator makes a change to the VTP server's VLAN database, this number increases by 1. The number increased by 1 each time a VLAN was added or named.

Step 5—Verify That the VTP Clients Have Received the New VLAN Information (2.1.2.6)

On S2, verify that the VLANs configured on S1 have been received and entered into the S2 VLAN database by using the `show vlan brief` command, as shown in Example 2-10.

Example 2-10 Verifying That the VTP Clients Have Received the New VLAN Information

```
S2# show vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Fa0/2, Fa0/3, Fa0/4, Fa0/5 Fa0/6, Fa0/7, Fa0/8, Fa0/9 Fa0/10, Fa0/11, Fa0/12, Fa0/13 Fa0/14, Fa0/15, Fa0/16, Fa0/17 Fa0/18, Fa0/19, Fa0/20, Fa0/21 Fa0/22, Fa0/23, Fa0/24, Gi0/1 Gi0/2
10	SALES	active	
20	MARKETING	active	
30	ACCOUNTING	active	
1002	fddi-default	act/unsup	
1003	token-ring-default	act/unsup	
1004	fddinet-default	act/unsup	
1005	trnet-default	act/unsup	

```
S2#
```

As expected, the VLANs configured on the VTP server have propagated to S2. Verify the VTP status on S2, as shown in Example 2-11.

Example 2-11 Verifying the VTP Status on S2

```
S2# show vtp status
```

VTP Version capable	: 1 to 3
VTP version running	: 1
VTP Domain Name	: CCNA
VTP Pruning Mode	: Disabled
VTP Traps Generation	: Disabled
Device ID	: b07d.4729.2400
Configuration last modified by	0.0.0.0 at 3-1-93 02:02:45

```

Feature VLAN:
-----
VTP Operating Mode           : Client
Maximum VLANs supported locally : 255
Number of existing VLANs     : 8
Configuration Revision       : 6
MD5 digest                   : 0xFE 0x8D 0x2D 0x21 0x3A 0x30 0x99 0xC8
                               0xDB 0x29 0xBD 0xE9 0x48 0x70 0xD6 0xB6
S2#

```

Notice that the configuration revision number on S2 is the same as the number on the VTP server.

Because S2 is operating in VTP client mode, attempts to configure VLANs are not allowed, as shown in Example 2-12.

Example 2-12 Attempting to Configure a VLAN on a Client

```

S2(config)# vlan 99
VTP VLAN configuration not allowed when device is in CLIENT mode.
S2(config)#

```

Extended VLANs (2.1.3)

All Catalyst switches can create normal-range VLANs. Some switches can also use extended-range VLANs.

The focus of this topic is on how to configure extended VLANs.

VLAN Ranges on Catalyst Switches (2.1.3.1)

Different Cisco Catalyst switches support various numbers of VLANs. The number of supported VLANs is typically large enough to accommodate the needs of most organizations. For example, the Catalyst 2960 and 3560 Series switches support more than 4000 VLANs. Normal-range VLANs on these switches are numbered 1 to 1005, and extended-range VLANs are numbered 1006 to 4094.

Example 2-13 displays the available VLANs on a Catalyst 2960 switch running Cisco IOS Release 15.x.

Example 2-13 Verifying VLANs on a Catalyst 2960 Switch

```

Switch# show vlan brief

VLAN Name                Status    Ports
-----
1      default                active    Fa0/2, Fa0/3, Fa0/4, Fa0/5
                                           Fa0/6, Fa0/7, Fa0/8, Fa0/9
                                           Fa0/10, Fa0/11, Fa0/12, Fa0/13
                                           Fa0/14, Fa0/15, Fa0/16, Fa0/17
                                           Fa0/18, Fa0/19, Fa0/20, Fa0/21
                                           Fa0/22, Fa0/23, Fa0/24, Gi0/1
                                           Gi0/2
1002  fddi-default            act/unsup
1003  token-ring-default      act/unsup
1004  fddinet-default         act/unsup
1005  trnet-default           act/unsup
Switch#

```

Table 2-6 shows the features of normal-range and extended-range VLANs.

Table 2-6 Types of VLANs

Type	Definition
Normal-range VLANs	<p>Used in small and medium-sized business and enterprise networks.</p> <p>Identified by VLAN IDs between 1 and 1005.</p> <p>IDs 1 and 1002 to 1005 are automatically created and cannot be removed. (IDs 1002 through 1005 are reserved for Token Ring and Fiber Distributed Data Interface [FDDI] VLANs.)</p> <p>Configurations are stored within a VLAN database file called vlan.dat, which is stored in flash memory.</p>
Extended-range VLANs	<p>Used by service providers and large organizations to extend their infrastructure to a greater number of customers.</p> <p>Identified by VLAN IDs between 1006 and 4094.</p> <p>Support fewer VLAN features than normal-range VLANs.</p> <p>Configurations are saved in the running configuration file.</p>

VLAN Trunking Protocol (VTP), which helps manage VLAN configurations between switches, can learn and store only normal-range VLANs. VTP does not function with extended-range VLANs.

Note

4096 is the upper boundary for the number of VLANs available on Catalyst switches because there are 12 bits in the VLAN ID field of the IEEE 802.1Q header.

Creating a VLAN (2.1.3.2)

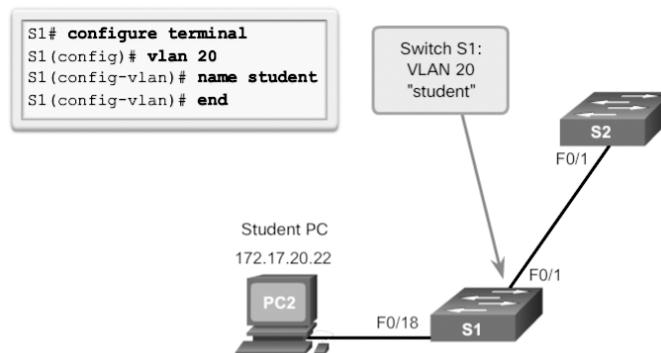
When configuring normal-range VLANs, the configuration details are stored in flash memory on the switch, in a file called **vlan.dat**. Flash memory is persistent and does not require the **copy running-config startup-config** command. However, because other details are often configured on a Cisco switch at the same time that VLANs are created, it is good practice to save running configuration changes to the startup configuration file.

Table 2-7 shows the Cisco IOS command syntax used to add a VLAN to a switch and give it a name. Naming each VLAN is considered a best practice in switch configuration.

Table 2-7 Command Syntax for Creating a VLAN

Command	Description
S1(config)# vlan <i>vlan-id</i>	Create a VLAN with a valid ID number.
S1(config-vlan)# name <i>vlan-name</i>	Specify a unique name to identify the VLAN.

Figure 2-5 shows how the student VLAN (VLAN 20) is configured on switch S1. In the topology example, notice that the student computer (PC2) has been assigned an IP address that is appropriate for VLAN 20, but the port to which the PC attaches has not been associated with a VLAN yet.

**Figure 2-5** Sample VLAN Configuration

The `vlan vlan-id` command can be used to create several VLANs at once. To do so, enter a series of VLAN IDs separated by commas. You also can enter a range of VLAN IDs separated by hyphens. For example, the following command would create VLANs 100, 102, 105, 106, and 107.

```
S1(config)# vlan 100,102,105-107
```

Assigning Ports to VLANs (2.1.3.3)

After creating a VLAN, the next step is to assign ports to the VLAN. An access port can belong to only one VLAN at a time; one exception to this rule is a port connected to an IP phone, in which case there are two VLANs associated with the port: one for voice and one for data.

Table 2-8 shows the syntax for defining a port to be an access port and assigning it to a VLAN. The `switchport mode access` command is optional but strongly recommended as a security best practice. With this command, the interface changes to permanent access mode.

Table 2-8 Command Syntax for Assigning Ports to VLANs

Command	Description
S1(config)# <code>interface <i>interface_id</i></code>	Enter interface configuration mode.
S1(config-if)# <code>switchport mode access</code>	Set the port to access mode.
S1(config-if)# <code>switchport access vlan <i>vlan_id</i></code>	Assign the port to a VLAN.

Note

Use the `interface range` command to simultaneously configure multiple interfaces.

In the example in Figure 2-6, VLAN 20 is assigned to port F0/18 on switch S1; therefore, the student computer (PC2) is in VLAN 20. When VLAN 20 is configured on other switches, the network administrator knows to configure the other student computers to be in the same subnet as PC2 (172.17.20.0/24).

The `switchport access vlan` command forces the creation of a VLAN if it does not already exist on the switch. For example, VLAN 30 is not present in the `show vlan brief` output of the switch. If the `switchport access vlan 30` command is entered on any interface with no previous configuration, the switch displays the following:

```
% Access VLAN does not exist. Creating vlan 30
```

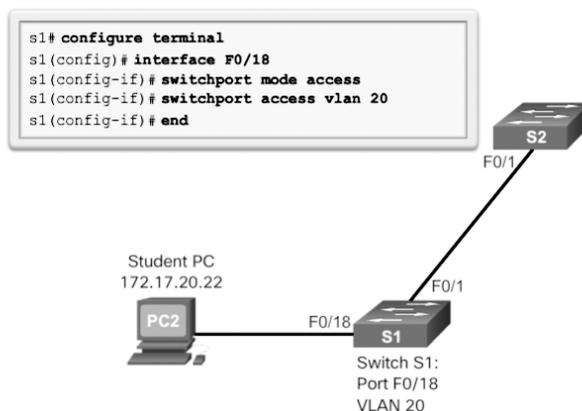


Figure 2-6 Sample VLAN Port Assignment Configuration

Verifying VLAN Information (2.1.3.4)

After a VLAN is configured, VLAN configurations can be validated using Cisco IOS `show` commands.

Table 2-9 shows the `show vlan` command options.

```
show vlan [brief | id vlan-id | name vlan-name | summary]
```

Table 2-9 The `show vlan` Command

brief	Display one line for each VLAN with the VLAN name, status, and its ports.
id <i>vlan-id</i>	Display information about a single VLAN identified by VLAN ID number. For <i>vlan-id</i> , the range is 1 to 4094.
name <i>vlan-name</i>	Display information about a single VLAN identified by VLAN name. <i>vlan-name</i> is an ASCII string from 1 to 32 characters.
summary	Display VLAN summary information.

Table 2-10 shows the `show interfaces` command options.

```
show interfaces [interface-id | vlan vlan-id] | switchport
```

Table 2-10 The **show interfaces** Command

<i>interface-id</i>	Display information about a specific interface. Valid interfaces include physical ports (including type, module, and port number) and port channels. The port channel range is 1 to 6.
vlan <i>vlan-id</i>	Display information about a specific VLAN. The <i>vlan-id</i> range is 1 to 4094.
switchport	Display the administrative and operational status of a switching port, including port blocking and port protection settings.

In Example 2-14, the **show vlan name student** command displays information that would also be found in the **show vlan brief** command, but only for VLAN 20, the student VLAN.

Example 2-14 Using the **show vlan** Command

```

S1# show vlan name student

VLAN Name                Status    Ports
-----
20    student                active    Fa0/11, Fa0/18

VLAN Type  SAID      MTU    Parent  RingNo BridgeNo  Stp   BrdgMode Trans1 Trans2
-----
20    enet    100020   1500   -       -       -     -       -       0     0

Remote SPAN VLAN
-----
Disabled

Primary Secondary Type                Ports
-----

S1# show vlan summary
Number of existing VLANs           : 7
Number of existing VTP VLANs       : 7
Number of existing extended VLANs  : 0

S1#

```

Example 2-14 indicates that the status is active, and specifies which switch ports are assigned to the VLAN. The **show vlan summary** command displays the count of all configured VLANs. The output in Example 2-14 shows seven VLANs.

The `show interfaces vlan vlan-id` command displays details about the VLAN. In the second line, it indicates whether the VLAN is up or down, as shown in Example 2-15.

Example 2-15 Using the `show interfaces vlan` Command

```
S1# show interfaces vlan 99
Vlan99 is up, line protocol is up

  Hardware is EtherSVI, address is 0cd9.96e2.3d41 (bia 0cd9.96e2.3d41)
  Internet address is 192.168.99.1/24
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:35, output 00:01:01, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    1 packets input, 60 bytes, 0 no buffer
    Received 0 broadcasts (0 IP multicasts)
    0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
    1 packets output, 64 bytes, 0 underruns
    0 output errors, 1 interface resets
    0 unknown protocol drops
    0 output buffer failures, 0 output buffers swapped out

S1#
```

Configuring Extended VLANs (2.1.3.5)

Extended-range VLANs are identified by a VLAN ID between 1006 and 4094. Example 2-16 shows that, by default, a Catalyst 2960 Plus Series switch does not support extended VLANs.

Example 2-16 Extended VLAN Failure

```
S1# conf t
Enter configuration commands, one per line. End with CNTL/Z.
S1(config)# vlan 2000
S1(config-vlan)# exit
```

```

% Failed to create VLANs 2000
Extended VLAN(s) not allowed in current VTP mode.
%Failed to commit extended VLAN(s) changes.

S1(config)#
*Mar  1 00:51:48.893: %SW_VLAN-4-VLAN_CREATE_FAIL: Failed to create VLANs 2000:
  extended VLAN(s) not allowed in current VTP mode

```

In order to configure an extended VLAN on a 2960 switch, it must be set to VTP transparent mode. Example 2-17 shows how to create an extended-range VLAN on the Catalyst 2960 Plus Series switch.

Example 2-17 Configuring an Extended VLAN on a 2960 Switch

```

S1(config)# vtp mode transparent
Setting device to VTP Transparent mode for VLANs.
S1(config)# vlan 2000
S1(config-vlan)# end
S1#

```

The **show vlan brief** command is used to verify that a VLAN was created, as shown in Example 2-18. This output confirms that the extended VLAN 2000 has been configured and is active.

Example 2-18 Verifying an Extended VLAN Configuration

```

S1# show vlan brief

VLAN Name                Status    Ports
-----
1    default                 active    Fa0/3, Fa0/4, Fa0/5, Fa0/6
                                           Fa0/7, Fa0/8, Fa0/9, Fa0/10
                                           Fa0/11, Fa0/12, Fa0/13, Fa0/14
                                           Fa0/15, Fa0/16, Fa0/17, Fa0/18
                                           Fa0/19, Fa0/20, Fa0/21, Fa0/22
                                           Fa0/23, Fa0/24, Gi0/1, Gi0/2

1002 fddi-default         act/unsup
1003 token-ring-default   act/unsup
1004 fddinet-default      act/unsup
1005 trnet-default        act/unsup
2000 VLAN2000            active
S1#

```

Note

A Cisco Catalyst 2960 switch can support up to 255 normal-range and extended-range VLANs. However, the number of VLANs configured affects the performance of the switch hardware.

Dynamic Trunking Protocol (2.1.4)

DTP simplifies the negotiation of trunk links between two switches. The focus of this topic is on how to configure DTP.

Introduction to DTP (2.1.4.1)

Ethernet trunk interfaces support different trunking modes. An interface can be set to trunking or non-trunking, or it can be set to negotiate trunking with the neighbor interface. Trunk negotiation is managed by DTP, which operates on a point-to-point basis only, between network devices.

DTP is a Cisco proprietary protocol that is automatically enabled on Catalyst 2960 and Catalyst 3560 Series switches. Switches from other vendors do not support DTP. DTP manages trunk negotiation only if the port on the neighbor switch is configured in a trunk mode that supports DTP.

Caution

Some internetworking devices might forward DTP frames improperly, which can cause misconfigurations. To avoid this, turn off DTP on interfaces on a Cisco switch connected to devices that do not support DTP.

The default DTP configuration for Cisco Catalyst 2960 and 3560 switches is dynamic auto, as shown in Figure 2-7 on interface F0/3 of switches S1 and S3.

To enable trunking from a Cisco switch to a device that does not support DTP, use the **switchport mode trunk** and **switchport nonegotiate** interface configuration mode commands. This causes the interface to become a trunk but not generate DTP frames.

In Figure 2-8, the link between switches S1 and S2 becomes a trunk because the F0/1 ports on switches S1 and S2 are configured to ignore all DTP advertisements and to come up in and stay in trunk port mode.

The F0/3 ports on switches S1 and S3 are set to dynamic auto, so the negotiation results in the access mode state. This creates an inactive trunk link. When configuring a port to be in trunk mode, use the **switchport mode trunk** command. Then there is no ambiguity about which state the trunk is in; it is always on.

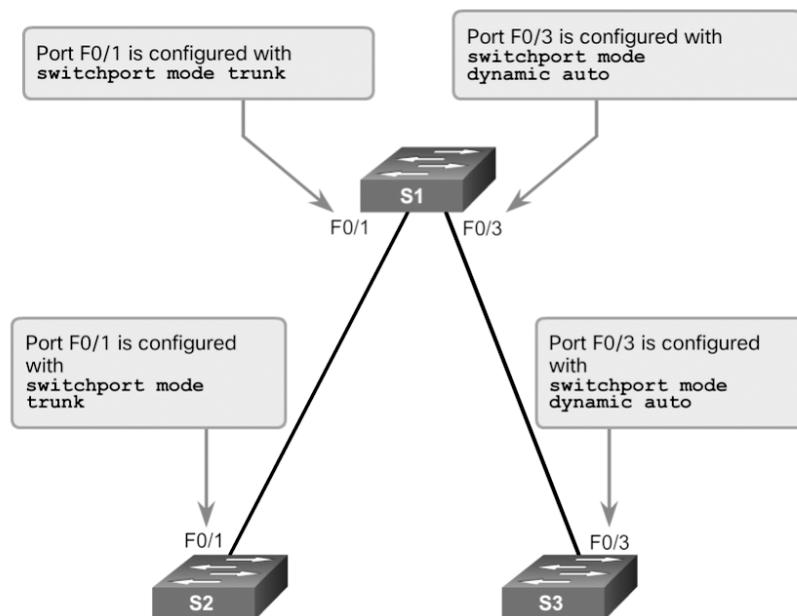


Figure 2-7 Initial DTP Configuration

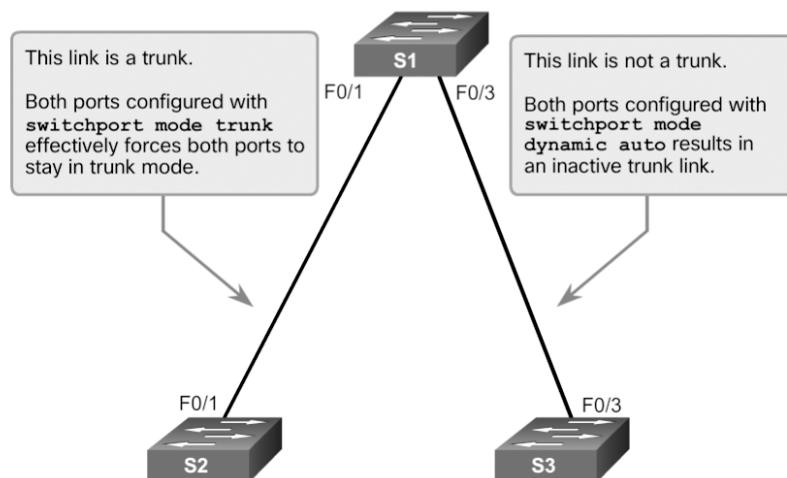


Figure 2-8 DTP Interaction Results

Negotiated Interface Modes (2.1.4.2)

Ethernet interfaces on Catalyst 2960 and Catalyst 3560 Series switches support different trunking modes with the help of DTP:

- **switchport mode access**—Puts the interface (access port) into permanent non-trunking mode and negotiates to convert the link into a non-trunk link. The interface becomes an access port, regardless of whether the neighboring interface is a trunk port.
- **switchport mode dynamic auto**—This is the default switchport mode for all Ethernet interfaces. It makes the port able to convert the link to a trunk link. The port becomes a trunk if the neighboring interface is set to trunk or desirable mode. It does not trunk if the interface is also set to dynamic auto.
- **switchport mode dynamic desirable**—Makes the interface actively attempt to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to trunk, desirable, or dynamic auto mode. Note that this is the default switchport mode on older Catalyst switches, such as the Catalyst 2950 and 3550 Series switches.
- **switchport mode trunk**—Puts the interface into permanent trunking mode and negotiates to convert the neighboring link into a trunk link. The interface becomes a trunk interface even if the neighboring interface is not a trunk interface.
- **switchport nonegotiate**—Prevents the interface from generating DTP frames. You can use this command only when the interface switchport mode is access or trunk. You must manually configure the neighboring interface as a trunk interface to establish a trunk link.

Table 2-11 illustrates the results of the DTP configuration options on opposite ends of a trunk link connected to Catalyst 2960 switch ports.

Table 2-11 DTP–Negotiated Interface Modes

	dynamic auto	dynamic desirable	trunk	access
dynamic auto	Access	Trunk	Trunk	Access
dynamic desirable	Trunk	Trunk	Trunk	Access
trunk	Trunk	Trunk	Trunk	Limited connectivity
access	Access	Access	Limited connectivity	Access

Configure trunk links statically whenever possible. The default DTP mode is dependent on the Cisco IOS Software version and on the platform. To determine the current DTP mode, issue the **show dtp interface** command, as shown in Example 2-19.

Example 2-19 Verifying DTP Mode

```

S1# show dtp interface f0/1
DTP information for FastEthernet0/1:
TOS/TAS/TNS:                TRUNK/ON/TRUNK
TOT/TAT/TNT:                802.1Q/802.1Q/802.1Q
Neighbor address 1:         0CD996D23F81
Neighbor address 2:         000000000000
Hello timer expiration (sec/state): 12/RUNNING
Access timer expiration (sec/state): never/STOPPED
Negotiation timer expiration (sec/state): never/STOPPED
Multidrop timer expiration (sec/state): never/STOPPED
FSM state:                  S6:TRUNK
# times multi & trunk       0
Enabled:                    yes
In STP:                     no

<output omitted>

```

Note

A general best practice is to set the interface to **trunk** and **nonegotiate** when a trunk link is required. On links where trunking is not intended, DTP should be turned off.

Interactive Graphic**Activity 2.1.4.3: Predict DTP Behavior**

Refer to the online course to complete this activity.

Packet Tracer Activity**Packet Tracer 2.1.4.4: Configure VTP and DTP**

In this activity, you will configure a switched environment in which trunks are negotiated and formed via DTP, and VLAN information is propagated automatically through a VTP domain.

**Lab 2.1.4.5: Configure Extended VLANs, VTP, and DTP**

Refer to *Scaling Networks v6 Labs & Study Guide* and the online course to complete this activity.

In this lab, you will complete the following objectives:

- Build the Network and Configure Basic Device Settings
- Use Dynamic Trunking Protocol (DTP) to Form Trunk Links
- Configure VLAN Trunking Activity 1.0.1.2: Do We Really Need a Map?

Troubleshoot Multi-VLAN Issues (2.2)

VLANs are susceptible to specific types of problems in a campus LAN. Most of these problems are related to *inter-VLAN routing* configuration issues, IP addressing issues, VTP issues, and DTP issues.

In this section, you will learn how to troubleshoot issues in an inter-VLAN routing environment.

Inter-VLAN Configuration Issues (2.2.1)

The focus of this topic is on how to troubleshoot common inter-VLAN configuration issues.

Deleting VLANs (2.2.1.1)

On occasion, you have to remove a VLAN from the VLAN database. When deleting a VLAN from a switch that is in VTP server mode, the VLAN is removed from the VLAN database for all switches in the VTP domain. When you delete a VLAN from a switch that is in VTP transparent mode, the VLAN is deleted only on that specific switch or switch stack.

Note

You cannot delete the default VLANs (that is, VLANs 1 and 1002 through 1005).

The following scenario illustrates how to delete a VLAN. Assume that S1 has VLANs 10, 20, and 99 configured, as shown in Example 2-20. Notice that VLAN 99 is assigned to ports Fa0/18 through Fa0/24.

Example 2-20 Verifying the VLAN Configuration on S1

```
S1# show vlan brief
```

VLAN Name	Status	Ports
1 default	active	Fa0/1, Fa0/2, Fa0/3, Fa0/4 Fa0/5, Fa0/6, Fa0/7, Fa0/8 Fa0/9, Fa0/10, Fa0/11, Fa0/12 Fa0/13, Fa0/14, Fa0/15, Fa0/16 Fa0/17, Gig0/1, Gig0/2
10 VLAN0010	active	
20 VLAN0020	active	
99 VLAN0099	active	Fa0/18, Fa0/19, Fa0/20, Fa0/21 Fa0/22, Fa0/23, Fa0/24


```

Fa0/9, Fa0/10, Fa0/11, Fa0/12
Fa0/13, Fa0/14, Fa0/15, Fa0/16
Fa0/17, Gig0/1, Gig0/2

10 VLAN0010 active
20 VLAN0020 active
1002 fddi-default active
1003 token-ring-default active
1004 fddinet-default active
1005 trnet-default active
S1#

```

Switch Port Issues (2.2.1.2)

Several common switch misconfigurations can arise when configuring routing between multiple VLANs.

When configuring a *legacy inter-VLAN routing* solution (also referred to as traditional inter-VLAN routing), ensure that the switch ports that connect to the router interfaces are configured with the correct VLANs. If a switch port is not configured for the correct VLAN, devices on that VLAN are unable to send data to the other VLANs.

For example, refer to the topology in Figure 2-9.

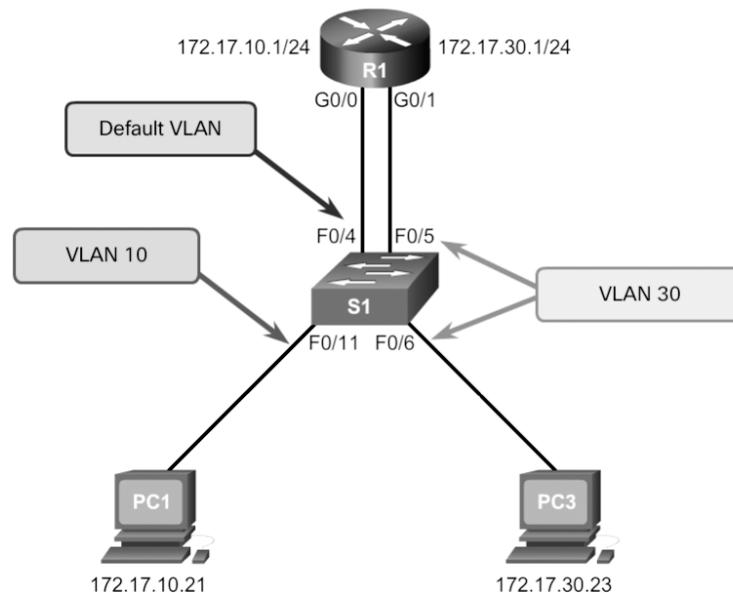


Figure 2-9 Legacy Inter-VLAN Routing Issue—Scenario 1

PC1 and router R1 interface G0/0 are configured to be on the same logical subnet, as indicated by their IPv4 address assignment. However, the S1 F0/4 port that connects to the R1 G0/0 interface has not been configured and therefore remains in the default VLAN. Because R1 is on a different VLAN than PC1, they are unable to communicate.

To correct this problem, port F0/4 on switch S1 must be in access mode (**switchport access mode**) and assigned to VLAN 20 (**switchport access vlan 20**). When this is configured, PC1 can communicate with the R1 G0/0 interface and be routed to other VLANs connected to R1.

When a *router-on-a-stick inter-VLAN routing* solution is implemented, ensure that interconnecting interfaces are configured properly as trunks. For example, refer to the topology in Figure 2-10. R1 has been configured with subinterfaces and trunking enabled. However, the F0/5 port on S1 has not been configured as a trunk and is left in the default VLAN. As a result, the router is unable to route between VLANs because each of its configured subinterfaces is unable to send or receive VLAN-tagged traffic.

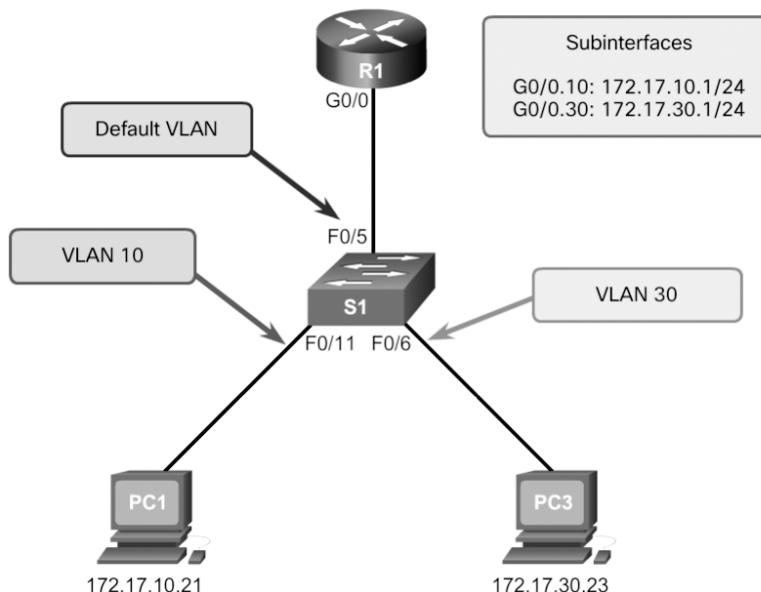


Figure 2-10 Router-on-a-Stick Inter-VLAN Routing Issue—Scenario 2

To correct this problem, issue the **switchport mode trunk** interface configuration mode command on the F0/5 interface of S1. This converts the interface to a trunk port, allowing a trunk to be established between R1 and S1. When the trunk is successfully established, devices connected to each of the VLANs are able to

communicate with the subinterface assigned to their VLAN, thereby enabling inter-VLAN routing.

Another VLAN issue is if a link goes down or fails. A downed interswitch link disrupts the inter-VLAN routing process.

For example, refer to the topology in Figure 2-11. Notice that the trunk link between S1 and S2 is down. Because there is no redundant connection or path between the devices, all devices connected to S2 are unable to reach router R1. Therefore, all devices connected to S2 are unable to route to other VLANs through R1.

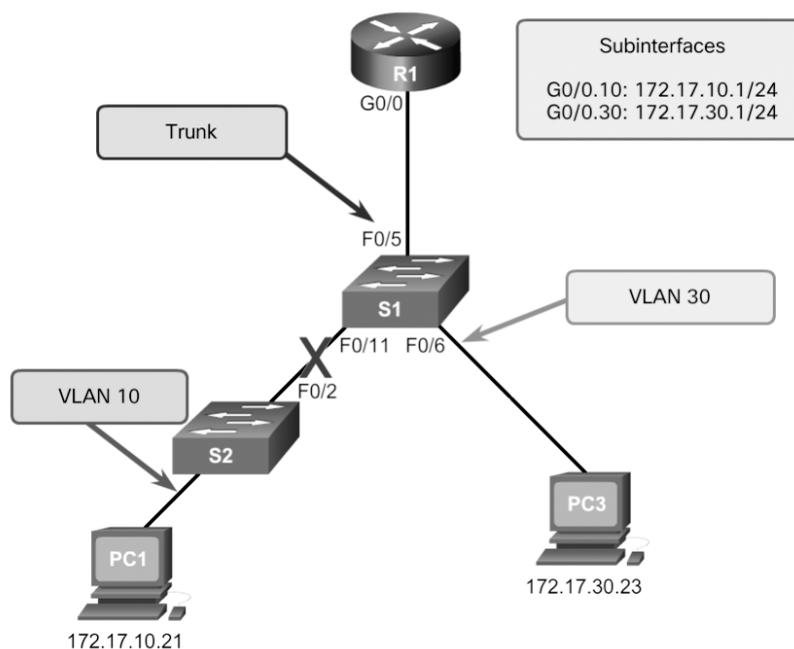


Figure 2-11 Failed Interswitch Link Issue—Scenario 3

The solution for this problem is not configuration related, but instead is a LAN design issue. The network design should include redundant links and alternate paths to reduce the risk of failed interswitch links.

Verify Switch Configuration (2.2.1.3)

When an inter-VLAN problem is suspected with a switch configuration, use verification commands to examine the configuration and identify the problem. Knowing the right verification commands to use helps you quickly identify issues.

The **show interfaces *interface-id* switchport** command is useful for identifying VLAN assignment and port configuration issues.

For example, assume that the Fa0/4 port on switch S1 should be an access port configured in VLAN 10. To verify the correct port settings, use the **show interfaces interface-id switchport** command, as shown in Example 2-22.

Example 2-22 Verifying the Current Interface Settings

```
S1# show interfaces FastEthernet 0/4 switchport
Name: Fa0/4
Switchport: Enabled
Administrative Mode: static access
Operational Mode: up
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
<output omitted>
S1#
```

The top highlighted area confirms that port F0/4 on switch S1 is in access mode. The bottom highlighted area confirms that port F0/4 is not set to VLAN 10 but instead is still set to the default VLAN. To correct this issue, the F0/4 port would have to be configured with the **switchport access vlan 10** command.

The **show running-config interface** is a useful command for identifying how an interface is configured. For example, assume that a device configuration changed, and the trunk link between R1 and S1 has stopped. Example 2-23 displays the output of the **show interfaces interface_id switchport** and **show running-config interface** verification commands.

Example 2-23 Switch IOS Commands

```
S1# show interface f0/4 switchport
Name: Fa0/4
Switchport: Enabled
Administrative Mode: static access
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
<output omitted>
S1#
S1# show run interface fa0/4
interface FastEthernet0/4
switchport mode access
S1#
```

The top highlighted area reveals that port F0/4 on switch S1 is in access mode. It should be in trunk mode. The bottom highlighted area also confirms that port F0/4 has been configured for access mode.

To correct this issue, the Fa0/4 port must be configured with the **switchport mode trunk** command.

Interface Issues (2.2.1.4)

Many inter-VLAN issues are physical layer (Layer 1) errors. For example, one of the most common configuration errors is to connect the physical router interface to the wrong switch port.

Refer to the legacy inter-VLAN solution in Figure 2-12. The R1 G0/0 interface is connected to the S1 F0/9 port. However, the F0/9 port is configured for the default VLAN, not VLAN 10. This prevents PC1 from being able to communicate with its default gateway, the router interface. Therefore, the PC is unable to communicate with any other VLANs, such as VLAN 30.

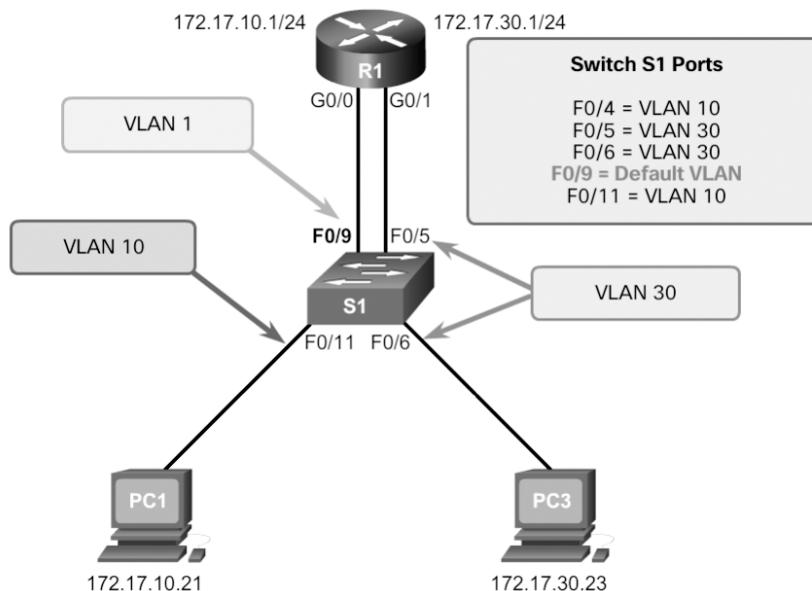


Figure 2-12 Layer 1 Issue

The port should have been connected to the Fa0/4 port on S1. Connecting the R1 interface G0/0 to switch S1 port F0/4 puts the interface in the correct VLAN and allows inter-VLAN routing.

Note that an alternative solution would be to change the VLAN assignment of port F0/9 to VLAN 10.

Verify Routing Configuration (2.2.1.5)

With router-on-a-stick configurations, a common problem is assigning the wrong VLAN ID to the subinterface.

For example, as shown in Figure 2-13, router R1 subinterface G0/0.10 has been configured in VLAN 100 instead of VLAN 10. This prevents devices configured on VLAN 10 from communicating with subinterface G0/0.10 and from being able to send data to other VLANs on the network.

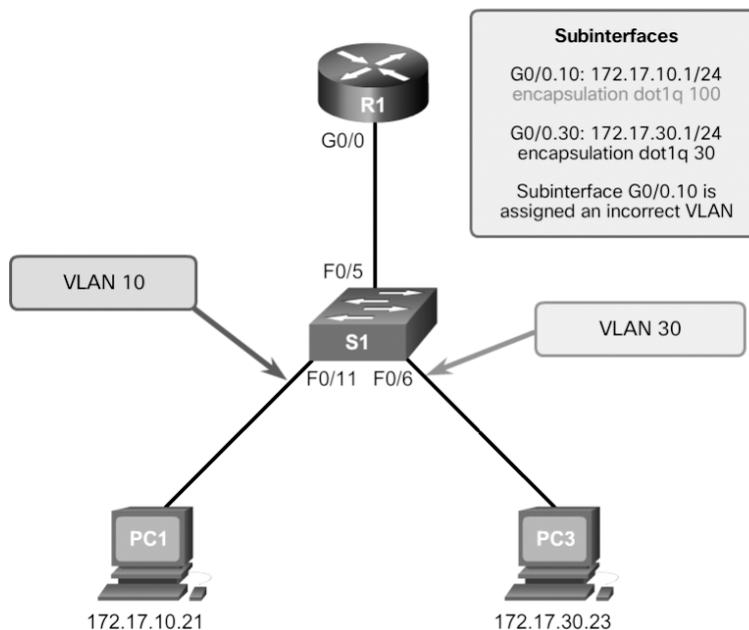


Figure 2-13 Router Configuration Issue

Using the `show interfaces` and the `show running-config interface` commands can be useful in troubleshooting this type of issue, as shown in Example 2-24.

Example 2-24 Verifying Router Configuration

```
R1# show interface G0/0.10
GigabitEthernet0/0.10 is up, line protocol is down (disabled)
Encapsulation 802.1Q Virtual LAN, Vlan ID 100
ARP type: ARPA, ARP Timeout 04:00:00,
Last clearing of "show interface" counters never

<Output omitted>
```

```
R1#  
R1# show run interface G0/0.10  
interface GigabitEthernet0/0.10  
encapsulation dot1Q 100  
ip address 172.17.10.1 255.255.255.0  
R1#
```

The **show interfaces** command produces a lot of output, sometimes making it difficult to see the problem. However, the top highlighted section of Example 2-24 shows that the subinterface G0/0.10 on router R1 uses VLAN 100.

The **show running-config** command confirms that subinterface G0/0.10 on router R1 has been configured to allow access to VLAN 100 traffic and not VLAN 10.

To correct this problem, configure subinterface G0/0.10 to be on the correct VLAN by using the **encapsulation dot1q 10** subinterface configuration mode command. Once this is configured, the subinterface performs inter-VLAN routing to users on VLAN 10.

IP Addressing Issues (2.2.2)

VLAN issues could also be caused by misconfigured network or IP address information. The focus of this topic is on how to troubleshoot common IP addressing issues in an inter-VLAN routed environment.

Errors with IP Addresses and Subnet Masks (2.2.2.1)

VLANs correspond to unique subnets on the network. For inter-VLAN routing to operate, a router must be connected to all VLANs, either by separate physical interfaces or by subinterfaces.

Each interface or subinterface must be assigned an IP address that corresponds to the subnet to which it is connected. This permits devices on the VLAN to communicate with the router interface and enables the routing of traffic to other VLANs connected to the router.

The following are examples of possible inter-VLAN routing problems related to IP addressing errors.

In Figure 2-14, router R1 has been configured with an incorrect IPv4 address on interface G0/0, preventing PC1 from being able to communicate with router R1 on VLAN 10.

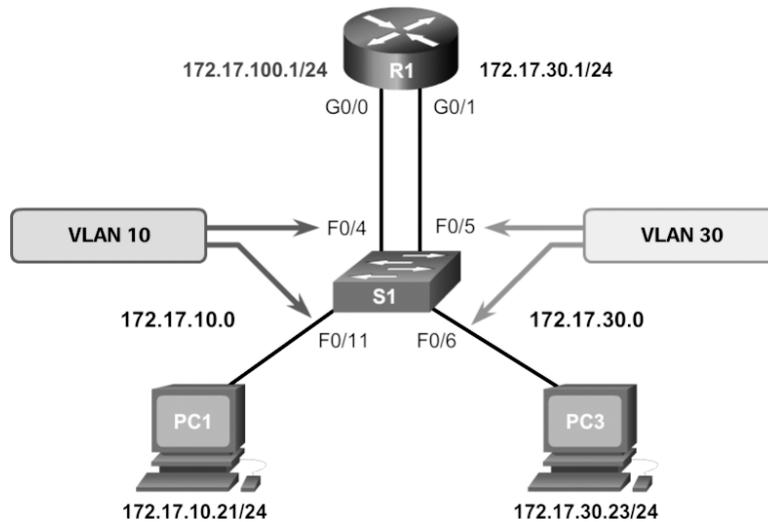


Figure 2-14 IP Addressing Issues—Scenario 1

To correct this problem, configure the `ip address 172.17.10.1 255.255.255.0` command on the R1 G0/0 interface. Once this is configured, PC1 can use the router interface as a default gateway for accessing other VLANs.

Another problem is illustrated in Figure 2-15. In this example, PC1 has been configured with an incorrect IPv4 address for the subnet associated with VLAN 10, preventing it from being able to communicate with router R1 on VLAN 10.

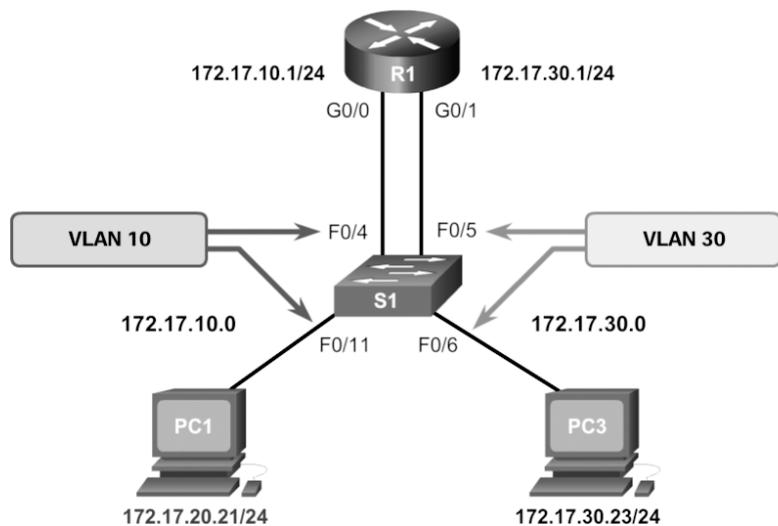


Figure 2-15 IP Addressing Issues—Scenario 2

To correct this problem, assign the correct IPv4 address to PC1.

Another problem is shown in Figure 2-16. In this example, PC1 cannot send traffic to PC3. The reason is that PC1 has been configured with the incorrect subnet mask, /16, instead of the correct /24 mask. The /16 mask makes PC1 assume that PC3 is on the same subnet. Therefore, PC1 never forwards traffic destined to PC3 to its default gateway, R1.

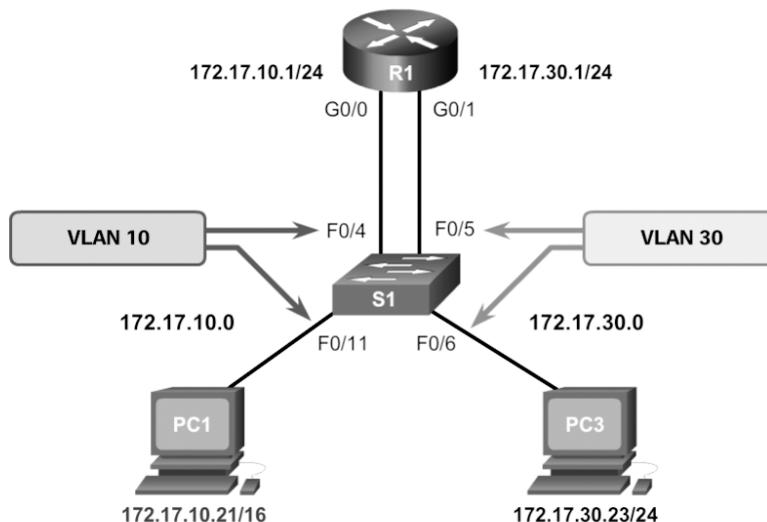


Figure 2-16 IP Addressing Issues—Scenario 3

To correct this problem, change the subnet mask on PC1 to 255.255.255.0.

Verifying IP Address and Subnet Mask Configuration Issues (2.2.2.2)

When troubleshooting addressing issues, ensure that the subinterface is configured with the correct address for that VLAN. Each interface or subinterface must be assigned an IP address corresponding to the subnet to which it is connected. A common error is to incorrectly configure an IP address on a subinterface.

Example 2-25 displays the output of the **show running-config** and **show ip interface** commands. The highlighted areas show that subinterface G0/0.10 on router R1 has IPv4 address 172.17.20.1. However, this is the wrong IP address for this subinterface, and instead it should be configured for VLAN 10.

Example 2-25 Using Commands to Discover the Configuration Issues

```
R1# show run
Building configuration...
<output omitted>
!
interface GigabitEthernet0/0
  no ip address
  duplex auto
  speed auto
!
interface GigabitEthernet0/0.10
  encapsulation dot1Q 10
  ip address 172.17.20.1 255.255.255.0
!
interface GigabitEthernet0/0.30
<output omitted>
R1#
R1# show ip interface
<output omitted>
GigabitEthernet0/0.10 is up, line protocol is up
  Internet address is 172.17.20.1/24
  Broadcast address is 255.255.255.255
<output omitted>
R1#
```

To correct this problem, change the IP address of subinterface G0/0.10 to 172.17.10.1/24.

Sometimes it is the end-user device that is improperly configured. For example, Figure 2-17 displays the IPv4 configuration of PC1. The configured IPv4 address is 172.17.20.21/24. However, in this scenario, PC1 should be in VLAN 10, with address 172.17.10.21/24.

To correct this problem, correct the IP address of PC1.

Note

In the examples in this chapter, the subinterface numbers always match the VLAN assignment. This is not a configuration requirement but instead has been done intentionally to make it easier to manage inter-VLAN configuration.

**Interactive
Graphic****Activity 2.2.2.3: Identify the Troubleshooting Command for an Inter-VLAN Routing Issue**

Refer to the online course to complete this activity.

```
Packet Tracer PC Command Line 1.0
PC1> ip config
Invalid Command.

PC1> ipconfig

IP Address.....: 172.17.20.21
Subnet Mask.....: 255.255.255.0
Default Gateway.....: 172.17.10.1

PC1>
```

This PC1 should be in the VLAN 10 subnet
So this should be: 172.17.10.21 with a subnet mask of
255.255.255.0

Figure 2-17 PC IP Addressing Issue

Packet Tracer
Activity

Packet Tracer 2.2.2.4: Troubleshooting Inter-VLAN Routing

In this activity, you will troubleshoot connectivity problems caused by improper configurations related to VLANs and inter-VLAN routing.



Lab 2.2.2.5: Troubleshooting Inter-VLAN Routing

Refer to *Scaling Networks v6 Labs & Study Guide* and the online course to complete this activity.

In this lab, you will complete the following objectives:

- Part 1: Build the Network and Load Device Configurations
- Part 2: Troubleshoot the Inter-VLAN Routing Configuration
- Part 3: Verify VLAN Configuration, Port Assignment, and Trunking
- Part 4: Test Layer 3 Connectivity

VTP and DTP Issues (2.2.3)

The focus of this topic is on how to troubleshoot common VTP and DTP issues in an inter-VLAN routed environment.

Troubleshoot VTP Issues (2.2.3.1)

Several issues can arise from an invalid VTP configuration. Common problems with VTP are listed in Table 2-12.

Table 2-12 Common VTP-Related Issues

VTP Problem	Description
Incompatible VTP versions	<p>VTP versions are incompatible with each other.</p> <p>Ensure that all switches are capable of supporting the required VTP version.</p>
Incorrect VTP domain name	<p>An improperly configured VTP domain affects VLAN synchronization between switches, and if a switch receives the wrong VTP advertisement, the switch discards the message.</p> <p>To avoid incorrectly configuring a VTP domain name, set the VTP domain name on only one VTP server switch.</p> <p>All other switches in the same VTP domain will accept and automatically configure their VTP domain name when they receive the first VTP summary advertisement.</p>
Incorrect VTP mode	<p>If all switches in the VTP domain are set to client mode, you cannot create, delete, or manage VLANs.</p> <p>To avoid losing all VLAN configurations in a VTP domain, configure two switches as VTP servers.</p>
Invalid VTP authentication	<p>If VTP authentication is enabled, switches must all have the same password configured to participate in VTP.</p> <p>Ensure that the password is manually configured on all switches in the VTP domain.</p>
Incorrect configuration revision number	<p>If a switch with the same VTP domain name but a higher configuration number is added to the domain, invalid VLANs can be propagated and/or valid VLANs can be deleted.</p> <p>The solution is to reset each switch to an earlier configuration and then reconfigure the correct VLANs.</p> <p>Before adding a switch to a VTP-enabled network, reset the revision number on the switch to 0 by assigning it to a false VTP domain and then reassigning it to the correct VTP domain name.</p>

Troubleshoot DTP Issues (2.2.3.2)

Trunking issues are associated with incorrect configurations. As outlined Table 2-13, three common problems are associated with trunks.

Table 2-13 Common Trunk-Related Issues

DTP Issues	Description
Trunk mode mismatches	<p>One trunk port is configured with trunk mode “off” and the other with trunk mode “on.”</p> <p>This configuration error causes the trunk link to stop working.</p> <p>Correct the situation by shutting down the interface, correcting the DTP mode settings, and re-enabling the interface.</p>
Invalid allowed VLANs on trunks	<p>The list of allowed VLANs on a trunk has not been updated with the current VLAN trunking requirements.</p> <p>In this situation, unexpected traffic or no traffic is being sent over the trunk.</p> <p>Configure the correct VLANs that are allowed on the trunk.</p>
Native VLAN mismatches	<p>When native VLANs do not match, the switches will generate informational messages letting you know about the problem.</p> <p>Ensure that both sides of a trunk link are using the same native VLAN.</p>

Packet Tracer
 Activity

Packet Tracer 2.2.3.3: Troubleshoot VTP and DTP Issues

In this activity, you will troubleshoot a switched environment where trunks are negotiated and formed via DTP and where VLAN information is propagated automatically through a VTP domain.

Layer 3 Switching (2.3)

A router-on-a-stick inter-VLAN solution is relatively easy to configure and suitable in a smaller network. An alternative solution is to use Layer 3 switches to perform inter-VLAN routing.

In this section, you will learn how to implement inter-VLAN routing using Layer 3 switching to forward data in a small to medium-sized business LAN.

Layer 3 Switching Operation and Configuration (2.3.1)

The focus of this topic is on how to configure inter-VLAN routing using Layer 3 switching.

Introduction to Layer 3 Switching (2.3.1.1)

Inter-VLAN routing using the router-on-a-stick method was simple to implement because routers were usually available in every network. However, as shown in Figure 2-18, most modern enterprise networks use a *Layer 3 inter-VLAN routing* solution. This requires the use of multilayer switches to achieve high packet-processing rates using hardware-based switching.

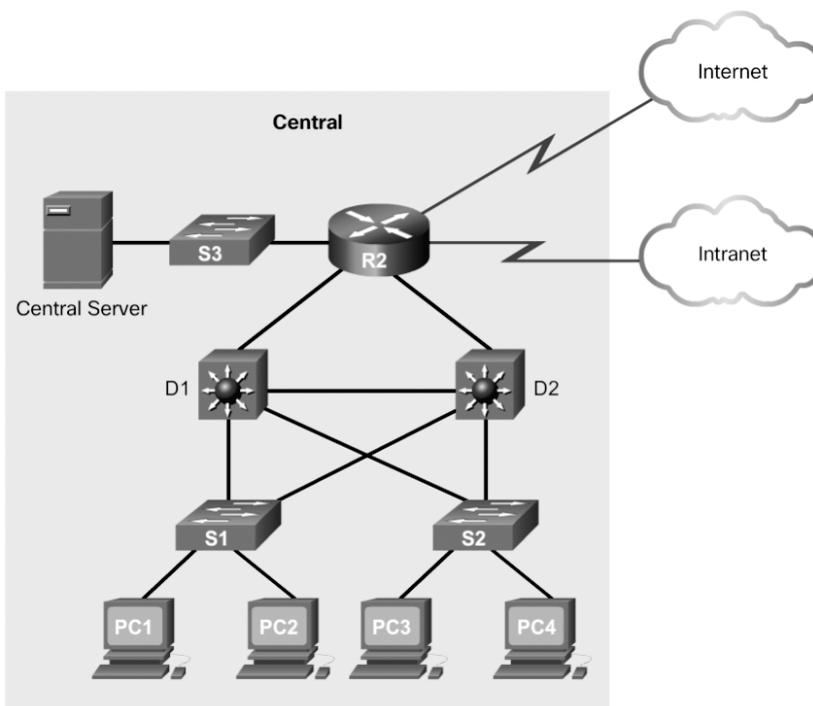


Figure 2-18 Layer 3 Switching Topology

Layer 3 switches usually have packet-switching throughputs in the millions of packets per second (pps), whereas traditional routers provide packet switching in the range of 100,000 pps to more than 1 million pps.

All Catalyst multilayer switches support the following types of Layer 3 interfaces:

- **Routed port**—A pure Layer 3 interface similar to a physical interface on a Cisco IOS router.

- *Switch virtual interface (SVI)*—A virtual VLAN interface for inter-VLAN routing. In other words, SVIs are virtual-routed VLAN interfaces.

High-performance switches, such as the Catalyst 6500 and Catalyst 4500, perform almost every function involving OSI Layer 3 and higher using hardware-based switching that is based on *Cisco Express Forwarding*.

All Layer 3 Cisco Catalyst switches support routing protocols, but several models of Catalyst switches require enhanced software for specific routing protocol features.

Note

Catalyst 2960 Series switches running IOS Release 12.2(55) or later support static routing.

Layer 3 Catalyst switches use different default settings for interfaces. For example:

- Catalyst 3560, 3650, and 4500 families of distribution layer switches use Layer 2 interfaces by default.
- Catalyst 6500 and 6800 families of core layer switches use Layer 3 interfaces by default.

Depending on which Catalyst family of switches is used, the **switchport** or **no switchport** interface configuration mode command might be present in the running config or startup configuration files.

Inter-VLAN Routing with Switch Virtual Interfaces (2.3.1.2)

In the early days of switched networks, switching was fast (often at hardware speed, meaning the speed was equivalent to the time it took to physically receive and forward frames onto other ports), and routing was slow (because it had to be processed in software). This prompted network designers to extend the switched portion of the network as much as possible. Access, distribution, and core layers were often configured to communicate at Layer 2. This topology created loop issues. To solve these issues, spanning-tree technologies were used to prevent loops while still enabling flexibility and redundancy in interswitch connections.

However, as network technologies have evolved, routing has become faster and cheaper. Today, routing can be performed at wire speed. One consequence of this evolution is that routing can be transferred to the core and the distribution layers (and sometimes even the access layer) without impacting network performance.

Many users are in separate VLANs, and each VLAN is usually a separate subnet. Therefore, it is logical to configure the distribution switches as Layer 3 gateways for the users of each access switch VLAN. This implies that each distribution switch must have IP addresses matching each access switch VLAN. This can be achieved by SVIs and routed ports.

For example, refer to the topology in Figure 2-19.

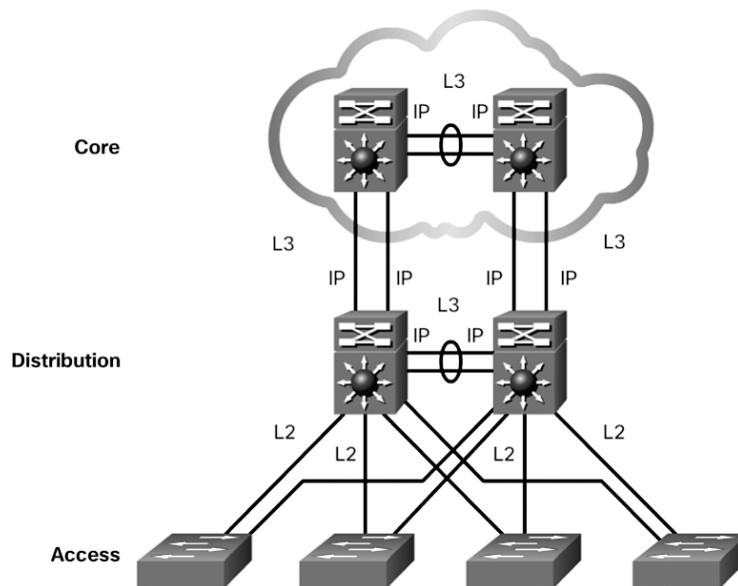


Figure 2-19 Switched Network Design

Layer 3 (routed) ports are normally implemented between the distribution layer and the core layer. Therefore, the core layer and distribution layer switches in the figure are interconnected using Layer 3 IP addressing.

The distribution layer switches are connected to the access layer switches using Layer 2 links. The network architecture depicted is not dependent on the spanning tree protocol (STP) because there are no physical loops in the Layer 2 portion of the topology.

Inter-VLAN Routing with Switch Virtual Interfaces (Con't.) (2.3.1.3)

The topologies in Figure 2-20 compare configuring inter-VLAN routing on a router and on a Layer 3 switch.

An SVI is a virtual interface that is configured within a multilayer switch, as shown in the figure. An SVI can be created for any VLAN that exists on the switch. An SVI is considered to be virtual because there is no physical port dedicated to the interface. It can perform the same functions for the VLAN as a router interface would, and it can be configured in much the same way as a router interface (that is, IP address, inbound/outbound ACLs, and so on). The SVI for the VLAN provides Layer 3 processing for packets to or from all switch ports associated with that VLAN.

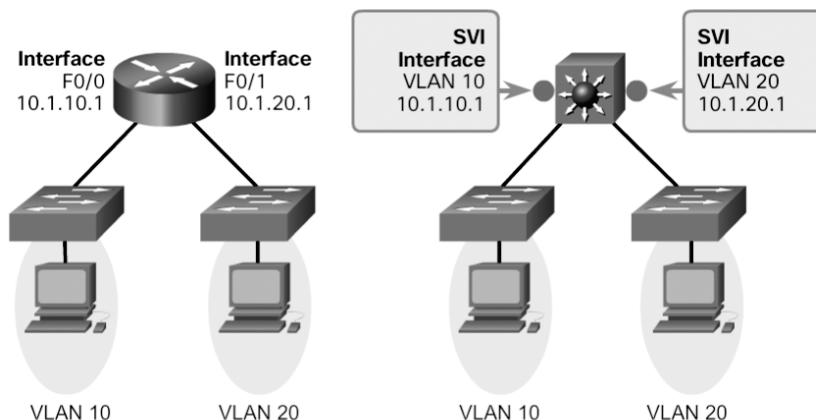


Figure 2-20 Switch Virtual Interface

By default, an SVI is created for the default VLAN (VLAN 1) to permit remote switch administration. Additional SVIs must be explicitly created. SVIs are created the first time the VLAN interface configuration mode is entered for a particular VLAN SVI, such as when the **interface vlan 10** command is entered. The VLAN number used corresponds to the VLAN tag associated with data frames on an 802.1Q encapsulated trunk or to the VLAN ID (VID) configured for an access port. When creating an SVI as a gateway for VLAN 10, name the SVI interface VLAN 10. Configure and assign an IP address to each VLAN SVI.

Whenever the SVI is created, ensure that the particular VLAN is present in the VLAN database. For the example shown in Figure 2-20, the switch should have VLAN 10 and VLAN 20 present in the VLAN database; otherwise, the SVI interface stays down.

The following are some of the reasons to configure SVI:

- To provide a gateway for a VLAN so that traffic can be routed into or out of that VLAN
- To provide Layer 3 IP connectivity to the switch
- To support routing protocol and bridging configurations

The only disadvantage of SVIs is that multilayer switches are expensive. The following are some of the advantages of SVIs:

- It is much faster than router-on-a-stick because everything is hardware switched and routed.
- There is no need for external links from the switch to the router for routing.

- It is not limited to one link. Layer 2 EtherChannels can be used between the switches to get more bandwidth.
- Latency is much lower because data does not need to leave the switch in order to be routed to a different network.

Inter-VLAN Routing with Routed Ports (2.3.1.4)

Routed Ports and Access Ports on a Switch

A routed port is a physical port that acts similarly to an interface on a router. Unlike an access port, a routed port is not associated with a particular VLAN. A routed port behaves like a regular router interface. Also, because Layer 2 functionality has been removed, Layer 2 protocols, such as STP, do not function on a routed interface. However, some protocols, such as LACP and EtherChannel, do function at Layer 3.

Unlike Cisco IOS routers, routed ports on a Cisco IOS switch do not support subinterfaces.

Routed ports are used for point-to-point links. Routed ports can be used for connecting WAN routers and security devices, for example. In a switched network, routed ports are mostly configured between switches in the core and distribution layers. Figure 2-21 illustrates an example of routed ports in a campus switched network.

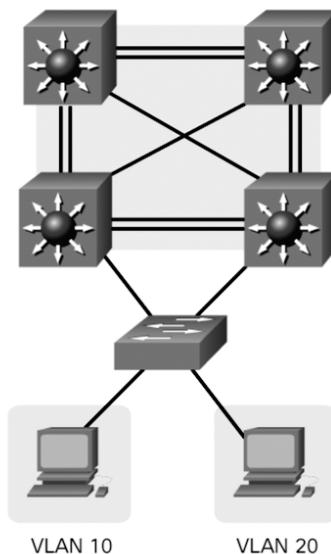


Figure 2-21 Routed Ports

To configure routed ports, use the **no switchport** interface configuration mode command on the appropriate ports. For example, the default configuration of the interfaces on Catalyst 3560 switches is as Layer 2 interfaces, so they must be manually configured as routed ports. In addition, assign an IP address and other Layer 3 parameters as necessary. After assigning the IP address, verify that IP routing is globally enabled and that applicable routing protocols are configured.

Note

Routed ports are not supported on Catalyst 2960 Series switches.

Packet Tracer Activity

Packet Tracer 2.3.1.5: Configure Layer 3 Switching and Inter-VLAN Routing

In this activity, you will configure Layer 3 switching and inter-VLAN routing on a Cisco 3560 switch.

Troubleshoot Layer 3 Switching (2.3.2)

The focus of this topic is on how to troubleshoot inter-VLAN routing in a Layer 3 switched environment.

Layer 3 Switch Configuration Issues (2.3.2.1)

The issues common to legacy inter-VLAN routing and router-on-a-stick inter-VLAN routing also manifest in the context of Layer 3 switching.

Table 2-14 lists items that should be checked for accuracy when troubleshooting inter-VLAN routing issues.

Table 2-14 Common Layer 3 Switching Issues

Check	Description
VLANs	VLANs must be defined across all the switches. VLANs must be enabled on the trunk ports. Ports must be in the right VLANs.
SVIs	SVIs must have the correct IP addresses or subnet masks. SVIs must be up. Each SVI must match the VLAN number.

Check	Description
Routing	Routing must be enabled. Each interface or network should be added to the routing protocol or static routes entered, where appropriate.
Hosts	Hosts must have the correct IP address or subnet mask. Hosts must have a default gateway associated with an SVI or a routed port.

To troubleshoot the Layer 3 switching problems, be familiar with the implementation and design layout of the topology, such as the one shown in Figure 2-22.

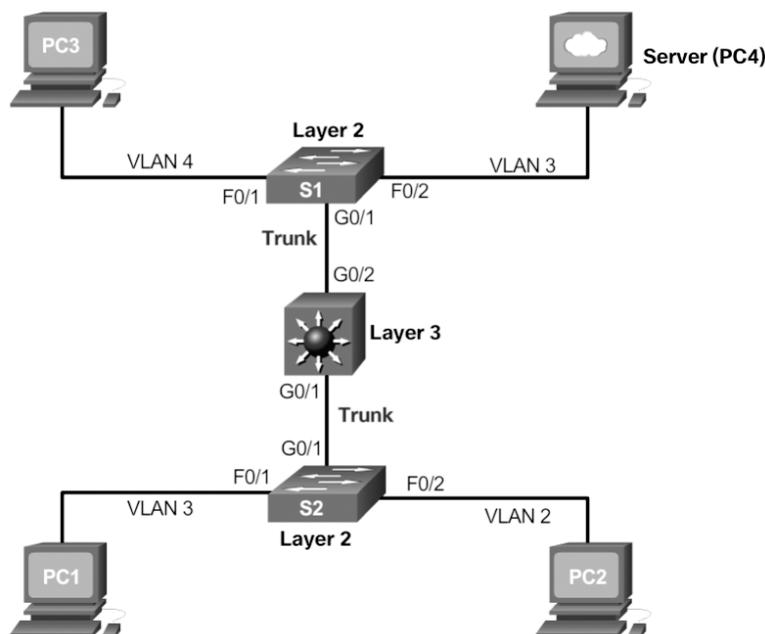


Figure 2-22 Layer 3 Switch Configuration Issues Topology

Example: Troubleshooting Layer 3 Switching (2.3.2.2)

Company XYZ is adding a new floor, floor 5, to the network (see Figure 2-23).

The current requirement is to make sure the users on floor 5 can communicate with users on other floors. Currently, users on floor 5 cannot communicate with users on other floors. The following is an implementation plan to install a new VLAN for users on floor 5 and to ensure the VLAN is routing to other VLANs.

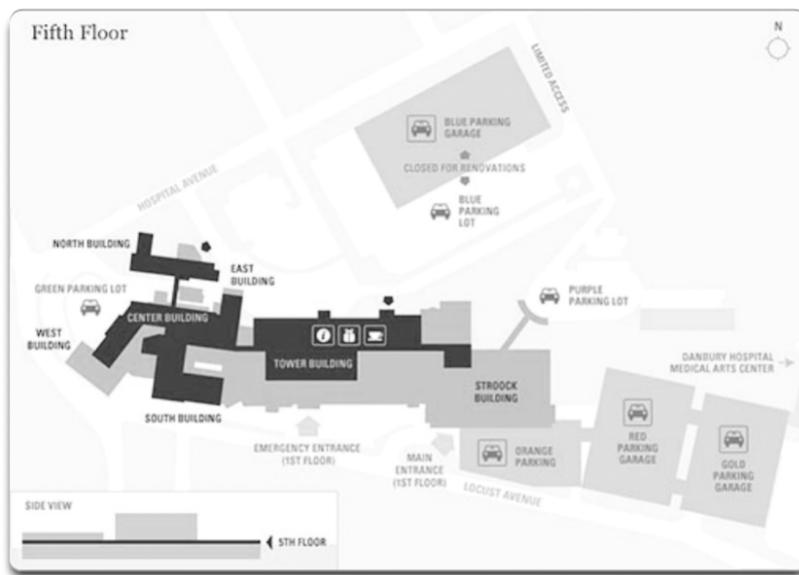


Figure 2-23 Company XYZ Floor Plan – Fifth Floor

There are four steps to implementing the new VLAN:

- Step 1.** Create a new VLAN on the fifth floor switch and on the distribution switches. Name this VLAN 500.
- Step 2.** Identify the ports needed for the users and switches. Set the **switchport access vlan** command to **500** and ensure that the trunk between the distribution switches is properly configured and that VLAN 500 is allowed on the trunk.
- Step 3.** Create an SVI interface on the distribution switches and ensure that IP addresses are assigned.
- Step 4.** Verify connectivity.

The troubleshooting plan checks for the following:

- Step 1.** Verify that all VLANs have been created:
 - Was the VLAN created on all the switches?
 - Verify with the **show vlan** command.
- Step 2.** Ensure that ports are in the right VLANs and that trunking is working as expected:
 - Did all access ports have the **switchport access VLAN 500** command added?

- Should any other ports have been added? If so, make those changes.
- Were these ports previously used? If so, ensure that there are no extra commands enabled on these ports that can cause conflicts. If not, are the ports enabled?
- Are any user ports set to trunks? If so, issue the **switchport mode access** command.
- Are the trunk ports set to trunk mode?
- Is manual pruning of VLANs configured? VTP pruning prevents flooded traffic from propagating to switches that do not have members in specific VLANs. If manual pruning is enabled, ensure that the trunks necessary to carry VLAN 500 traffic have the VLAN in the allowed statements.

Step 3. Verify SVI configurations (if necessary):

- Is the SVI already created with the correct IP address and subnet mask?
- Is it enabled?
- Is routing enabled?

**Interactive
Graphic**

Activity 2.3.2.3: Troubleshoot Layer 3 Switching Issues

Refer to the online course to complete this activity.

Summary (2.4)

VLAN Trunking Protocol (VTP) reduces administration of VLANs in a switched network. A switch configured as the VTP server distributes and synchronizes VLAN information over trunk links to VTP-enabled switches throughout the domain.

The three VTP modes are server, client, and transparent.

The configuration revision number is used when determining whether a VTP switch should keep or whether to update its existing VLAN database. A switch overwrites its existing VLAN database if it receives a VTP update from another switch in the same domain with a higher configuration revision number. Therefore, when a switch is being added to a VTP domain, it must have the default VTP configuration or a lower configuration revision number than the VTP server.

Troubleshooting VTP can involve dealing with errors caused by incompatible VTP versions and incorrectly configured domain names or passwords.

Trunk negotiation is managed by Dynamic Trunking Protocol (DTP), which operates on a point-to-point basis between network devices. DTP is a Cisco proprietary protocol that is automatically enabled on Catalyst 2960 and Catalyst 3560 Series switches. A general best practice when a trunk link is required is to set the interface to **trunk** and **nonegotiate**. On links where trunking is not intended, DTP should be turned off.

When troubleshooting DTP, problems can be related to trunk mode mismatches, allowed VLANs on a trunk, and native VLAN mismatches.

Layer 3 switching using switch virtual interfaces (SVI) is a method of inter-VLAN routing that can be configured on Catalyst 2960 switches. An SVI with appropriate IP addressing is configured for each VLAN, and provides Layer 3 processing for packets to or from all switch ports associated with those VLANs.

Another method of Layer 3 inter-VLAN routing is using routed ports. A routed port is a physical port that acts similarly to an interface on a router. Routed ports are mostly configured between switches in the core and distribution layers.

Troubleshooting inter-VLAN routing with a router and with a Layer 3 switch are similar. Common errors involve VLAN, trunk, Layer 3 interface, and IP address configurations.

Practice

The following activities provide practice with the topics introduced in this chapter. The Labs and Class Activities are available in the companion *Scaling Networks v6 Labs & Study Guide* (ISBN 978-1-58713-433-3). The Packet Tracer activity instructions are also in the *Labs & Study Guide*. The PKA files are found in the online course.

**Labs**

Lab 2.1.4.5: Configure Extended VLANs, VTP, and DTP

Lab 2.2.2.5: Troubleshooting Inter-VLAN Routing

**Packet Tracer Activities**

Packet Tracer 2.1.4.4: Configure VTP and DTP

Packet Tracer 2.2.2.4: Troubleshooting Inter-VLAN Routing

Packet Tracer 2.2.3.3: Troubleshoot VTP and DTP Issues

Packet Tracer 2.3.1.5: Configure Layer 3 Switching and Inter-VLAN Routing

Check Your Understanding Questions

Complete all the review questions listed here to test your understanding of the topics and concepts in this chapter. The appendix “Answers to ‘Check Your Understanding’ Questions” lists the answers.

1. Which statement is true when VTP is configured on a switched network that incorporates VLANs?
 - A. VTP adds to the complexity of managing a switched network.
 - B. VTP allows a switch to be configured to belong to more than one VTP domain.
 - C. VTP dynamically communicates VLAN changes to all switches in the same VTP domain.
 - D. VTP is only compatible with the 802.1Q standard.
2. What are two features of VTP client mode operation? (Choose two.)
 - A. VTP clients can add VLANs of local significance.
 - B. VTP clients can forward VLAN information to other switches in the same VTP domain.
 - C. VTP clients can only pass VLAN management information without adopting changes.
 - D. VTP clients can forward broadcasts out all ports with no respect to VLAN information.
 - E. VTP clients are unable to add VLANs.

3. What does a client mode switch in a VTP management domain do when it receives a summary advertisement with a revision number higher than its current revision number?
 - A. It deletes the VLANs not included in the summary advertisement.
 - B. It increments the revision number and forwards it to other switches.
 - C. It issues an advertisement request for new VLAN information.
 - D. It issues summary advertisements to advise other switches of status changes.
 - E. It suspends forwarding until a subset advertisement update arrives.
4. What causes a VTP-configured switch to issue a summary advertisement?
 - A. A five-minute update timer has elapsed.
 - B. A new host has been attached to a switch in the management domain.
 - C. A port on the switch has been shut down.
 - D. The switch is changed to transparent mode.
5. Which three VTP parameters must be identical on all switches to participate in the same VTP domain? (Choose three.)
 - A. Domain name
 - B. Domain password
 - C. Mode
 - D. Pruning
 - E. Revision number
 - F. Version number
6. Which two statements describe VTP transparent mode operation? (Choose two.)
 - A. Transparent mode switches can add VLANs of local significance only.
 - B. Transparent mode switches can adopt VLAN management changes that are received from other switches.
 - C. Transparent mode switches can create VLAN management information.
 - D. Transparent mode switches originate updates about the status of their VLANs and inform other switches about that status.
 - E. Transparent mode switches pass any VLAN management information they receive to other switches.
7. Which two statements are true about the implementation of VTP? (Choose two.)
 - A. Switches must be connected via trunks.
 - B. Switches that use VTP must have the same switch name.

- C. The VTP domain name is case sensitive.
 - D. The VTP password is mandatory and case sensitive.
 - E. Transparent mode switches cannot be configured with new VLANs.
8. A network administrator is replacing a failed switch with a switch that was previously on the network. What precautionary step should the administrator take on the replacement switch to avoid incorrect VLAN information from propagating through the network?
- A. Change all the interfaces on the switch to access ports.
 - B. Change the VTP domain name.
 - C. Change the VTP mode to client.
 - D. Enable VTP pruning.
9. Which two events cause the VTP revision number on a VTP server to change? (Choose two.)
- A. Adding VLANs
 - B. Changing interface VLAN designations
 - C. Changing the switch to a VTP client
 - D. Changing the VTP domain name
 - E. Rebooting the switch
10. How are VTP messages sent between switches in a domain?
- A. Layer 2 broadcast
 - B. Layer 2 multicast
 - C. Layer 2 unicast
 - D. Layer 3 broadcast
 - E. Layer 3 multicast
 - F. Layer 3 unicast
11. A router has two FastEthernet interfaces and needs to connect to four VLANs in the local network. How can this be accomplished using the fewest number of physical interfaces without unnecessarily decreasing network performance?
- A. Add a second router to handle the inter-VLAN traffic.
 - B. Implement a router-on-a-stick configuration.
 - C. Interconnect the VLANs via the two additional FastEthernet interfaces.
 - D. Use a hub to connect the four VLANs with a FastEthernet interface on the router.

-
12. What distinguishes traditional legacy inter-VLAN routing from router-on-a-stick?
- A. Traditional routing is only able to use a single switch interface, while router-on-a-stick can use multiple switch interfaces.
 - B. Traditional routing requires a routing protocol, while router-on-a-stick needs to route only directly connected networks.
 - C. Traditional routing uses one port per logical network, while router-on-a-stick uses subinterfaces to connect multiple logical networks to a single router port.
 - D. Traditional routing uses multiple paths to the router and therefore requires STP, while router-on-a-stick does not provide multiple connections and therefore eliminates the need for STP.
13. What two statements are true regarding the use of subinterfaces for inter-VLAN routing? (Choose two.)
- A. Fewer router Ethernet ports are required than in traditional inter-VLAN routing.
 - B. The physical connection is less complex than in traditional inter-VLAN routing.
 - C. More switch ports are required than in traditional inter-VLAN routing.
 - D. Layer 3 troubleshooting is simpler than with traditional inter-VLAN routing.
 - E. Subinterfaces have no contention for bandwidth.
14. What is important to consider while configuring the subinterfaces of a router when implementing inter-VLAN routing?
- A. The IP address of each subinterface must be the default gateway address for each VLAN subnet.
 - B. The **no shutdown** command must be run on each subinterface.
 - C. The physical interface must have an IP address configured.
 - D. The subinterface numbers must match the VLAN ID number.
15. What steps must be completed in order to enable inter-VLAN routing using router-on-a-stick?
- A. Configure the physical interfaces on the router and enable a routing protocol.
 - B. Create the VLANs on the router and define the port membership assignments on the switch.
 - C. Create the VLANs on the switch to include port membership assignment and enable a routing protocol on the router.
 - D. Create the VLANs on the switch to include port membership assignment and configure subinterfaces on the router matching the VLANs.

Objectives

Upon completion of this chapter, you will be able to answer the following questions:

- What are common problems in a redundant switched network?
- How do different varieties of spanning-tree protocols operate?
- How do you implement PVST+ and Rapid PVST+ in a switched LAN environment?
- How are switch stacking and chassis aggregation implemented in a small switched LAN?

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary.

Layer 2 loop Page 107

broadcast storm Page 109

time to live (TTL) Page 109

bridge protocol data unit (BPDU) Page 115

Spanning Tree Algorithm (STA) Page 117

root bridge Page 117

bridge ID (BID) Page 117

extended system ID Page 117

root port Page 118

designated port Page 118

alternate port Page 118

backup port Page 118

blocking state Page 118

spanning-tree instance (STP instance) Page 119

bridge priority Page 119

root path cost Page 121

default port cost Page 121

topology change (TC) bit Page 129

topology change acknowledgment (TCA) bit Page 129

STP Page 141

802.1D Page 141

Common Spanning Tree (CST) Page 141

Per-VLAN Spanning Tree (PVST+) Page 141

Rapid Spanning Tree Protocol (RSTP) Page 141

IEEE 802.1w Page 141

Rapid Per-VLAN Spanning Tree (Rapid PVST+) Page 141

Multiple Spanning Tree Protocol (MSTP) Page 141

IEEE 802.1s Page 141

Multiple Spanning Tree (MST) Page 141

PortFast Page 142

BPDU guard Page 142

BPDU filter Page 142

root guard Page 142

loop guard Page 142

listening state Page 145

learning state Page 145

forwarding state Page 145

disabled state Page 145

edge port Page 150

point-to-point Page 152

STP diameter Page 171

Introduction (3.0.1.1)

Network redundancy is a key to maintaining network reliability. Multiple physical links between devices provide redundant paths. The network can then continue to operate when a single link or port has failed. Redundant links can also share the traffic load and increase capacity.

Multiple paths need to be managed so that *Layer 2 loops* are not created. The best paths are chosen, and an alternate path is immediately available in case a primary path fails. Spanning Tree Protocol (STP) is used to create one path through a Layer 2 network.

This chapter focuses on the protocols used to manage these forms of redundancy. It also covers some of the potential redundancy problems and their symptoms.



Class Activity 3.0.1.2: Stormy Traffic

Refer to *Scaling Networks v6 Labs & Study Guide* and the online course to complete this activity.

It is your first day on the job as a network administrator for a small to medium-sized business. The previous network administrator left suddenly after a network upgrade took place for the business.

During the upgrade, a new switch was added. Since the upgrade, many employees have complained that they are having trouble accessing the Internet and servers on the network. In fact, most of them cannot access the network at all. Your corporate manager asks you to immediately research what could be causing these connectivity problems and delays.

So you take a look at the equipment operating on the network at your main distribution facility in the building. You notice that the network topology seems to be visually correct and that cables have been connected correctly, routers and switches are powered on and operational, and switches are connected together to provide backup or redundancy.

However, you notice that all of your switches' status lights are constantly blinking at a very fast pace, to the point that they almost appear solid. You think you have found the problem with the connectivity issues your employees are experiencing.

Use the Internet to research STP. As you research, take notes and describe:

- Broadcast storm
- Switching loops
- The purpose of STP
- Variations of STP

Complete the reflection questions that accompany the PDF file for this activity. Save your work and be prepared to share your answers with the class.

Spanning Tree Concepts (3.1)

In this section, you will learn how to build a simple switched network with redundant links.

Purpose of Spanning Tree (3.1.1)

The focus of this topic is to describe how Spanning Tree Protocol can solve common looping problems in a redundant switched network.

Redundancy at OSI Layers 1 and 2 (3.1.1.1)

The three-tier hierarchical network design that uses core, distribution, and access layers with redundancy attempts to eliminate single points of failure on the network. Multiple cabled paths between switches provide physical redundancy in a switched network. This improves the reliability and availability of the network. Having alternate physical paths for data to traverse the network makes it possible for users to access network resources, despite path disruption.

The following steps explain how redundancy works in the topology shown in Figure 3-1:

1. PC1 is communicating with PC4 over a redundant network topology.
2. When the network link between S1 and S2 is disrupted, the path between PC1 and PC4 is automatically adjusted by STP to compensate for the disruption.

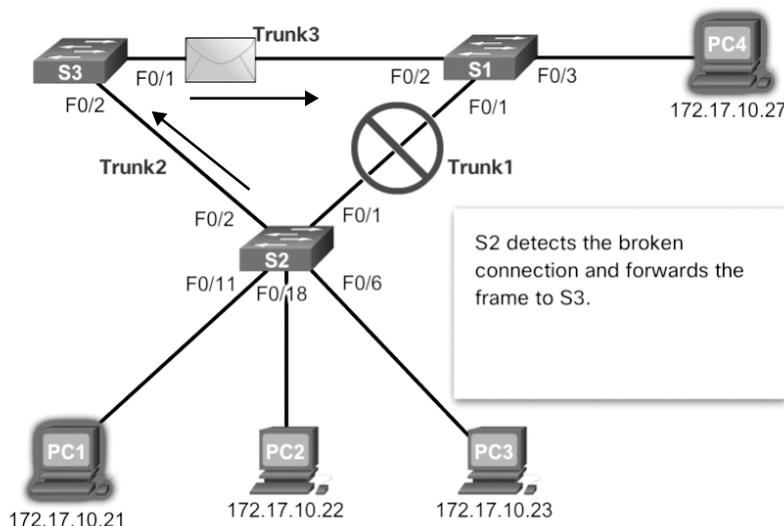


Figure 3-1 Redundancy in a Hierarchical Network

3. When the network connection between S1 and S2 is restored, the path is readjusted by STP to route traffic directly from S2 to S1 to get to PC4.

Note

To view an animation of these steps, refer to the online course.

For many organizations, the availability of the network is essential to supporting business needs. Therefore, the network infrastructure design is a critical business element. Path redundancy provides the necessary availability of multiple network services by eliminating the possibility of a single point of failure.

Note

OSI Layer 1 redundancy is illustrated using multiple links and devices, but more than just physical planning is required to complete the network setup. For the redundancy to work in a systematic way, the use of OSI Layer 2 protocols, such as STP, is also required.

Redundancy is an important part of the hierarchical design for preventing disruption of network services to users. Redundant networks require the addition of physical paths, but logical redundancy must also be part of the design. However, redundant paths in a switched Ethernet network may cause both physical and logical Layer 2 loops.

Logical Layer 2 loops may occur due to the natural operation of switches—specifically the learning and forwarding process. When multiple paths exist between two devices on a network, and there is no spanning-tree implementation on the switches, a Layer 2 loop occurs. A Layer 2 loop can result in three primary issues:

- **MAC database instability**—Instability in the content of the MAC address table results from copies of the same frame being received on different ports of the switch. Data forwarding can be impaired when the switch consumes the resources that are coping with instability in the MAC address table.
- **Broadcast storm**—Without some loop-avoidance process, each switch may flood broadcasts endlessly. This situation is commonly called a broadcast storm.
- **Multiple-frame transmission**—Multiple copies of unicast frames may be delivered to destination stations. Many protocols expect to receive only a single copy of each transmission. Multiple copies of the same frame can cause unrecoverable errors.

Issues with Layer 1 Redundancy: MAC Database Instability (3.1.1.2)

Ethernet frames do not have a *time to live (TTL)* attribute. As a result, if there is no mechanism enabled to block continued propagation of these frames on a switched network, they continue to propagate between switches endlessly, or until a link

is disrupted and breaks the loop. This continued propagation between switches can result in MAC database instability. This can occur due to broadcast frames forwarding.

Broadcast frames are forwarded out all switch ports except the original ingress port. This ensures that all devices in a broadcast domain are able to receive the frame. If there is more than one path through which the frame can be forwarded, an endless loop can result. When a loop occurs, it is possible for the MAC address table on a switch to constantly change with the updates from the broadcast frames, which results in MAC database instability.

The following sequence of events demonstrate the MAC database instability issue:

1. PC1 sends a broadcast frame to S2. S2 receives the broadcast frame on F0/11. When S2 receives the broadcast frame, it updates its MAC address table to record that PC1 is available on port F0/11.
2. Because it is a broadcast frame, S2 forwards the frame out all ports, including Trunk1 and Trunk2. When the broadcast frame arrives at S3 and S1, the switches update their MAC address tables to indicate that PC1 is available out port F0/1 on S1 and out port F0/2 on S3.
3. Because it is a broadcast frame, S3 and S1 forward the frame out all ports except the ingress port. S3 sends the broadcast frame from PC1 to S1. S1 sends the broadcast frame from PC1 to S3. Each switch updates its MAC address table with the incorrect port for PC1.
4. Each switch forwards the broadcast frame out all of its ports except the ingress port, which results in both switches forwarding the frame to S2.
5. When S2 receives the broadcast frames from S3 and S1, the MAC address table is updated with the last entry received from the other two switches.
6. S2 forwards the broadcast frame out all ports except the last received port. The cycle starts again.

Note

To view an animation of these sequence of events, refer to the online course.

Figure 3-2 shows a snapshot during sequence 6. Notice that S2 now thinks PC1 is reachable out the F0/1 interface.

This process repeats over and over again until the loop is broken by physically disconnecting the connections that are causing the loop or powering down one of the switches in the loop. This creates a high CPU load on all switches caught in the loop. Because the same frames are constantly being forwarded back and forth between all switches in the loop, the CPU of the switch must process a lot of data. This slows down performance on the switch when legitimate traffic arrives.

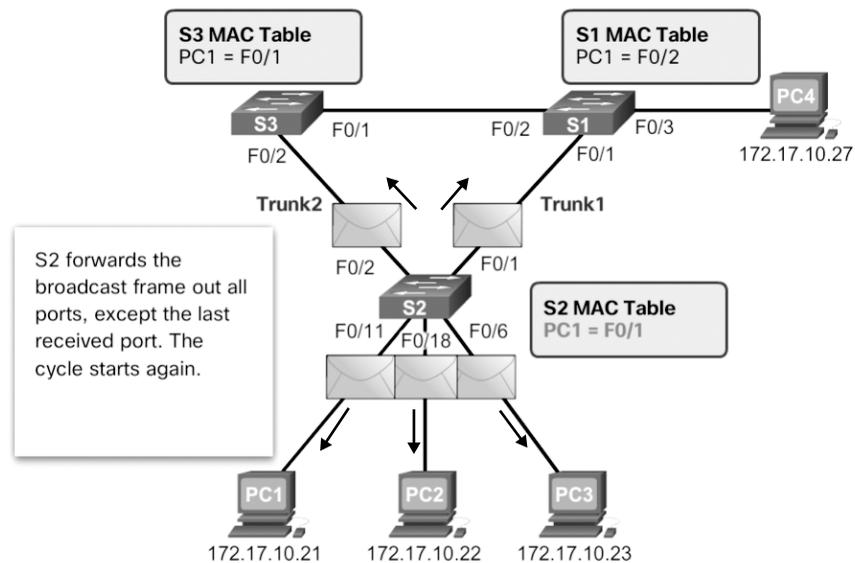


Figure 3-2 MAC Database Instability Example

A host caught in a network loop is not accessible to other hosts on the network. In addition, due to the constant changes in the MAC address table, the switch does not know which port to use to forward unicast frames. In this example just shown, the switches will have the incorrect ports listed for PC1. Any unicast frame destined for PC1 loops around the network, just as the broadcast frames do. More and more frames looping around the network eventually creates a broadcast storm.

Issues with Layer 1 Redundancy: Broadcast Storms (3.1.1.3)

A broadcast storm occurs when there are so many broadcast frames caught in a Layer 2 loop that all available bandwidth is consumed. Consequently, no bandwidth is available for legitimate traffic, and the network becomes unavailable for data communication. This is an effective denial of service (DoS).

Broadcast storms are inevitable on a looped network. As more devices send broadcasts over the network, more traffic is caught in the loop, consuming resources. This eventually creates a broadcast storm that causes the network to fail.

There are other consequences of broadcast storms. Because broadcast traffic is forwarded out every port on a switch, all connected devices have to process all the broadcast traffic that is being flooded endlessly around the looped network. This can cause the end device to malfunction because of the processing requirements needed to sustain such a high traffic load on the NIC.

The following sequence of events demonstrate the broadcast storm issue:

1. PC1 sends a broadcast frame out onto the looped network.
2. The broadcast frame loops between all the interconnected switches on the network.
3. PC4 also sends a broadcast frame out onto the looped network.
4. The PC4 broadcast frame gets caught in the loop between all the interconnected switches, just like the PC1 broadcast frame.
5. As more devices send broadcasts over the network, more traffic is caught in the loop, consuming resources. This eventually creates a broadcast storm that causes the network to fail.
6. When the network is fully saturated with broadcast traffic that is looping between the switches, the switch discards new traffic because it is unable to process it. Figure 3-3 displays the resulting broadcast storm.

Note

To view an animation of these sequence of events, refer to the online course.

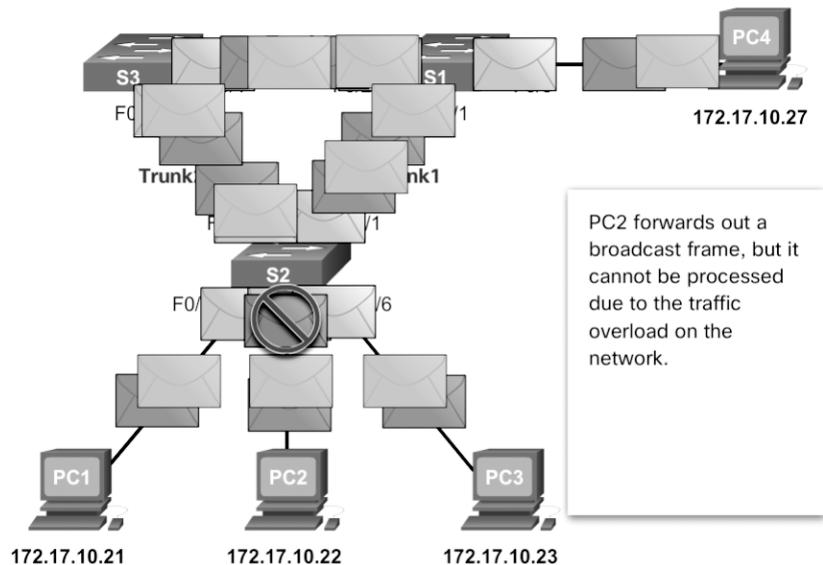


Figure 3-3 Broadcast Storm Example

A broadcast storm can develop in seconds because devices connected to a network regularly send out broadcast frames, such as ARP requests. As a result, when a loop is created, the switched network is quickly brought down.

Issues with Layer 1 Redundancy: Duplicate Unicast Frames (3.1.1.4)

Broadcast frames are not the only type of frames that are affected by loops.

Unknown unicast frames sent onto a looped network can result in duplicate frames arriving at the destination device. An unknown unicast frame occurs when the switch does not have the destination MAC address in its MAC address table and must forward the frame out all ports except the ingress port.

The following sequence of events demonstrate the duplicate unicast frames issue:

1. PC1 sends a unicast frame destined for PC4.
2. S2 does not have an entry for PC4 in its MAC table. In an attempt to find PC4, it floods the unknown unicast frame out all switch ports except the port that received the traffic.
3. The frame arrives at switches S1 and S3.
4. S1 has a MAC address entry for PC4, so it forwards the frame out to PC4.
5. S3 has an entry in its MAC address table for PC4, so it forwards the unicast frame out Trunk3 to S1.
6. S1 receives the duplicate frame and forwards the frame out to PC4.
7. PC4 has now received the same frame twice.

Figure 3-4 shows a snapshot during sequences 5 and 6.

Note

To view an animation of these sequence of events, refer to the online course.

Most upper-layer protocols are not designed to recognize duplicate transmissions. In general, protocols that make use of a sequence-numbering mechanism assume that the transmission has failed and that the sequence number has recycled for another communication session. Other protocols attempt to hand the duplicate transmission to the appropriate upper-layer protocol to be processed and possibly discarded.

Layer 2 LAN protocols, such as Ethernet, do not include a mechanism to recognize and eliminate endlessly looping frames. Some Layer 3 protocols implement a TTL mechanism that limits the number of times a Layer 3 networking device can retransmit a packet. Layer 2 devices do not have this mechanism, so they continue to retransmit looping traffic indefinitely. STP, a Layer 2 loop-avoidance mechanism, was developed to address these problems.

To prevent these issues from occurring in a redundant network, some type of spanning tree must be enabled on the switches. Spanning tree is enabled by default on Cisco switches to prevent Layer 2 loops from occurring.

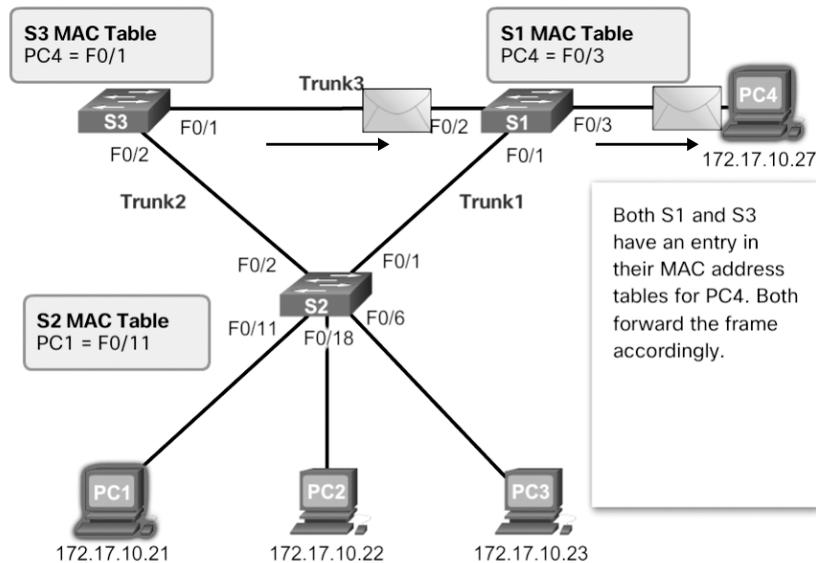


Figure 3-4 S1 and S3 Send a Duplicate Frame to PC4

Packet Tracer
Activity

Packet Tracer 3.1.1.5: Examining a Redundant Design

Background/Scenario

In this activity, you will observe how STP operates by default, and how it reacts when faults occur. Switches have been added to the network “out of the box.” Cisco switches can be connected to a network without any additional action required by the network administrator. For the purpose of this activity, the bridge priority (covered later in the chapter) was modified.

STP Operation (3.1.2)

The focus of this topic is to learn how to build a simple switched network using STP.

Spanning Tree Algorithm: Introduction (3.1.2.1)

Redundancy increases the availability of the network topology by protecting the network from a single point of failure, such as a failed network cable or switch. When physical redundancy is introduced into a design, loops and duplicate frames occur. Loops and duplicate frames have severe consequences for a switched network. Spanning Tree Protocol (STP) was developed to address these issues.

STP ensures that there is only one logical path between all destinations on the network by intentionally blocking redundant paths that could cause a loop. A port is

considered blocked when user data is prevented from entering or leaving that port. This does not include *bridge protocol data unit (BPDU)* frames that are used by STP to prevent loops. Blocking the redundant paths is critical to preventing loops on the network. The physical paths still exist to provide redundancy, but these paths are disabled to prevent the loops from occurring. If the path is ever needed to compensate for a network cable or switch failure, STP recalculates the paths and unblocks the necessary ports to allow the redundant path to become active.

Figure 3-5 illustrates normal STP operation when all switches have STP enabled:

1. PC1 sends a broadcast out onto the network.

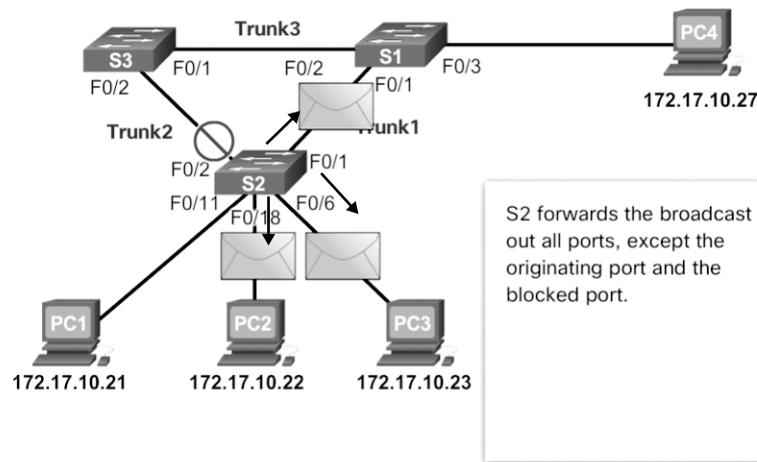


Figure 3-5 Normal STP Operation

2. S2 is configured with STP and has set the port for Trunk2 to a blocking state. The blocking state prevents ports from being used to forward user data, which prevents a loop from occurring. S2 forwards a broadcast frame out all switch ports except the originating port from PC1 and the port for Trunk2.
3. S1 receives the broadcast frame and forwards it out all of its switch ports, where it reaches PC4 and S3. S3 forwards the frame out the port for Trunk2, and S2 drops the frame. The Layer 2 loop is prevented.

Note

To view an animation of these steps, refer to the online course.

Figure 3-6 shows how STP recalculates the path when a failure occurs:

1. PC1 sends a broadcast out onto the network.

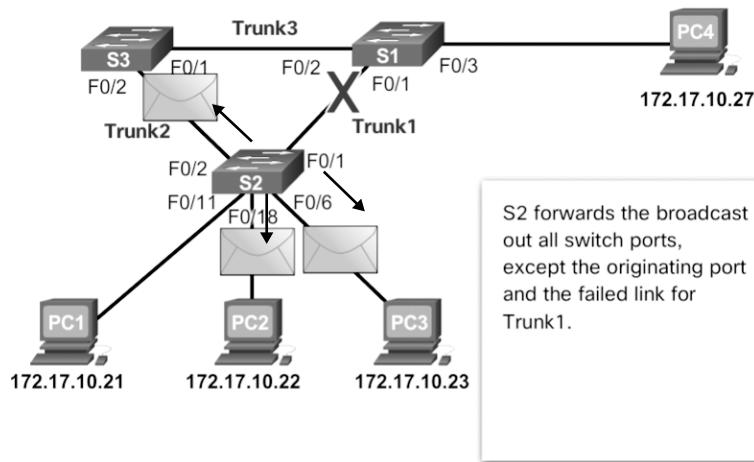


Figure 3-6 STP Compensates for Network Failure

2. The broadcast is then forwarded around the network.
3. As shown in the figure, the trunk link between S2 and S1 fails, resulting in the previous path being disrupted.
4. S2 unblocks the previously blocked port for Trunk2 and allows the broadcast traffic to traverse the alternate path around the network, permitting communication to continue. If this link comes back up, STP reconverges, and the port on S2 is again blocked.

Note

To view an animation of these sequence of events, refer to the online course.

STP prevents loops from occurring by configuring a loop-free path through the network using strategically placed “blocking-state” ports. The switches running STP are able to compensate for failures by dynamically unblocking the previously blocked ports and permitting traffic to traverse the alternate paths.

Up to now, we have used the terms *Spanning Tree Protocol* and *STP*. However, these terms can be misleading. Many professionals generically use these to refer to various implementations of spanning tree, such as Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP).

To communicate spanning tree concepts correctly, it is important to refer to the particular implementation or standard in context. The latest IEEE documentation on spanning tree (IEEE-802-1D-2004) says, “STP has now been superseded by the

Rapid Spanning Tree Protocol (RSTP).” The IEEE uses “STP” to refer to the original implementation of spanning tree and “RSTP” to describe the version of spanning tree specified in IEEE-802.1D-2004. In this curriculum, when the original Spanning Tree Protocol is the context of a discussion, the phrase “original 802.1D spanning tree” is used to avoid confusion. Because the two protocols share much of the same terminology and methods for the loop-free path, the primary focus is on the current standard and the Cisco proprietary implementations of STP and RSTP.

Note

STP is based on an algorithm that Radia Perlman invented while working for Digital Equipment Corporation and published in the 1985 paper “An Algorithm for Distributed Computation of a Spanning Tree in an Extended LAN.”

Spanning Tree Algorithm: Port Roles (3.1.2.2)

IEEE 802.1D STP and RSTP use *Spanning Tree Algorithm (STA)* to determine which switch ports on a network must be put in blocking state to prevent loops from occurring. STA designates a single switch as the *root bridge* and uses it as the reference point for all path calculations. In Figure 3-7, the root bridge (switch S1) is chosen through an election process. All switches that are participating in STP exchange BPDU frames to determine which switch has the lowest *bridge ID (BID)* on the network. The switch with the lowest BID automatically becomes the root bridge for the STA calculations.

Note

For simplicity, assume until otherwise indicated that all ports on all switches are assigned to VLAN 1. Each switch has a unique MAC address associated with VLAN 1.

A BPDU is a messaging frame exchanged by switches for STP. Each BPDU contains a BID that identifies the switch that sent the BPDU. The BID contains a priority value, the MAC address of the sending switch, and an optional *extended system ID*. The lowest BID value is determined by the combination of these three fields.

After the root bridge has been determined, the STA calculates the shortest path to the root bridge. Each switch uses the STA to determine which ports to block. While the STA determines the best paths to the root bridge for all switch ports in the broadcast domain, traffic is prevented from being forwarded through the network. The STA considers both path and port costs when determining which ports to block. The path costs are calculated using port cost values associated with port speeds for each switch port along a given path. The sum of the port cost values determines the overall path cost to the root bridge. If there is more than one path to choose from, STA chooses the path with the lowest path cost.

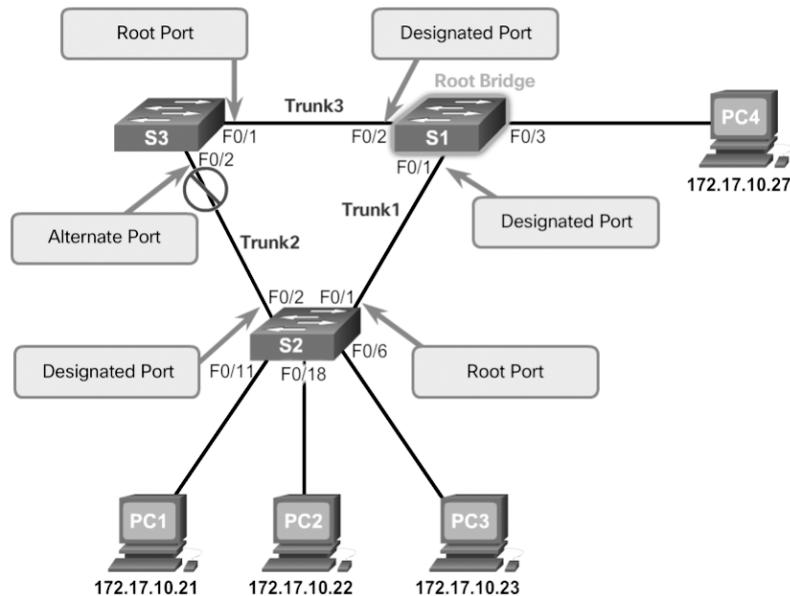


Figure 3-7 STP Algorithm—RSTP Port Roles

When the STA has determined which paths are most desirable relative to each switch, it assigns port roles to the participating switch ports. The port roles describe their relationship in the network to the root bridge and whether they are allowed to forward traffic:

- **Root port**—A root port is selected on all non-root bridge switches on a per-switch basis. Root ports are the switch ports closest to the root bridge, based on the overall cost to the root bridge. There can be only one root port per non-root switch. Root ports could be single-link interfaces or an EtherChannel port channel interface.
- **Designated port**—A designated port is a non-root port that is permitted to forward traffic. Designated ports are selected on a per-segment basis, based on the cost of each port on either side of the segment and the total cost calculated by STP for that port to get back to the root bridge. If one end of a segment is a root port, then the other end is a designated port. All ports on the root bridge are designated ports.
- **Alternate port** and **backup port**—An alternate port and a backup port are in a *blocking state* (or discarding state) to prevent loops. Alternate ports are selected only on links where neither end is a root port. Only one end of the segment is blocked, while the other end remains in forwarding state, allowing for a faster transition to the forwarding state when necessary.
- **Disabled ports**—A disabled port is a switch port that is shut down.

Note

The port roles displayed are those defined by RSTP. The role originally defined by the 802.1D STP for alternate and backup ports was non-designated.

For example, on the link between S2 and the root bridge S1 in Figure 3-7, the root port selected by STP is the F0/1 port on S2. The root port selected by STP on the link between S3 and S1 is the F0/1 port on S3. Because S1 is the root bridge, all of its ports (that is, F0/1 and F0/2) become designated ports.

Next, the interconnecting link between S2 and S3 must negotiate to see which port will become the designated port and which port will transition to alternate. In this scenario, the F0/2 port on S2 transitioned to a designated port, and the F0/2 port on S3 transitioned to an alternate port and is therefore blocking traffic.

Spanning Tree Algorithm: Root Bridge (3.1.2.3)

As shown in Figure 3-8, every *spanning-tree instance (STP instance)* has a switch designated as the root bridge. The root bridge serves as a reference point for all spanning-tree calculations to determine which redundant paths to block.

An election process determines which switch becomes the root bridge.

Figure 3-9 shows the BID fields. The BID is made up of a priority value, an extended system ID, and the MAC address of the switch. The *bridge priority* value is automatically assigned but can be modified. The extended system ID is used to specify a VLAN ID or a Multiple Spanning Tree Protocol (MSTP) instance ID. The MAC address field initially contains the MAC address of the sending switch.

All switches in the broadcast domain participate in the election process. After a switch boots, it begins to send out BPDU frames every two seconds. These BPDUs contain the switch BID and the root ID.

The switch with the lowest BID becomes the root bridge. At first, all switches declare themselves as the root bridge. But through the exchange of several BPDUs, the switches eventually agree on the root bridge.

Specifically, each switch forwards BPDU frames containing their BID and the root ID to adjacent switches in the broadcast domain. The receiving switch compares its current root ID with the received root ID identified in the received frames. If the received root ID is lower, the receiving switch updates its root ID with the lower root ID. It then forwards new BPDU frames containing the lower root ID to the other adjacent switches. Eventually, the switch with the lowest BID is identified as the root bridge for the spanning-tree instance.

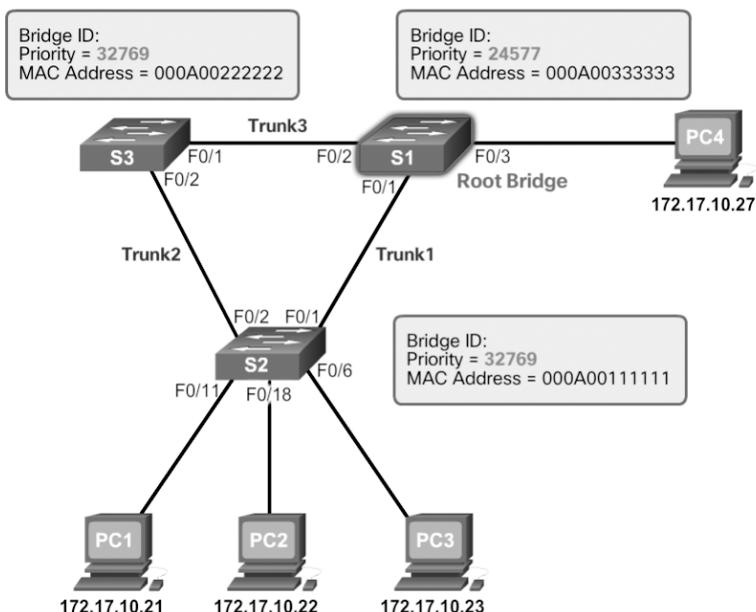


Figure 3-8 The Root Bridge

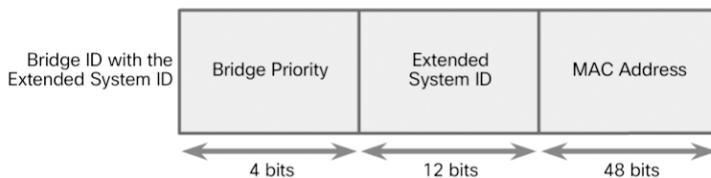


Figure 3-9 BID Fields

A root bridge is elected for each spanning-tree instance. It is possible to have multiple distinct root bridges for different sets of VLANs. If all ports on all switches are members of VLAN 1, then there is only one spanning-tree instance. The extended system ID includes the VLAN ID and plays a role in how spanning-tree instances are determined.

The BID consists of a configurable bridge priority number and a MAC address. Bridge priority is a value between 0 and 65,535. The default is 32,768. If two or more switches have the same priority, the switch with the lowest MAC address becomes the root bridge.

Note

The reason the bridge priority value in Figure 3-8 displays 32,769 instead of the default value 32,768 is that the STA also adds the default VLAN number (VLAN 1) to the priority value.

Spanning Tree Algorithm: Root Path Cost (3.1.2.4)

When the root bridge has been elected for the spanning-tree instance, STA starts determining the best paths to the root bridge.

Switches send BPDUs, which include the *root path cost*. This is the cost of the path from the sending switch to the root bridge. It is calculated by adding the individual port costs along the path from the switch to the root bridge. When a switch receives the BPDU, it adds the ingress port cost of the segment to determine its internal root path cost. It then advertises the new root path cost to its adjacent peers.

The *default port cost* is defined by the speed at which the port operates. As shown in Table 3-1, 10 Gbps Ethernet ports have a port cost of 2, 1 Gbps Ethernet ports have a port cost of 4, 100 Mbps Fast Ethernet ports have a port cost of 19, and 10 Mbps Ethernet ports have a port cost of 100.

Table 3-1 Revised IEEE Cost Values

Link Speed	Cost (Revised IEEE 802.1D-1998 Specification)
10 Gbps	2
1 Gbps	4
100 Mbps	19
10 Mbps	100

Note

The original IEEE specification did not account for links faster than 1 Gbps. Specifically, 1 Gbps links were assigned a port cost of 1, 100 Mbps link a cost of 10, and 10 Mbps links a cost of 100. Any link faster than 1 Gbps (i.e., 10 GE) was automatically assigned the same port cost of 1 Gbps links (i.e., port cost of 1).

Note

Modular switches such as the Catalyst 4500 and 6500 switches support higher port costs—specifically, 10 Gbps = 2000, 100 Gbps = 200, and 1 Tbps = 20 port costs.

As Ethernet technologies evolve, the port cost values may change to accommodate the different speeds available. The nonlinear numbers in the table accommodate some improvements to the older Ethernet standard.

Although switch ports have a default port cost associated with them, the port cost is configurable. The ability to configure individual port costs gives the administrator the flexibility to manually control the spanning-tree paths to the root bridge.

To configure the port cost of an interface, enter the **spanning-tree cost *value*** command in interface configuration mode. The value can be between 1 and 200,000,000.

Example 3-1 displays how to change the port cost of F0/1 to 25 by using the **spanning-tree cost 25** interface configuration mode command.

Example 3-1 Changing the Default Port Cost

```
S2(config)# interface f0/1
S2(config-if)# spanning-tree cost 25
S2(config-if)# end
```

Example 3-2 shows how to restore the port cost to the default value, 19, by entering the **no spanning-tree cost** interface configuration mode command.

Example 3-2 Restoring the Default Port Cost

```
S2(config)# interface f0/1
S2(config-if)# no spanning-tree cost
S2(config-if)# end
S2#
```

The internal root path cost is equal to the sum of all the port costs along the path to the root bridge. Paths with the lowest cost become preferred, and all other redundant paths are blocked.

In Figure 3-10, the internal root path cost from S2 to the root bridge S1 using Path 1 is 19 (based on Table 3-1), while the internal root path cost using Path 2 is 38.

Path 1 has a lower overall path cost to the root bridge and therefore becomes the preferred path. STP configures the redundant path to be blocked, which prevents a loop from occurring.

Use the **show spanning-tree** command as shown in Example 3-3 to verify the root ID and internal root path cost to the root bridge.

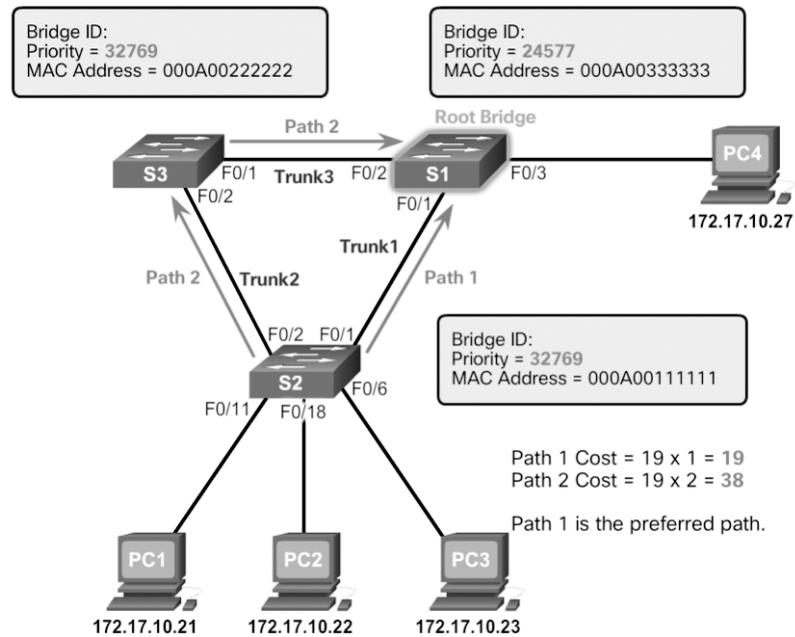


Figure 3-10 Root Path Cost Example

Example 3-3 Verifying the Root Bridge and Port Costs

```

S2# show spanning-tree
VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority    24577
            Address     000A.0033.0033
            Cost       19
            Port       1
            Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32769 (priority 32768 sys-id-ext 1)
            Address     000A.0011.1111
            Hello Time  2 sec  Max Age 20 sec  Forward Delay 15 sec
            Aging Time  15 sec

Interface    Role  Sts  Cost    Prio.Nbr  Type
-----
Fa0/1        Root  FWD  19      128.1     Edge P2p
Fa0/2        Desg  FWD  19      128.2     Edge P2p

```

The output generated identifies the root BID as 24577.000A0033003, with a root path cost of 19. The Cost field value changes depending on how many switch ports must be traversed to get to the root bridge. Also notice that each interface is assigned a port role and port cost of 19.

Port Role Decisions for RSTP (3.1.2.5)

After the root bridge is elected, the STA determines port roles on interconnecting links. The next seven figures help illustrate this process.

In Figure 3-11, switch S1 is the root bridge.

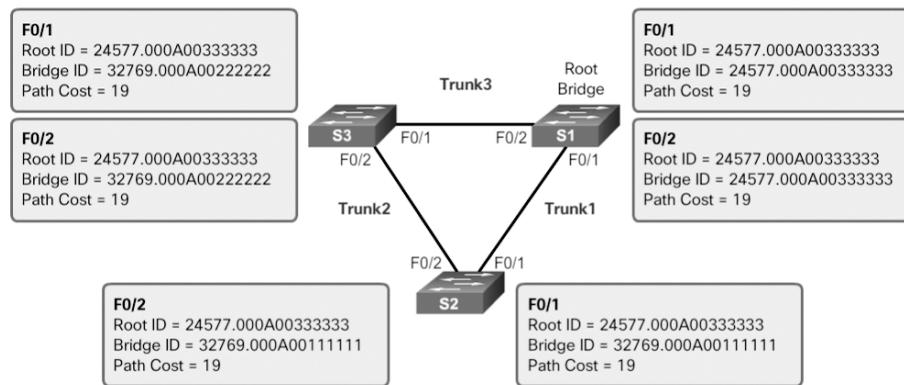


Figure 3-11 Port Role Decisions: Step 1

The root bridge always transitions its interconnecting links to designated port status. For example, in Figure 3-12, S1 configures both of its trunk ports connected to F0/1 and F0/2 as designated ports.

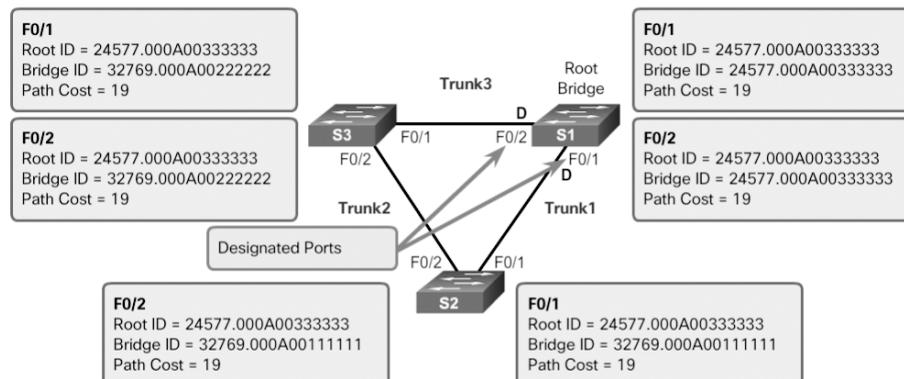


Figure 3-12 Port Role Decisions: Step 2

Non-root switches transition ports with the lowest root path cost to root ports. In Figure 3-13, S2 and S3 transition their F0/1 ports to root ports.

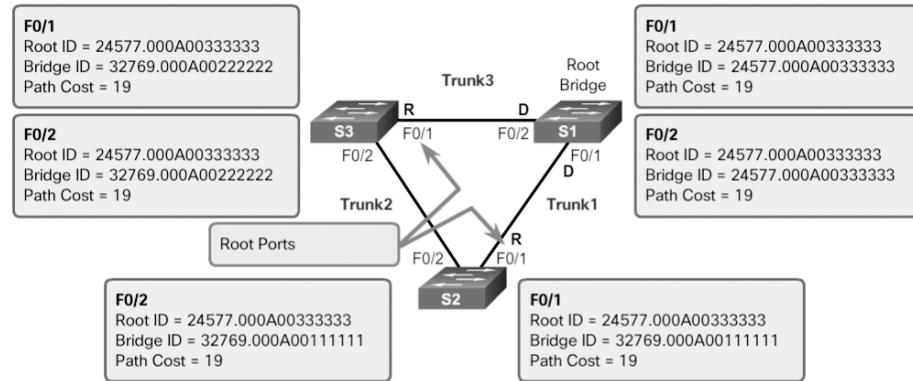


Figure 3-13 Port Role Decisions: Step 3

After the root ports are selected, the STA decides which ports will have the designated and alternate roles, as illustrated with the S2 to S3 link in Figure 3-14.

The root bridge already transitioned its ports to designated status. Non-root switches must transition their non-root ports to either designated or alternate port status.

The two non-root switches exchange BPDUs, as illustrated in Figure 3-15.

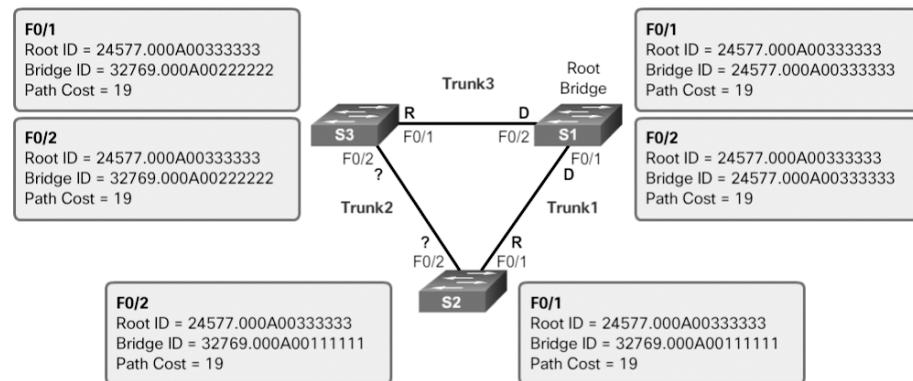


Figure 3-14 Port Role Decisions: Step 4

The incoming BPDUs include the BID of the sending switch. When a switch receives a BPDUs frame, it compares the BID in the BPDUs with its BID to see which one is higher. The switch advertising the higher BID transitions its port to alternate status.

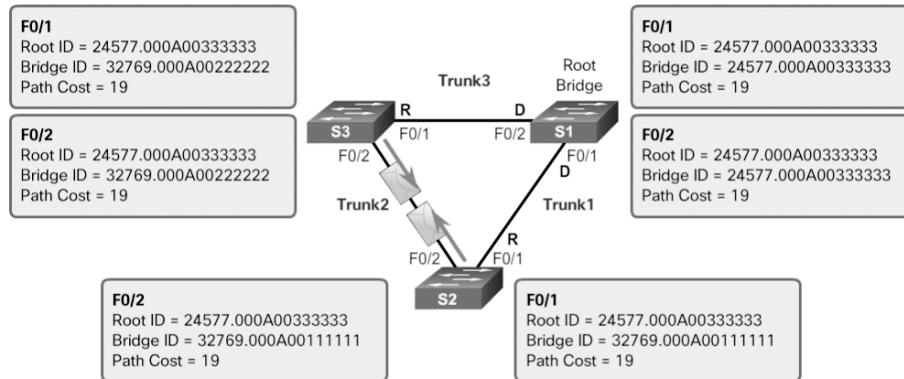


Figure 3-15 Port Role Decisions: Step 5

As illustrated in Figure 3-16, S3 has a higher BID (32769.000A00222222) compared to the BID of S2 (32769.000A00111111). Therefore, S3 transitions its F0/2 port to alternate status.

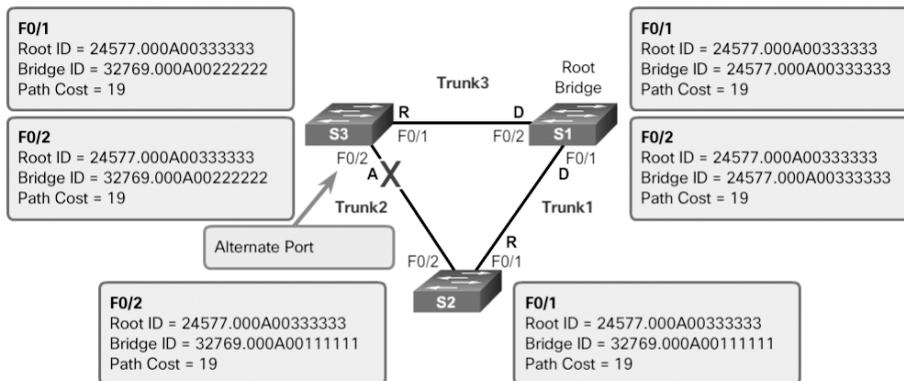


Figure 3-16 Port Role Decisions: Step 6

S2 has the lower BID and therefore transitions its port to designated status, as shown in Figure 3-17.

Keep in mind that the first priority is the lowest-path cost to the root bridge and that the sender's BID is used only if the port costs are equal.

Each switch determines which port roles are assigned to each of its ports to create the loop-free spanning tree.

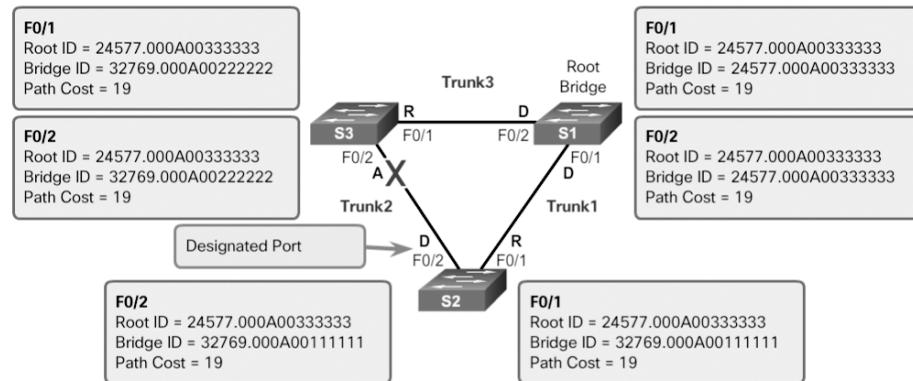


Figure 3-17 Port Role Decisions: Step 7

Designated and Alternate Ports (3.1.2.6)

When determining the root port on a switch, the switch compares the path costs on all switch ports participating in the spanning tree. The switch port with the lowest overall path cost to the root bridge is automatically assigned the root port role because it is closest to the root bridge. In a network topology of switches, all non-root bridge switches have a single root port chosen, and that port provides the lowest-cost path back to the root bridge.

A root bridge does not have any root ports. All ports on a root bridge are designated ports. A switch that is not the root bridge of a network topology has only one root port defined.

Figure 3-18 shows a topology with four switches.

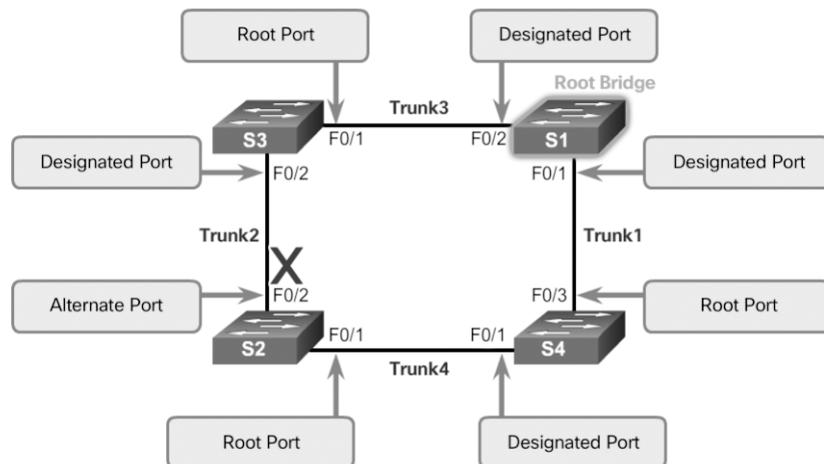


Figure 3-18 Determining Designated and Alternate Ports

Examine the port roles, and you see that port F0/1 on switch S3 and port F0/3 on switch S4 have been selected as root ports because they have the lowest-cost path (root path cost) to the root bridge for their respective switches.

S2 has two ports, F0/1 and F0/2, with equal-cost paths to the root bridge. In this case, the bridge IDs of the neighboring switches, S3 and S4, will be used to break the tie. This is known as the sender's BID. S3 has a BID of 24577.5555.5555.5555, and S4 has a BID of 24577.1111.1111.1111. Because S4 has a lower BID, S2's F0/1 port, the port connected to S4, becomes the root port.

Note

The BIDs are not shown in Figure 3-18.

Next, designated ports need to be selected on shared segments. S2 and S3 connect to the same LAN segment, and therefore, they exchange BPDU frames. STP determines whether S2's F0/2 port or S3's F0/2 port is the designated port for the shared segment. The switch with the lower-cost path to the root bridge (root path cost) has its port selected as the designated port. S3's F0/2 port has a lower-cost path to the root bridge, so it is the designated port for that segment.

S2 and S4 go through a similar process for their shared segment. S4's F0/1 port has the lower-cost path to the root bridge and becomes the designated port on this shared segment.

All STP port roles have been assigned except for S2's F0/2 port. S2's F0/1 port has already been selected as the root port for that switch. Because S3's F0/2 port is the designated port for this segment, S2's F0/2 port becomes an alternate port.

The designated port is the port that sends and receives traffic to and from that segment to the root bridge. This is the best port on that segment toward the root bridge. The alternate port does not send or receive traffic on that segment; this is the loop prevention part of STP.

802.1D BPDU Frame Format (3.1.2.7)

The STA depends on the exchange of BPDUs to determine a root bridge. As shown in Table 3-2, a BPDU frame contains 12 distinct fields that convey the path and priority information used to determine the root bridge and the paths to the root bridge:

Table 3-2 The BPDU Fields

Field Number	Bytes	Field	Description
1	2	Protocol ID	This field indicates the type of protocol being used. This field contains the value 0.
2	1	Version	This field indicates the version of the protocol. This field contains the value 0.
3	1	Message type	This field indicates the type of message. This field contains the value 0.
4	1	Flags	This field includes one of the following: <ul style="list-style-type: none"> ■ <i>Topology change (TC) bit</i>, which signals a topology change in the event that a path to the root bridge has been disrupted ■ <i>Topology change acknowledgment (TCA) bit</i>, which is set to acknowledge receipt of a configuration message with the TC bit set
5	8	Root ID	This field indicates the root bridge by listing its 2-byte priority followed by its 6-byte MAC address ID. When a switch first boots, the root ID is the same as the bridge ID. However, as the election process occurs, the lowest bridge ID replaces the local root ID to identify the root bridge switch.
6	4	Root Path Cost	This field indicates the cost of the path from the bridge sending the configuration message to the root bridge. The path cost field is updated by each switch along the path to the root bridge.
7	8	Bridge ID	This field indicates the priority, extended system ID, and MAC address ID of the bridge sending the message. This label allows the root bridge to identify where the BPDU originated and to identify the multiple paths from the switch to the root bridge. When the root bridge receives more than one BPDU from a switch with different path costs, it knows that there are two distinct paths and uses the path with the lower cost.
8	2	Port ID	This field indicates the port number from which the configuration message was sent. This field allows loops created by multiple attached bridges to be detected and corrected.

Field Number	Bytes	Field	Description
9	2	Message age	This field indicates the amount of time that has elapsed since the root sent the configuration message on which the current configuration message is based.
10	2	Max age	This field indicates when the current configuration message should be deleted. When the message age reaches the maximum age, the switch expires the current configuration and initiates a new election to determine a new root bridge because it assumes that it has been disconnected from the root bridge. This is 20 seconds by default but can be tuned to be between 6 and 40 seconds.
11	2	Hello time	This field indicates the time between root bridge configuration messages. The interval defines how long the root bridge waits between sending configuration message BPDUs. This is equal to 2 seconds by default but can be tuned to be between 1 and 10 seconds.
12	2	Forward delay	This field indicates the length of time bridges should wait before transitioning to a new state after a topology change. If a bridge transitions too soon, it is possible that not all network links will be ready to change their state, and loops can result. This is, by default, equal to 15 seconds for each state but can be tuned to be between 4 and 30 seconds.

The first four fields in the BPDU identify specifics about the type of BPDU message, including the protocol, version, message type, and status flags. The next four fields are used to identify the root bridge and the root path cost to the root bridge. The last four fields are all timer-related fields that determine how frequently BPDU messages are sent and how long the information received through the BPDU process is retained.

Figure 3-19 shows a BPDU frame that was captured using Wireshark. In this example, the BPDU frame contains more fields than previously described. The BPDU message is encapsulated in an Ethernet frame when it is transmitted across the network. The 802.3 header indicates the source and destination addresses of the BPDU frame. This frame has a destination MAC address of 01:80:C2:00:00:00, which is a multicast address for the spanning-tree group. When a frame is addressed with this MAC address, each switch that is configured for spanning tree accepts and reads the information from the frame. All other devices on the network disregard the frame.

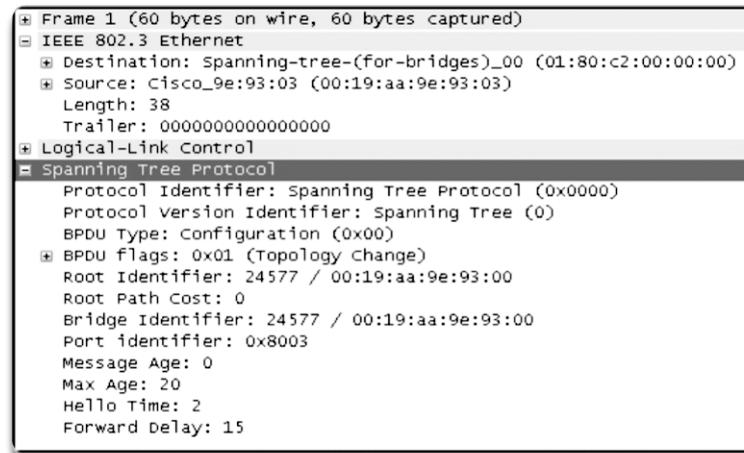


Figure 3-19 The BPDU Example

In Figure 3-19, the root ID and the BID are the same in the captured BPDU frame. This indicates that the frame was captured from a root bridge. The timers are all set to the default values.

802.1D BPDU Propagation and Process (3.1.2.8)

Each switch in a broadcast domain initially assumes that it is the root bridge for a spanning-tree instance, so the BPDU frames that are sent contain the BID of the local switch as the root ID. By default, BPDU frames are sent every two seconds after a switch is booted. The default value of the hello timer specified in the BPDU frame is two seconds. Each switch maintains local information about its own BID, the root ID, and the root path cost.

When adjacent switches receive a BPDU frame, they compare the root ID from the BPDU frame with the local root ID. If the root ID in the received BPDU is lower than the local root ID, the switch updates the local root ID and the ID in its BPDU messages. These messages indicate the new root bridge on the network. If the local root ID is lower than the root ID received in the BPDU frame, the BPDU frame is discarded.

The distance to the root bridge is indicated by the root path cost in the BPDU. The ingress port cost is then added to the root path cost in the BPDU to determine the internal root path cost from this switch to the root bridge. For example, if the BPDU was received on a Fast Ethernet switch port, the root path cost in the BPDU would be added to the ingress port cost of 19, for a cumulative internal root path cost. This is the cost from this switch to the root bridge.

After a root ID has been updated to identify a new root bridge, all subsequent BPDU frames sent from that switch contain the new root ID and updated root path cost. That way, all other adjacent switches are able to see the lowest root ID identified at all times. As the BPDU frames pass between other adjacent switches, the path cost is continually updated to indicate the total path cost to the root bridge. Each switch in the spanning tree uses its path costs to identify the best possible path to the root bridge.

The following figures summarize the BPDU process.

Note

Bridge priority is the initial deciding factor when electing a root bridge. If the bridge priorities of all the switches are the same, the device with the lowest MAC address becomes the root bridge.

In Figure 3-20, S2 forwards BPDU frames identifying itself as the root bridge out all switch ports.

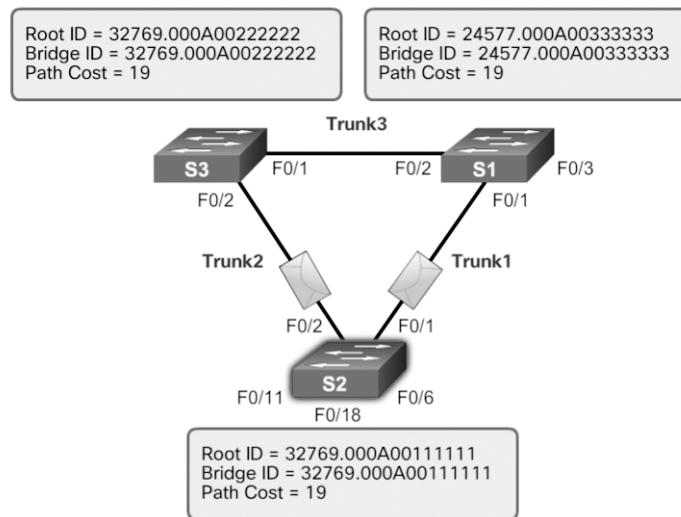


Figure 3-20 The BPDU Process: Step 1

In Figure 3-21, S3 receives the BPDU from S2 and compares its root ID with the BPDU frame it received. The priorities are equal, so S3 examines the MAC address portion. S2 has a lower MAC address value, so S3 updates its root ID with the S2 root ID. S3 now considers S2 the root bridge.

In Figure 3-22, S1 receives the BPDU from S2 and compares its root ID with the BPDU frame it received. S1 identifies its root ID as the lower value and discards the BPDU from S2.

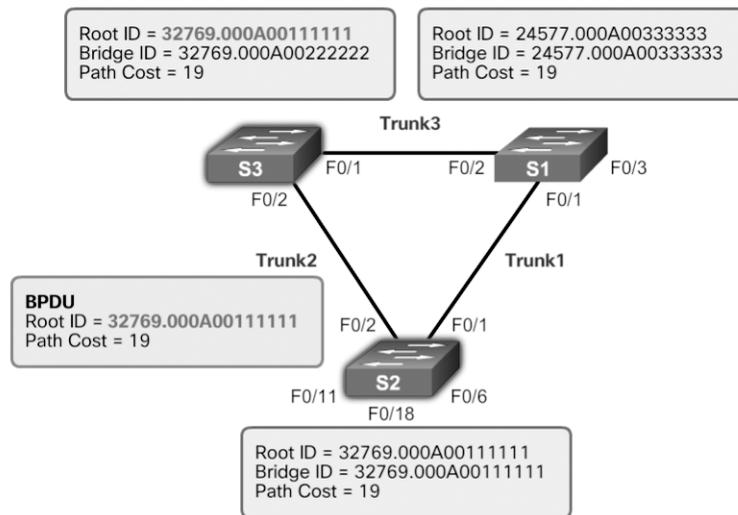


Figure 3-21 The BPDUs Process: Step 2

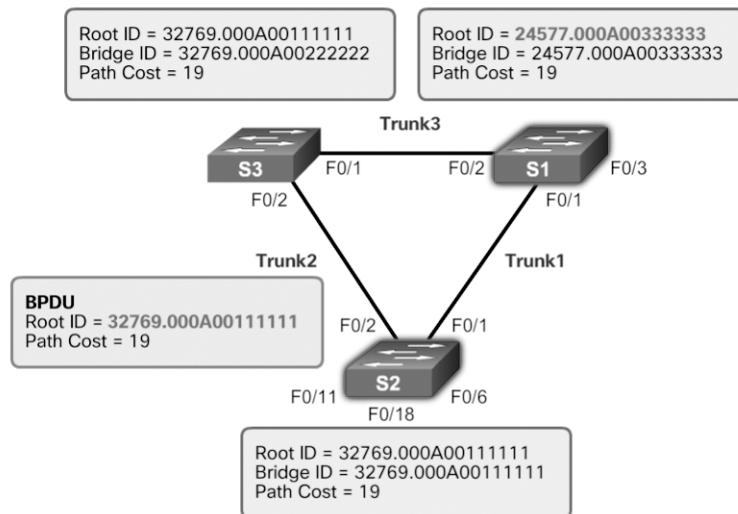


Figure 3-22 The BPDUs Process: Step 3

In Figure 3-23, S3 sends out BPDUs advertising its BID and the new root ID, which is that of S2.

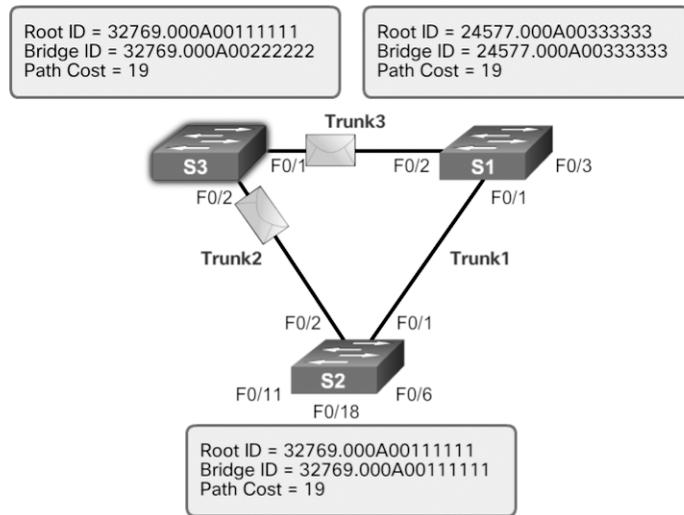


Figure 3-23 The BPDU Process: Step 4

In Figure 3-24, S2 receives the BPDUs from S3 and discards it after verifying that the root ID in the BPDUs matches its local root ID.

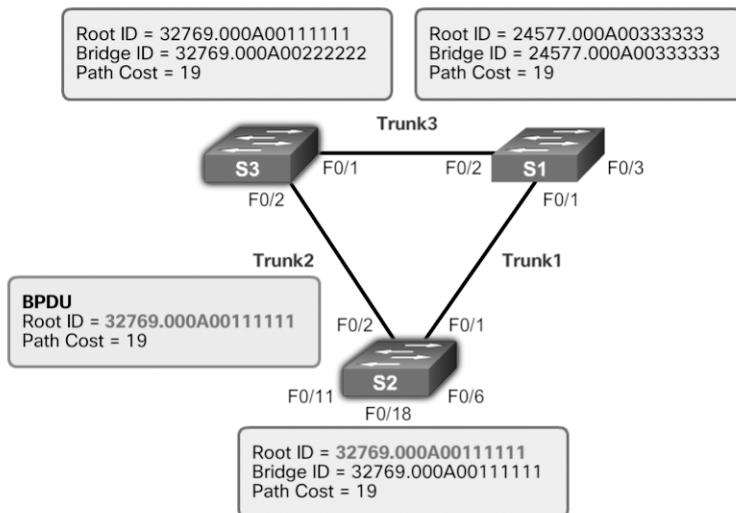


Figure 3-24 The BPDU Process: Step 5

In Figure 3-25, S1 receives the BPDUs from S3 and discards it because S1 has a lower priority value in its root ID.

In Figure 3-26, S1 sends out BPDUs frames advertising its BID and itself as the root ID.

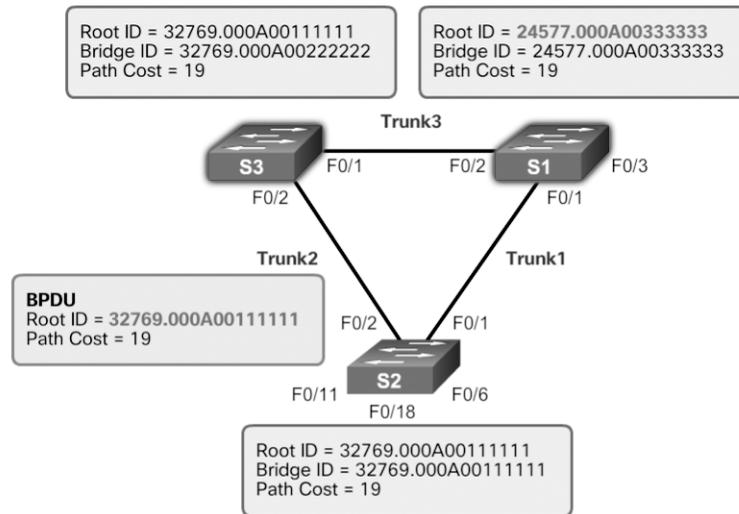


Figure 3-25 The BPDUs Process: Step 6

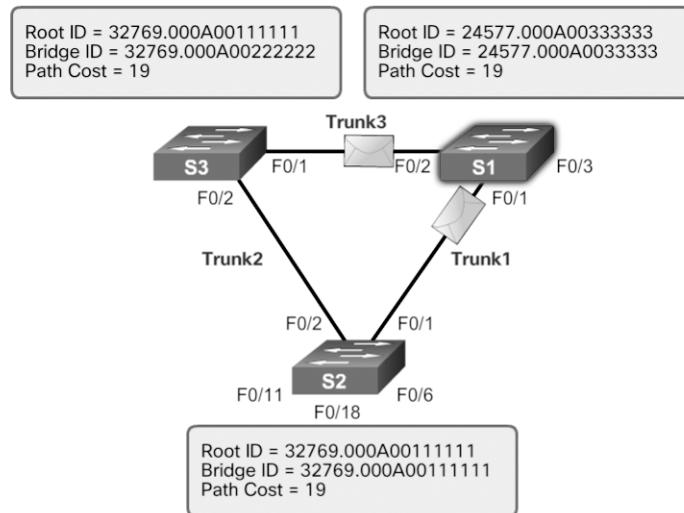


Figure 3-26 The BPDUs Process: Step 7

In Figure 3-27, S3 receives the BPDUs from S1 and compares its root ID with the BPDUs frame it received. S3 identifies the received root ID to be the lower value. Therefore, S3 updates its root ID values to indicate that S1 is now the root bridge.

In Figure 3-28, S2 receives the BPDUs from S1 and compares its root ID with the BPDUs frame it received. S2 identifies the received root ID to be the lower value. Therefore, S2 updates its root ID values to indicate that S1 is now the root bridge.

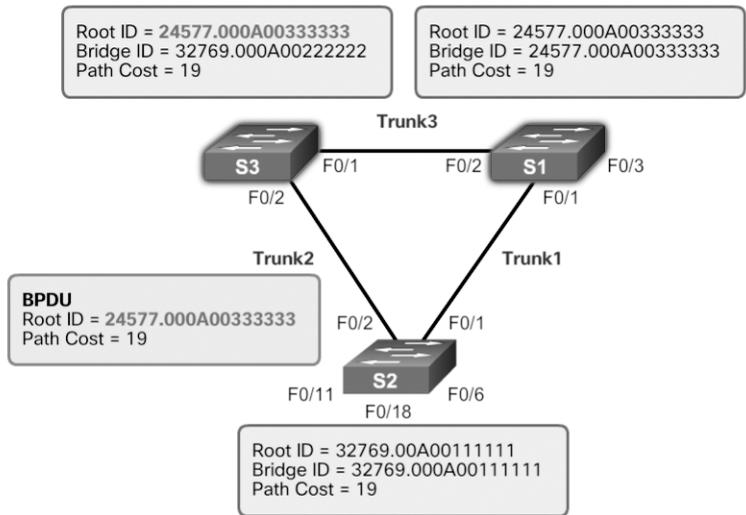


Figure 3-27 The BPDU Process: Step 8

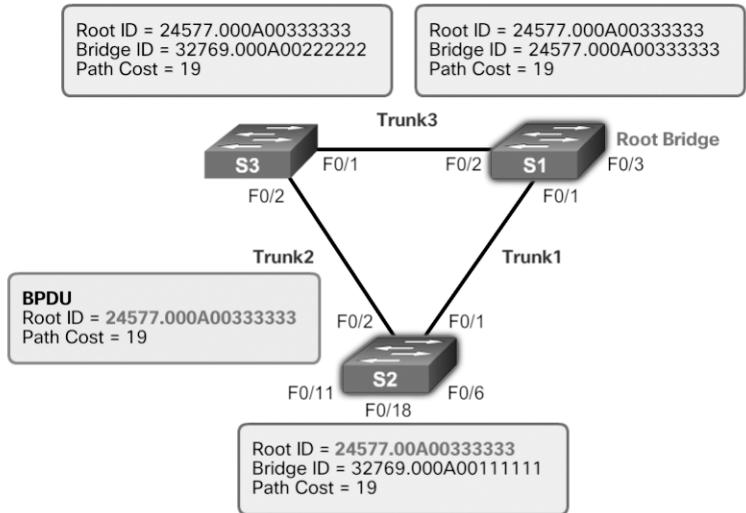


Figure 3-28 The BPDU Process: Step 9

Extended System ID (3.1.2.9)

The bridge ID (BID) is used to determine the root bridge on a network. The BID field of a BPDUs frame contains three separate fields:

- Bridge priority
- Extended system ID
- MAC address

Each of these fields is used during the root bridge election.

Bridge Priority

The bridge priority is a customizable value that can be used to influence which switch becomes the root bridge. The switch with the lowest priority, which implies the lowest BID, becomes the root bridge because a lower priority value takes precedence. For example, to ensure that a specific switch is always the root bridge, set the priority to a lower value than the rest of the switches on the network.

The default priority value for all Cisco switches is the decimal value 32768. The range is 0 to 61440, in increments of 4096. Therefore, valid priority values are 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. A bridge priority of 0 takes precedence over all other bridge priorities. All other values are rejected.

Extended System ID

Early implementations of IEEE 802.1D were designed for networks that did not use VLANs. There was a single common spanning tree across all switches. For this reason, in older Cisco switches, the extended system ID could be omitted in BPDU frames.

As VLANs became common for network infrastructure segmentation, 802.1D was enhanced to include support for VLANs, which required that the VLAN ID be included in the BPDU frame. VLAN information is included in the BPDU frame through the use of the extended system ID. All newer switches include the use of the extended system ID by default.

As shown in Figure 3-29, the bridge priority field is 2 bytes, or 16 bits, in length. The first 4 bits identify the bridge priority, and the remaining 12 bits identify the VLAN participating in this particular STP process.

Using these 12 bits for the extended system ID reduces the bridge priority to 4 bits. This process reserves the rightmost 12 bits for the VLAN ID and the far-left 4 bits for the bridge priority. This explains why the bridge priority value can be configured only in multiples of 4096, or 2^{12} . If the far-left bits are 0001, then the bridge priority is 4096. If the far-left bits are 1111, then the bridge priority is 61440 ($= 15 \times 4096$). The Catalyst 2960 and 3560 Series switches do not allow the configuration of a bridge priority of 65536 ($= 16 \times 4096$) because this priority assumes the use of a fifth bit that is unavailable due to the use of the extended system ID.

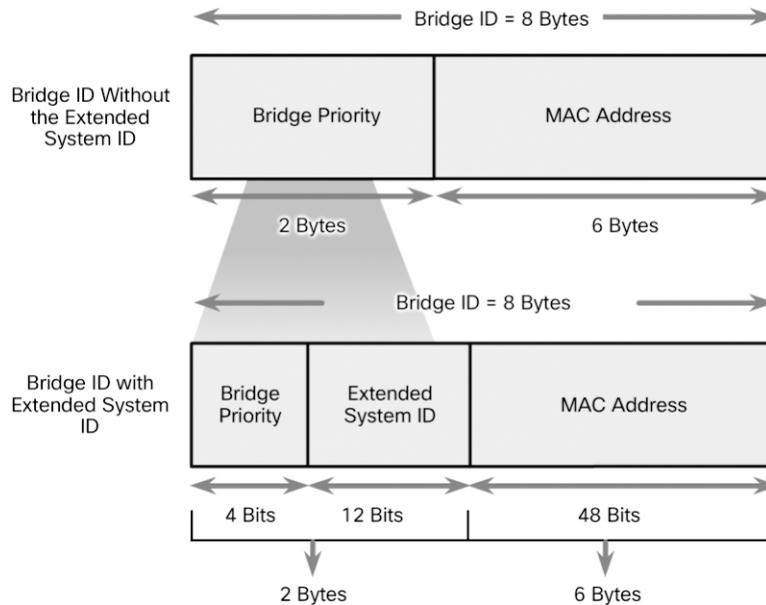


Figure 3-29 BID Fields

The extended system ID value is a decimal value added to the bridge priority value in the BID to identify the priority and VLAN of the BPDU frame.

When two switches are configured with the same priority and have the same extended system ID, the switch with the lowest MAC address has the lower BID. Initially, all switches are configured with the same default priority value. The MAC address is often the deciding factor in which switch becomes the root bridge.

To ensure that the root bridge decision best meets network requirements, it is recommended that the administrator configure the desired root bridge switch with a priority lower than 32768. This also ensures that the addition of new switches to the network does not trigger a new spanning-tree election, which can disrupt network communication while a new root bridge is being selected.

In Figure 3-30, S1 has been configured with a lower priority. Therefore, it is preferred as the root bridge for that spanning-tree instance.

What happens if all switches have the same priority, such as the default priority 32768? The lowest MAC address becomes the deciding factor in which switch becomes the root bridge.

In the scenario in Figure 3-31, S2 becomes the root bridge because it has the lowest MAC address.

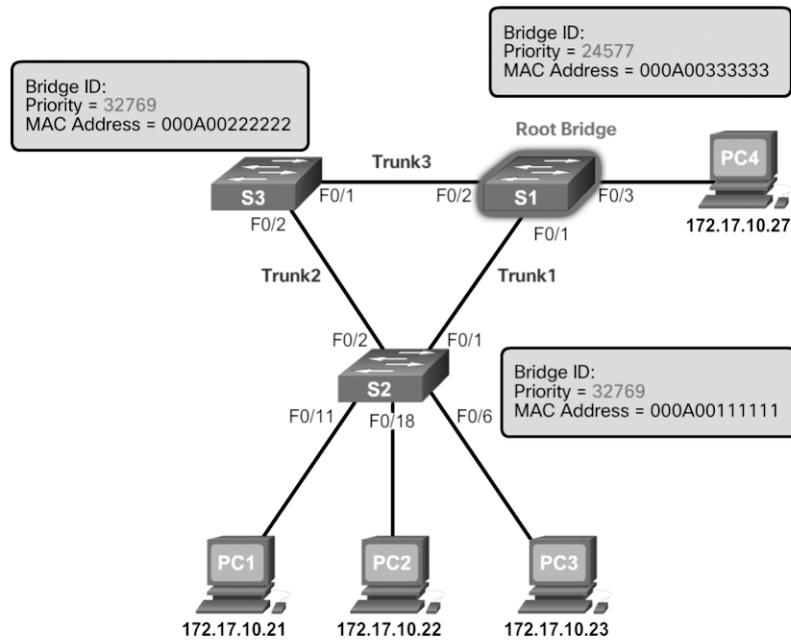


Figure 3-30 Priority-Based Decision

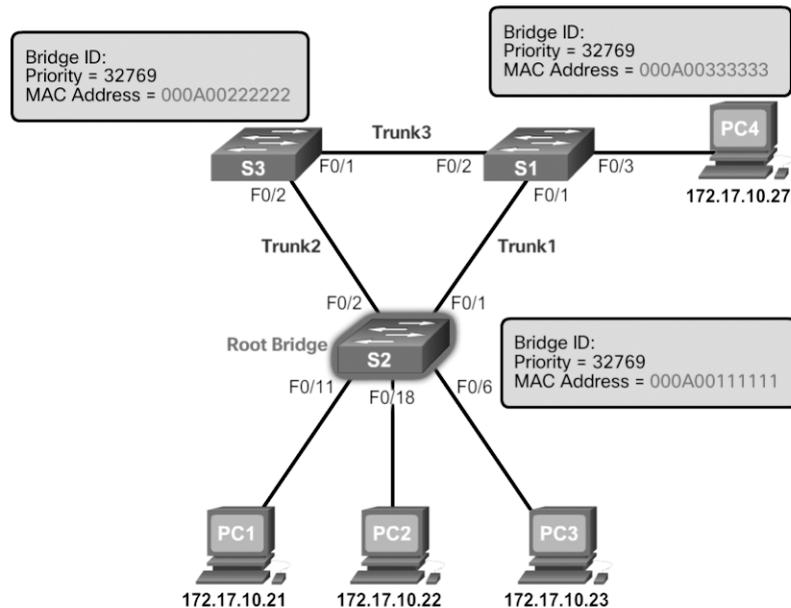


Figure 3-31 MAC Address-Based Decision

Note

In the example shown in Figure 3-31, the priority of all the switches is 32769. The value is based on the 32768 default priority and the VLAN 1 assignment associated with each switch (32768 + 1).

Interactive Graphic**Activity 3.1.2.10: Identify 802.1D Port Rules**

Refer to the online course to complete this activity.

Video**Video Demonstration 3.1.2.11: Observing Spanning Tree Protocol Operation**

Refer to the online course to view this video.

**Lab 3.1.2.12: Building a Switched Network with Redundant Links**

Refer to *Scaling Networks v6 Labs & Study Guide* and the online course to complete this activity.

In this lab, you will complete the following objectives:

- Part 1: Build the Network and Configure Basic Device Settings
 - Part 2: Determine the Root Bridge
 - Part 3: Observe STP Port Selection Based on Port Cost
 - Part 4: Observe STP Port Selection Based on Port Priority
-

Varieties of Spanning Tree Protocols (3.2)

There have been several implementations of STP. In this section, you will learn how different varieties of spanning-tree protocols operate.

Overview (3.2.1)

The focus of this topic is on the different spanning-tree varieties.

Types of Spanning Tree Protocols (3.2.1.1)

Several varieties of spanning-tree protocols have emerged since the original IEEE 802.1D.

The varieties of spanning-tree protocols include the following:

- **STP**—Defined in IEEE *802.1D*, this is the original standard that provided a loop-free topology in a network with redundant links. Also called *Common Spanning Tree (CST)*, it assumed one spanning-tree instance for the entire bridged network, regardless of the number of VLANs.
- **Per-VLAN Spanning Tree (PVST+)**—PVST+ is a Cisco enhancement of STP that provides a separate 802.1D spanning-tree instance for each VLAN configured in the network.
- **Rapid Spanning Tree Protocol (RSTP)**—RSTP is defined in *IEEE 802.1w*. It is an evolution of STP that provides faster convergence than STP.
- **Rapid Per-VLAN Spanning Tree (Rapid PVST+)**—Rapid PVST+ is a Cisco enhancement of RSTP that uses PVST+ and provides a separate instance of 802.1w for each VLAN.
- **Multiple Spanning Tree Protocol (MSTP)**—MSTP, defined in *IEEE 802.1s*, maps multiple VLANs into the same spanning-tree instance. The Cisco implementation of MSTP is often referred to as *Multiple Spanning Tree (MST)*.

A network professional whose duties include switch administration may be required to decide which type of spanning-tree protocol to implement.

Characteristics of the Spanning Tree Protocols (3.2.1.2)

Table 3-3 lists the characteristics of the various STP versions.

Table 3-3 Spanning Tree Protocol Characteristics

STP Version	Characteristics
STP	<ul style="list-style-type: none"> ■ IEEE 802.1D is the original standard. ■ STP creates one spanning-tree instance for the entire bridged network, regardless of the number of VLANs. ■ However, because there is only one root bridge, traffic for all VLANs flows over the same path, which can lead to suboptimal traffic flows. ■ This version is slow to converge. ■ The CPU and memory requirements are lower than for all other STP protocols.

STP Version	Characteristics
PVST+	<ul style="list-style-type: none"> ■ This is a Cisco enhancement of STP that provides a separate STP instance for each VLAN. ■ Each instance supports <i>PortFast</i>, <i>BPDU guard</i>, <i>BPDU filter</i>, <i>root guard</i>, and <i>loop guard</i>. ■ This design allows the spanning tree to be optimized for the traffic of each VLAN. ■ However, CPU and memory requirements are high due to maintaining separate STP instances per VLAN. ■ Convergence is per-VLAN and is slow, like 802.1D.
RSTP	<ul style="list-style-type: none"> ■ 802.1w is an evolution of 802.1D that addresses many convergence issues. ■ Like STP, it provides only a single instance of STP and therefore does not address suboptimal traffic flow issues. ■ The CPU and memory requirements are less than for Rapid PVST+ but more than for 802.1D.
Rapid PVST+	<ul style="list-style-type: none"> ■ This is a Cisco enhancement of RSTP. ■ Rapid PVST+ uses PVST+ and provides a separate instance of 802.1w for each VLAN. ■ Each instance supports PortFast, BPDU guard, BPDU filter, root guard, and loop guard. ■ This version addresses the convergence issues and the suboptimal traffic flow issues. ■ The CPU and memory requirements are the highest of all STP implementations.
MSTP	<ul style="list-style-type: none"> ■ IEEE 802.1s is based on the Cisco Multiple Instance Spanning-Tree Protocol (MISTP) which is often simply referred to as Multiple Spanning Tree (MST). ■ The Cisco implementation is often referred to as Multiple Spanning Tree (MST). ■ MSTP maps multiple VLANs into the same spanning-tree instance. ■ It supports up to 16 instances of RSTP. ■ Each instance supports PortFast, BPDU guard, BPDU filter, root guard, and loop guard. ■ The CPU and memory requirements are less than for Rapid PVST+ but more than for RSTP.

Table 3-4 summarizes the STP characteristics.

Table 3-4 Comparing Spanning Tree Protocols

Protocol	Standard	Resources Needed	Convergence	STP Tree Calculation
STP	IEEE 802.1D	Low	Slow	All VLANs
PVST+	Cisco	High	Slow	Per VLAN
RSTP	IEEE 802.1w	Medium	Fast	All VLANs
Rapid PVST+	Cisco	High	Fast	Per VLAN
MSTP (MST)	IEEE 802.1s, Cisco	Medium or high	Fast	Per instance

Cisco switches running IOS 15.0 or later run PVST+ by default.

Cisco Catalyst switches support PVST+, Rapid PVST+, and MSTP. However, only one version can be active at any time.

**Interactive
Graphic**

Activity 3.2.1.3: Identify Types of Spanning Tree Protocols

Refer to the online course to complete this activity.

PVST+ (3.2.2)

The focus of this topic is on how the default mode of PVST+ on Cisco Catalyst switches operates.

Overview of PVST+ (3.2.2.1)

The original IEEE 802.1D standard defines only one spanning-tree instance for the entire switched network, regardless of the number of VLANs. A network running 802.1D has these characteristics:

- No load sharing is possible. One uplink must block for all VLANs.
- The CPU is spared. Only one instance of spanning tree must be computed.

Cisco developed PVST+ so that a network can run an independent instance of the Cisco implementation of IEEE 802.1D for each VLAN in the network. A PVST+ topology is shown in Figure 3-32.

With PVST+, it is possible for one trunk port on a switch to block for a VLAN while forwarding for other VLANs. PVST+ can be used to manually implement Layer 2 load balancing. The switches in a PVST+ environment require greater CPU process and BPDU bandwidth consumption than a traditional STP because each VLAN runs a separate instance of STP.

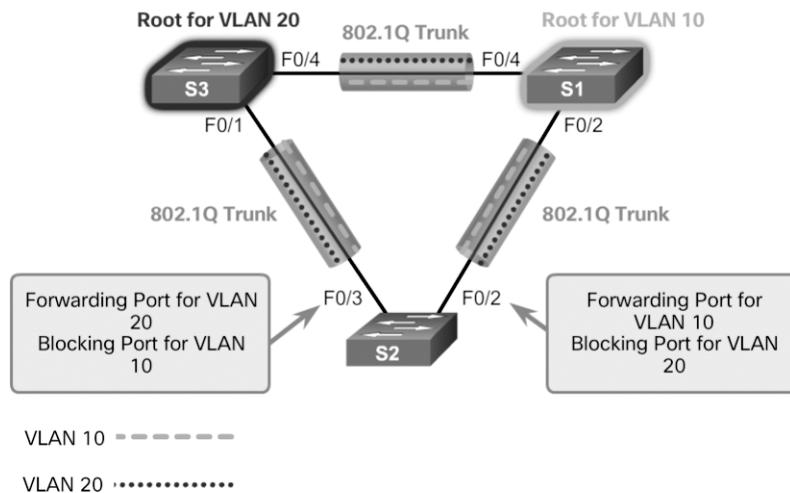


Figure 3-32 PVST+

In a PVST+ environment, spanning-tree parameters can be tuned so that half of the VLANs forward on each uplink trunk. In Figure 3-32, port F0/3 on S2 is the forwarding port for VLAN 20, and F0/2 on S2 is the forwarding port for VLAN 10. This is accomplished by configuring one switch to be elected the root bridge for half of the VLANs in the network and a second switch to be elected the root bridge for the other half of the VLANs. In the figure, S3 is the root bridge for VLAN 20, and S1 is the root bridge for VLAN 10. Having multiple STP root bridges per VLAN increases redundancy in the network.

Networks running PVST+ have these characteristics:

- Optimum load balancing can result.
- One spanning-tree instance for each VLAN maintained can mean a considerable waste of CPU cycles for all the switches in the network (in addition to the bandwidth that is used for each instance to send its own BPDU). This is problematic only if a large number of VLANs are configured.

Port States and PVST+ Operation (3.2.2.2)

STP facilitates the logical loop-free path throughout the broadcast domain. The spanning tree is determined through the information learned by the exchange of the BPDU frames between the interconnected switches. To facilitate the learning of the logical spanning tree, each switch port transitions through five possible port states and three BPDU timers.

The spanning tree is determined immediately after a switch is finished booting up. If a switch port transitions directly from the blocking state to the forwarding state

without information about the full topology during the transition, the port can temporarily create a data loop. For this reason, STP introduced five port states that PVST+ uses as well. Table 3-5 lists and explains the five port states.

Table 3-5 STP Port States

Port State	Characteristics
Blocking state	<ul style="list-style-type: none"> ■ The port is an alternate port and does not participate in frame forwarding. ■ The port receives BPDU frames to determine the location and root ID of the root bridge switch and which port roles each switch port should assume in the final active STP topology.
<i>Listening state</i>	<ul style="list-style-type: none"> ■ Listens for the path to the root. ■ STP has determined that the port can participate in frame forwarding according to the BPDU frames that the switch has received. ■ The switch port receives BPDU frames, transmits its own BPDU frames, and informs adjacent switches that the switch port is preparing to participate in the active topology.
<i>Learning state</i>	<ul style="list-style-type: none"> ■ Learns the MAC addresses. ■ The port prepares to participate in frame forwarding and begins to populate the MAC address table.
<i>Forwarding state</i>	<ul style="list-style-type: none"> ■ The port is considered part of the active topology. ■ It forwards data frames and sends and receives BPDU frames.
<i>Disabled state</i>	<ul style="list-style-type: none"> ■ The Layer 2 port does not participate in spanning tree and does not forward frames. ■ The disabled state is set when the switch port is administratively disabled.

Table 3-6 summarizes the port states which ensure that no loops are created during the creation of the logical spanning tree.

Table 3-6 Port States

Operation Allowed	Port State				
	Blocking	Listening	Learning	Forwarding	Disabled
Can receive and process BPDUs	Yes	Yes	Yes	No	No
Can forward data frames received on the interface	No	No	No	Yes	No

Operation Allowed	Port State				
	Blocking	Listening	Learning	Forwarding	Disabled
Can forward data frames switched from another interface	No	No	No	Yes	No
Can learn MAC addresses	No	No	Yes	Yes	No

Note that the number of ports in each of the various states (blocking, listening, learning, or forwarding) can be displayed with the **show spanning-tree summary** command.

For each VLAN in a switched network, PVST+ performs four steps to provide a loop-free logical network topology:

- Step 1.** It elects one root bridge. Only one switch can act as the root bridge (for a given VLAN). The root bridge is the switch with the lowest bridge ID. On the root bridge, all ports are designated ports (no root ports).
- Step 2.** It selects the root port on each non-root bridge. PVST+ establishes one root port on each non-root bridge for each VLAN. The root port is the lowest-cost path from the non-root bridge to the root bridge, which indicates the direction of the best path to the root bridge. Root ports are normally in the forwarding state.
- Step 3.** It selects the designated port on each segment. On each link, PVST+ establishes one designated port for each VLAN. The designated port is selected on the switch that has the lowest-cost path to the root bridge. Designated ports are normally in the forwarding state and forwarding traffic for the segment.
- Step 4.** It makes the remaining ports in the switched network alternate ports. Alternate ports normally remain in the blocking state to logically break the loop topology. When a port is in the blocking state, it does not forward traffic, but it can still process received BPDU messages.

Extended System ID and PVST+ Operation (3.2.2.3)

In a PVST+ environment, the extended system ID (see Figure 3-33) ensures that each switch has a unique BID for each VLAN.

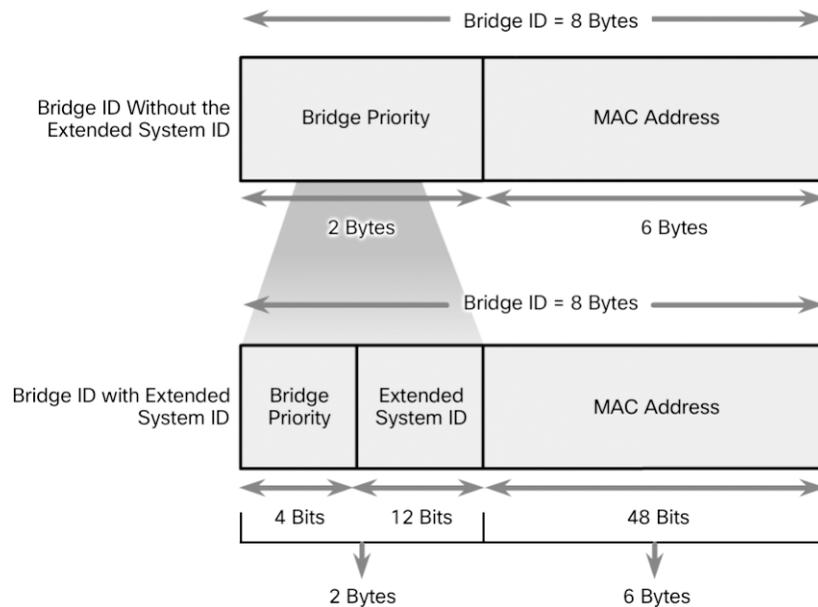


Figure 3-33 PVST+ and the Extended System ID

For example, the VLAN 2 default BID would be 32770 (priority 32768, plus the extended system ID 2). If no priority has been configured, every switch has the same default priority, and the election of the root bridge for each VLAN is based on the MAC address. Because the bridge ID is based on the lowest MAC address, the switch chosen to be root bridge might not be the most powerful or the most optimal switch.

In some situations, an administrator may want a specific switch to be selected as the root bridge. This may be for a variety of reasons, including the following:

- The switch is more optimally located within the LAN design in regards to the majority of traffic flow patterns for a particular VLAN.
- The switch has higher processing power.
- The switch is simply easier to access and manage remotely.

To manipulate the root-bridge election, assign a lower priority to the switch that should be selected as the root bridge for the desired VLAN(s).

Activity 3.2.2.4: Identifying PVST+ Operation

Refer to the online course to complete this activity.

Rapid PVST+ (3.2.3)

The focus of this topic is on how Rapid PVST+ operates.

Overview of Rapid PVST+ (3.2.3.1)

RSTP (IEEE 802.1w) is an evolution of the original 802.1D standard and is incorporated into the IEEE 802.1D-2004 standard. The 802.1w STP terminology remains primarily the same as the original IEEE 802.1D STP terminology. Most parameters have been left unchanged, so users who are familiar with STP can easily configure the new protocol. Rapid PVST+ is the Cisco implementation of RSTP on a per-VLAN basis. An independent instance of RSTP runs for each VLAN.

Figure 3-34 shows a network running RSTP. S1 is the root bridge, with two designated ports in a forwarding state. RSTP supports a new port type. Port F0/3 on S2 is an alternate port in discarding state. Notice that there are no blocking ports. RSTP does not have a blocking port state. RSTP defines port states as discarding, learning, or forwarding.

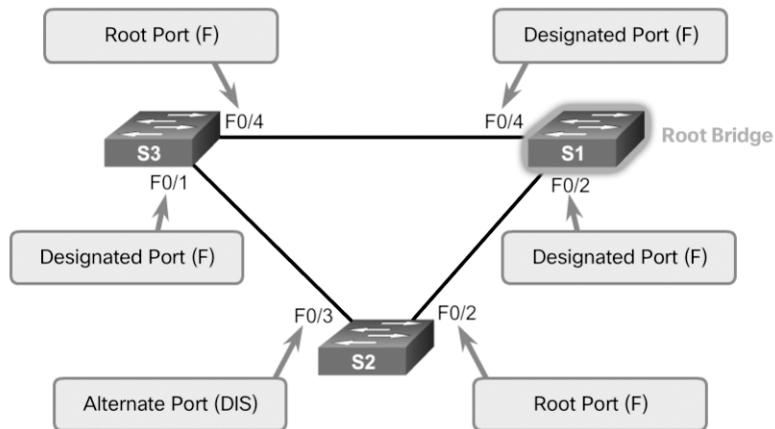


Figure 3-34 RSTP Topology

RSTP speeds the recalculation of the spanning tree when the Layer 2 network topology changes. RSTP can achieve much faster convergence in a properly configured network—sometimes in as little as a few hundred milliseconds.

RSTP redefines the types of ports and their states. If a port is configured to be an alternate port or a backup port, it can immediately change to a forwarding state without waiting for the network to converge.

The following is a brief description of RSTP characteristics:

- RSTP is the preferred protocol for preventing Layer 2 loops in a switched network environment. Many of the differences were established by Cisco proprietary enhancements to the original 802.1D. These enhancements, such as BPDUs carrying and sending information about port roles only to neighboring switches, require no additional configuration and generally perform better than the earlier Cisco proprietary versions. They are now transparent and integrated into the protocol's operation.
- RSTP (802.1w) supersedes the original 802.1D while retaining backward compatibility. Much of the original 802.1D terminology remains, and most parameters are unchanged. In addition, 802.1w is capable of reverting to legacy 802.1D to interoperate with legacy switches on a per-port basis. For example, the RSTP spanning-tree algorithm elects a root bridge in exactly the same way as the original 802.1D.
- RSTP keeps the same BPDU format as the original IEEE 802.1D, except that the version field is set to 2 to indicate RSTP, and the flags field uses all 8 bits.
- RSTP is able to actively confirm that a port can safely transition to the forwarding state without having to rely on a timer configuration.

RSTP BPDUs (3.2.3.2)

RSTP uses type 2, Version 2 BPDUs. The original 802.1D STP uses type 0, Version 0 BPDUs. However, a switch running RSTP can communicate directly with a switch running the original 802.1D STP. RSTP sends BPDUs and populates the flags byte in a slightly different manner than in the original 802.1D:

- Protocol information can be immediately aged on a port if hello packets are not received for three consecutive hello times (six seconds, by default) or if the max age timer expires.
- BPDUs are used as a keepalive mechanism. Therefore, three consecutively missed BPDUs indicate lost connectivity between a bridge and its neighboring root or designated bridge. The fast aging of the information allows failures to be detected quickly.

Note

As with STP, an RSTP switch sends a BPDU with its current information every hello time period (two seconds, by default), even if the RSTP switch does not receive BPDUs from the root bridge.

As shown in Figure 3-35, RSTP uses the flags byte of a Version 2 BPDU:

RSTP Version 2 BPDUs	
Field	Byte Length
Protocol ID=0x0000	2
Protocol Version ID=0x02	1
BPDUs Type=0x02	1
Flags	1
Root ID	8
Root Path Cost	4
Bridge ID	8
Port ID	2
Message Age	2
Max Age	2
Hello Time	2
Forward Delay	2

Flag Field	
Field Bit	Bit
Topology Change	0
Proposal	1
Port Role	2-3
Unknown Port	00
Alternate or Backup Port	01
Root Port	10
Designated Port	11
Learning	4
Forwarding	5
Agreement	6
Topology Change Acknowledgment	7

Figure 3-35 RSTP BPDUs Fields

- Bits 0 and 7 are used for topology change and acknowledgment. They are in the original 802.1D.
- Bits 1 and 6 are used for the proposal agreement process (used for rapid convergence).
- Bits 2 to 5 encode the role and state of the port.
- Bits 4 and 5 are used to encode the port role using a 2-bit code.

Edge Ports (3.2.3.3)

An RSTP *edge port* is a switch port that is never intended to be connected to another switch. It immediately transitions to the forwarding state when enabled.

The RSTP edge port concept corresponds to the PVST+ PortFast feature. An edge port is directly connected to an end station and assumes that no switch device is connected to it. RSTP edge ports should immediately transition to the forwarding state, thereby skipping the time-consuming original 802.1D listening and learning port states.

The Cisco RSTP implementation (Rapid PVST+) maintains the PortFast keyword, using the **spanning-tree portfast** command for edge port configuration. This makes the transition from STP to RSTP seamless.

Figure 3-36 shows examples of ports that can be configured as edge ports.

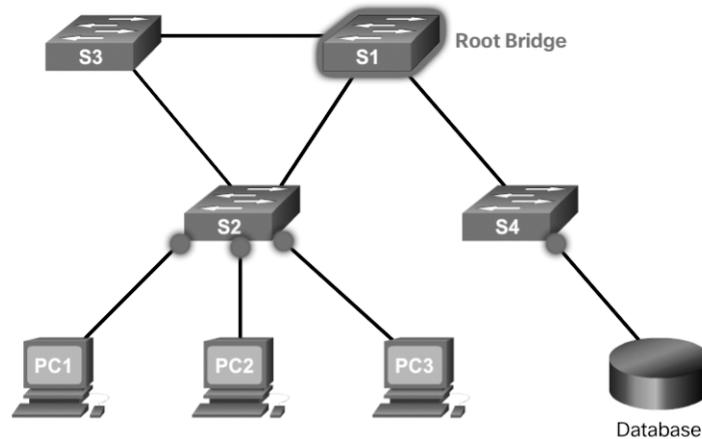


Figure 3-36 Edge Ports

Figure 3-37 shows examples of ports that are non-edge ports.

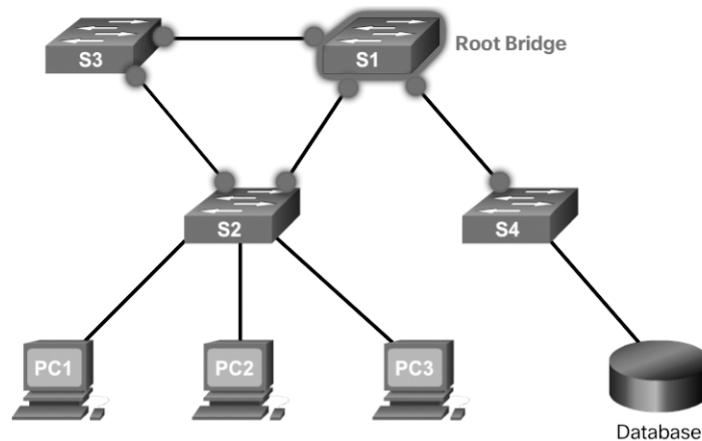


Figure 3-37 Non-Edge Ports

Note

Configuring an edge port to be attached to another switch is not recommended. It can have negative implications for RSTP because a temporary loop may result, possibly delaying the convergence of RSTP.

Link Types (3.2.3.4)

The link type provides a categorization for each port participating in RSTP by using the duplex mode on the port. Depending on what is attached to each port, two different link types can be identified:

- **Point-to-point**—A port operating in full-duplex mode typically connects a switch to a switch and is a candidate for a rapid transition to a forwarding state.
- **Shared**—A port operating in half-duplex mode connects a switch to a hub that attaches multiple devices.

Figure 3-38 displays the various RSTP port assignments.

The link type can determine whether the port can immediately transition to a forwarding state, assuming that certain conditions are met. These conditions are different for edge ports and non-edge ports. Non-edge ports are categorized into two link types: point-to-point and shared.

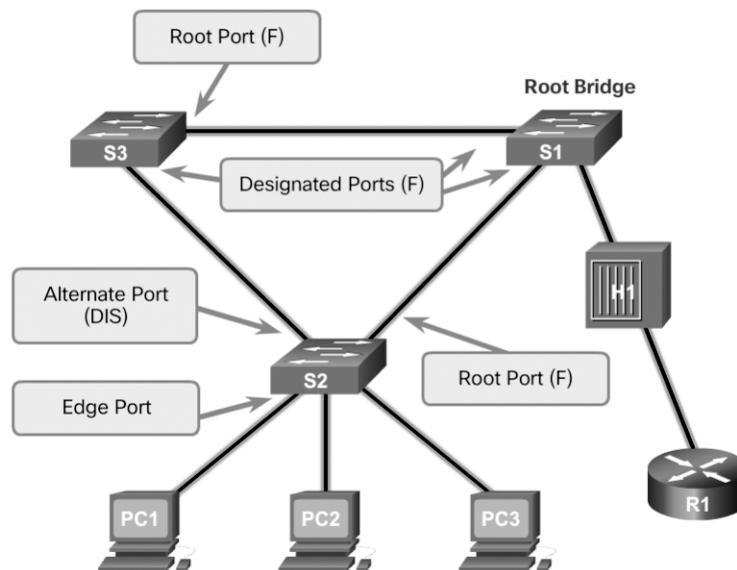


Figure 3-38 Link Types

The link type is automatically determined but can be overridden with an explicit port configuration, using the `spanning-tree link-type { point-to-point | shared }` command.

Characteristics of port roles, with regard to link types, include the following:

- Edge port connections and point-to-point connections are candidates for rapid transition to a forwarding state. However, before the `link-type` parameter is considered, RSTP must determine the port role.

- Root ports do not use the **link-type** parameter. Root ports are able to make a rapid transition to the forwarding state as soon as the port is in sync (that is, receives a BPDU from the root bridge).
- Alternate and backup ports do not use the **link-type** parameter in most cases.
- Designated ports make the most use of the **link-type** parameter. A rapid transition to the forwarding state for the designated port occurs only if the **link-type** parameter is set to **point-to-point**.

**Interactive
Graphic**

Activity 3.2.3.5: Identify Port Roles in Rapid PVST+

Refer to the online course to complete this activity.

**Interactive
Graphic**

Activity 3.2.3.6: Compare PVST+ and Rapid PVST+

Refer to the online course to complete this activity.

Spanning Tree Configuration (3.3)

In this section, you will learn how to implement PVST+ and Rapid PVST+ in a switched LAN environment.

PVST+ Configuration (3.3.1)

The focus of this topic is on how to configure PVST+ in a switched LAN environment.

Catalyst 2960 Default Configuration (3.3.1.1)

Table 3-7 shows the default spanning-tree configuration for a Cisco Catalyst 2960 Series switch. Notice that the default spanning-tree mode is PVST+.

Table 3-7 Default Switch Configuration

Feature	Default Setting
Enable state	Enabled on VLAN 1
Spanning-tree mode	PVST+ (Rapid PVST+ and MSTP are disabled.)
Switch priority	32768

Feature	Default Setting
Spanning-tree port priority (configurable on a per-interface basis)	128
Spanning-tree port cost (configurable on a per-interface basis)	1000 Mbps: 4
	100 Mbps: 19
	10 Mbps: 100
Spanning-tree VLAN port priority (configurable on a per-VLAN basis)	128
Spanning-tree VLAN port cost (configurable on a per-VLAN basis)	1000 Mbps: 4
	100 Mbps: 19
	10 Mbps: 100
Spanning-tree timers	Hello time: 2 seconds
	Forward-delay time: 15 seconds
	Maximum-aging time: 20 seconds
	Transmit hold count: 6 BPDUs

Configuring and Verifying the Bridge ID (3.3.1.2)

When an administrator wants a specific switch to become a root bridge, the bridge priority value must be adjusted to ensure that it is lower than the bridge priority values of all the other switches on the network. There are two different methods to configure the bridge priority value on a Cisco Catalyst switch.

Method 1

To ensure that a switch has the lowest bridge priority value, use the **spanning-tree vlan *vlan-id* root primary** command in global configuration mode. The priority for the switch is set to the predefined value of 24,576 or to the highest multiple of 4096 less than the lowest bridge priority detected on the network.

If an alternate root bridge is desired, use the **spanning-tree vlan *vlan-id* root secondary** global configuration mode command. This command sets the priority for the switch to the predefined value 28,672. This ensures that the alternate switch becomes the root bridge if the primary root bridge fails. This assumes that the rest of the switches in the network have the default 32,768 priority value defined.

In Figure 3-39, S1 has been assigned as the primary root bridge, using the **spanning-tree vlan 1 root primary** command, and S2 has been configured as the secondary root bridge, using the **spanning-tree vlan 1 root secondary** command.

Method 2

Another method for configuring the bridge priority value is by using the **spanning-tree vlan *vlan-id* priority *value*** global configuration mode command. This command gives more granular control over the bridge priority value. The priority value is configured in increments of 4096 between 0 and 61,440.

In the example in Figure 3-39, S3 has been assigned a bridge priority value of 24,576, using the **spanning-tree vlan 1 priority 24576** command.

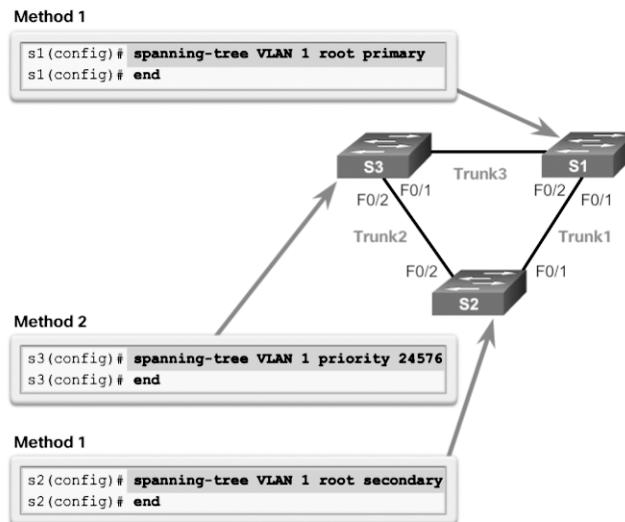


Figure 3-39 Configuring the Bridge ID

To verify the bridge priority of a switch, use the **show spanning-tree** command. In Example 3-4, the priority of the switch has been set to 24,576. Also notice that the switch is designated as the root bridge for the spanning-tree instance.

Example 3-4 Verifying the Root Bridge and BID

```
S3# show spanning-tree
VLAN0001
  Spanning tree enabled protocol ieee
  Root ID    Priority    24577
            Address     000A.0033.0033
            This bridge is the root
  Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
```

Bridge ID	Priority	24577 (priority 24576 sys-id-ext 1)			
	Address	000A.0033.3333			
	Hello Time	2 sec	Max Age 20 sec	Forward Delay 15 sec	
	Aging Time	300			
Interface	Role	Sts	Cost	Prio.Nbr	Type

Fa0/1	Desg	FWD	4	128.1	P2p
Fa0/2	Desg	FWD	4	128.2	P2p

PortFast and BPDU Guard (3.3.1.3)

PortFast is a Cisco feature for PVST+ environments. When a switch port is configured with PortFast, that port transitions from blocking to forwarding state immediately, bypassing the usual 802.1D STP transition states (the listening and learning states). As shown in Figure 3-40, you can use PortFast on access ports to allow these devices to connect to the network immediately rather than wait for IEEE 802.1D STP to converge on each VLAN. Access ports are ports that are connected to a single workstation or to a server.

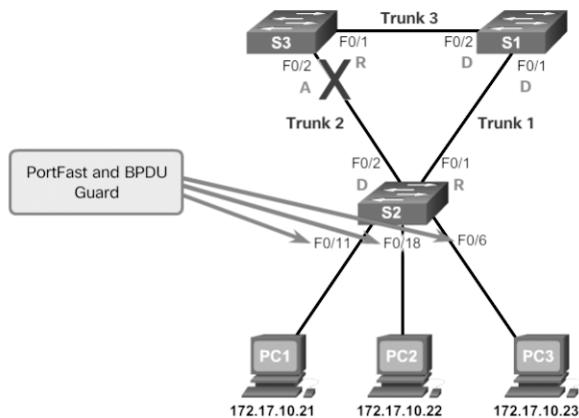


Figure 3-40 PortFast and BPDU Guard Topology

In a valid PortFast configuration, BPDUs should never be received because that would indicate that another bridge or switch is connected to the port, potentially causing a spanning-tree loop. Cisco switches support a feature called BPDU guard. When it is enabled, BPDU guard puts the port in an errdisabled (error-disabled) state on receipt of a BPDU. This effectively shuts down the port. The BPDU guard feature provides a secure response to invalid configurations because you must manually put the interface back into service.

Cisco PortFast technology is useful for DHCP. Without PortFast, a PC can send a DHCP request before the port is in forwarding state, denying the host from getting a usable IP address and other information. Because PortFast immediately changes the state to forwarding, the PC always gets a usable IP address (if the DHCP server has been configured correctly and communication with the DHCP server has occurred).

Note

Because the purpose of PortFast is to minimize the time that access ports must wait for spanning tree to converge, it should be used only on access ports. If you enable PortFast on a port connecting to another switch, you risk creating a spanning-tree loop.

To configure PortFast on a switch port, enter the **spanning-tree portfast** interface configuration mode command on each interface on which PortFast is to be enabled, as shown in Example 3-5.

Example 3-5 Configuring PortFast

```
S2(config)# interface FastEthernet 0/11
S2(config-if)# spanning-tree portfast
%Warning: portfast should only be enabled on ports connected to a single
          host. Connecting hubs, concentrators, switches, bridges, etc... to this
          interface when portfast is enabled, can cause temporary bridging loops.
          Use with CAUTION

%Portfast has been configured on FastEthernet0/11 but will only
          have effect when the interface is in a non-trunking mode.

S2(config-if)#
```

The **spanning-tree portfast default** global configuration mode command enables PortFast on all non-trunking interfaces.

To configure BPDU guard on a Layer 2 access port, use the **spanning-tree bpduguard enable** interface configuration mode command, as shown in Example 3-6.

Example 3-6 Configuring and Verifying BPDU Guard

```
S2(config-if)# spanning-tree bpduguard enable
S2(config-if)# end
S2#
S2# show running-config interface f0/11
interface FastEthernet0/11
spanning-tree portfast
spanning-tree bpduguard enable

S2#
```

The **spanning-tree portfast bpduguard default** global configuration command enables BPDU guard on all PortFast-enabled ports.

Notice in Example 3-6 how the **show running-config interface** command can be used to verify that PortFast and BPDU guard have been enabled for a switch port. PortFast and BPDU guard are disabled, by default, on all interfaces.

PVST+ Load Balancing (3.3.1.4)

The topology in Figure 3-41 shows three switches with 802.1Q trunks connecting them.

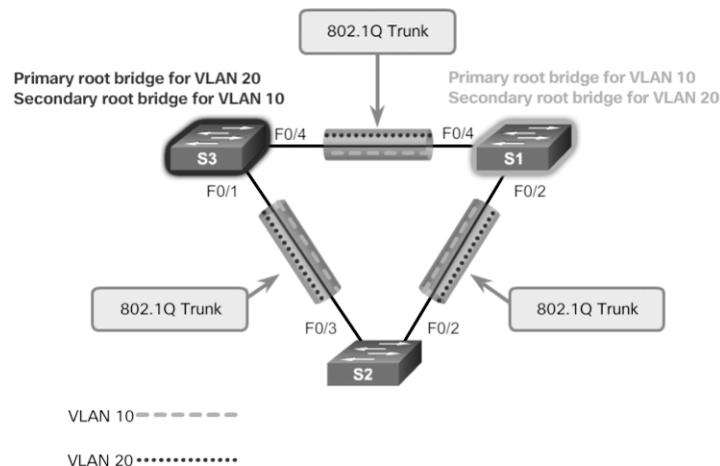


Figure 3-41 PVST+ Configuration Topology

Two VLANs, 10 and 20, are being trunked across these links. The goal is to configure S3 as the root bridge for VLAN 20 and S1 as the root bridge for VLAN 10. Port F0/3 on S2 is the forwarding port for VLAN 20 and the blocking port for VLAN 10. Port F0/2 on S2 is the forwarding port for VLAN 10 and the blocking port for VLAN 20.

In addition to establishing a root bridge, it is also possible to establish a secondary root bridge. A secondary root bridge is a switch that may become the root bridge for a VLAN if the primary root bridge fails. Assuming that the other bridges in the VLAN retain their default STP priority, this switch becomes the root bridge if the primary root bridge fails.

Configuring PVST+ on this topology involves the following steps:

- Step 1.** Select the switches you want for the primary and secondary root bridges for each VLAN. For example, in Figure 3-41, S3 is the primary bridge for VLAN 20, and S1 is the secondary bridge for VLAN 20.

Step 2. As shown in Example 3-7, configure S3 to be a primary bridge for VLAN 10 and the secondary bridge for VLAN 20 by using the `spanning-tree vlan number root { primary | secondary }` command.

Example 3-7 Configuring Primary and Secondary Root Bridges for Each VLAN on S3

```
S3(config)# spanning-tree vlan 20 root primary
S3(config)# spanning-tree vlan 10 root secondary
```

Step 3. As shown in Example 3-8, configure S1 to be a primary bridge for VLAN 20 and the secondary bridge for VLAN 10.

Example 3-8 Configuring Primary and Secondary Root Bridges for Each VLAN on S1

```
S1(config)# spanning-tree vlan 10 root primary
S1(config)# spanning-tree vlan 20 root secondary
```

Another way to specify the root bridge is to set the spanning-tree priority on each switch to the lowest value so that the switch is selected as the primary bridge for its associated VLAN, as shown in Example 3-9.

Example 3-9 Configuring the Lowest Possible Priority to Ensure That a Switch Is Root

```
S3(config)# spanning-tree vlan 20 priority 4096
S1(config)# spanning-tree vlan 10 priority 4096
```

The switch priority can be set for any spanning-tree instance. This setting affects the likelihood that a switch is selected as the root bridge. A lower value increases the probability that the switch is selected. The range is 0 to 61,440, in increments of 4096; all other values are rejected. For example, a valid priority value is $4096 \times 2 = 8192$.

As shown in Example 3-10, the `show spanning-tree active` command displays spanning-tree configuration details for the active interfaces only.

The output shown is for S1 configured with PVST+. A number of Cisco IOS command parameters are associated with the `show spanning-tree` command.

In Example 3-11, the output shows that the priority for VLAN 10 is 4096, the lowest of the three respective VLAN priorities.

Example 3-10 Verifying STP Active Interfaces

```

S1# show spanning-tree active
<output omitted>
VLAN0010
  Spanning tree enabled protocol ieee
  Root ID      Priority      4106
  Address      ec44.7631.3880
  This bridge is the root
  Hello Time   2 sec    Max Age 20 sec    Forward Delay 15 sec

  Bridge ID    Priority      4106    (priority 4096 sys-id-ext 10)
  Address      ec44.7631.3880
  Hello Time   2 sec    Max Age 20 sec    Forward Delay 15 sec
  Aging Time   300 sec

Interface          Role Sts Cost          Prio.Nbr Type
-----
Fa0/3              Desg FWD 19            128.5   P2p
Fa0/4              Desg FWD 19            128.6   P2p

```

Example 3-11 Verifying the S1 STP Configuration

```

S1# show running-config | include span
spanning-tree mode pvst
spanning-tree extend system-id
spanning-tree vlan 1 priority 24576
spanning-tree vlan 10 priority 4096
spanning-tree vlan 20 priority 28672

```

Packet Tracer
 Activity
Packet Tracer 3.3.1.5: Configuring PVST+

In this activity, you will configure VLANs and trunks and examine and configure the Spanning Tree Protocol primary and secondary root bridges. You will also optimize the switched topology by using PVST+, PortFast, and BPDU guard.

Rapid PVST+ Configuration (3.3.2)

Rapid PVST+ is the Cisco implementation of RSTP. It supports RSTP on a per-VLAN basis. The focus of this topic is on how to configure Rapid PVST+ in a switched LAN environment.

Spanning Tree Mode (3.3.2.1)

Rapid PVST+ commands control the configuration of VLAN spanning-tree instances. A spanning-tree instance is created when an interface is assigned to a VLAN, and is removed when the last interface is moved to another VLAN. In addition, you can configure STP switch and port parameters before a spanning-tree instance is created. These parameters are applied when a spanning-tree instance is created.

Use the **spanning-tree mode rapid-pvst** global configuration mode command to enable Rapid PVST+. Optionally, you can also identify interswitch links as point-to-point links by using the **spanning-tree link-type point-to-point** interface configuration command. When specifying an interface to configure, valid interfaces include physical ports, VLANs, and port channels.

To reset and reconverge STP, use the **clear spanning-tree detected-protocols** privileged EXEC mode command.

To illustrate how to configure Rapid PVST+, refer to the topology in Figure 3-42.

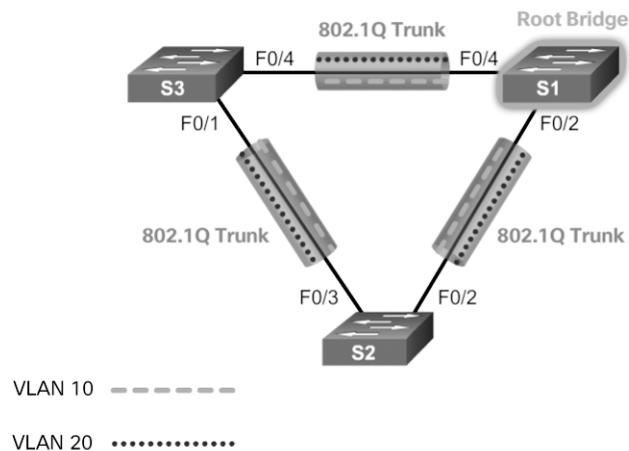


Figure 3-42 Rapid PVST+ Topology

Note

The default spanning-tree configuration on a Catalyst 2960 Series switch is PVST+. A Catalyst 2960 switch supports PVST+, Rapid PVST+, and MST, but only one version can be active for all VLANs at any time.

Example 3-12 displays the commands to configure Rapid PVST+ on S1.

Example 3-12 Configuring Rapid PVST+ on S1

```

S1# configure terminal
S1(config)# spanning-tree mode rapid-pvst
S1(config)# spanning-tree vlan 1 priority 24576
S1(config)# spanning-tree vlan 10 priority 4096
S1(config)# spanning-tree vlan 20 priority 28672
S1(config)# interface f0/2
S1(config-if)# spanning-tree link-type point-to-point
S1(config-if)# end
S1# clear spanning-tree detected-protocols

```

In Example 3-13, the **show spanning-tree vlan 10** command shows the spanning-tree configuration for VLAN 10 on switch S1.

Example 3-13 Verifying That VLAN 10 Is Using RSTP

```

S1# show spanning-tree vlan 10

VLAN0010
Spanning tree enabled protocol rstp
  Root ID    Priority    4106
             Address    ec44.7631.3880
             This bridge is the root
             Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    4106  (priority 4096 sys-id-ext 10)
             Address    ec44.7631.3880
             Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300 sec

Interface                Role Sts Cost          Prio.Nbr Type
-----
Fa0/3                    Desg FWD 19           128.5  P2p Peer (STP)
Fa0/4                    Desg FWD 19           128.6  P2p Peer (STP)

```

In the output, the statement “Spanning tree enabled protocol rstp” indicates that S1 is running Rapid PVST+. Notice that the BID priority is set to 4096. Because S1 is the root bridge for VLAN 10, all of its interfaces are designated ports.

In Example 3-14, the **show running-config** command is used to verify the Rapid PVST+ configuration on S1.

Example 3-14 Verifying the Rapid PVST+ Configuration

```
S1# show running-config | include span
spanning-tree mode rapid-pvst
spanning-tree extend system-id
spanning-tree vlan 1 priority 24576
spanning-tree vlan 10 priority 4096
spanning-tree vlan 20 priority 28672
spanning-tree link-type point-to-point
```

Note

Generally, it is unnecessary to configure the **point-to-point link-type** parameter for Rapid PVST+ because it is unusual to have a shared link type. In most cases, the only difference between configuring PVST+ and Rapid PVST+ is the **spanning-tree mode rapid-pvst** command.

Packet Tracer
Activity**Packet Tracer 3.3.2.2: Configuring Rapid PVST+**

In this activity, you will configure VLANs and trunks and examine and configure the spanning-tree primary and secondary root bridges. You will also optimize it by using rapid PVST+, PortFast, and BPDU guard.

**Lab 3.3.2.3: Configuring Rapid PVST+, PortFast, and BPDU Guard**

Refer to *Scaling Networks v6 Labs & Study Guide* and the online course to complete this activity.

In this lab, you will complete the following objectives:

- Part 1: Build the Network and Configure Basic Device Settings
- Part 2: Configure VLANs, Native VLAN, and Trunks
- Part 3: Configure the Root Bridge and Examine PVST+ Convergence
- Part 4: Configure Rapid PVST+, PortFast, BPDU Guard, and Examine Convergence

STP Configuration Issues (3.3.3)

The focus of this topic is on how to analyze common STP configuration issues.

Analyzing the STP Topology (3.3.3.1)

To analyze the STP topology, follow these steps, as shown in the logic diagram in Figure 3-43:

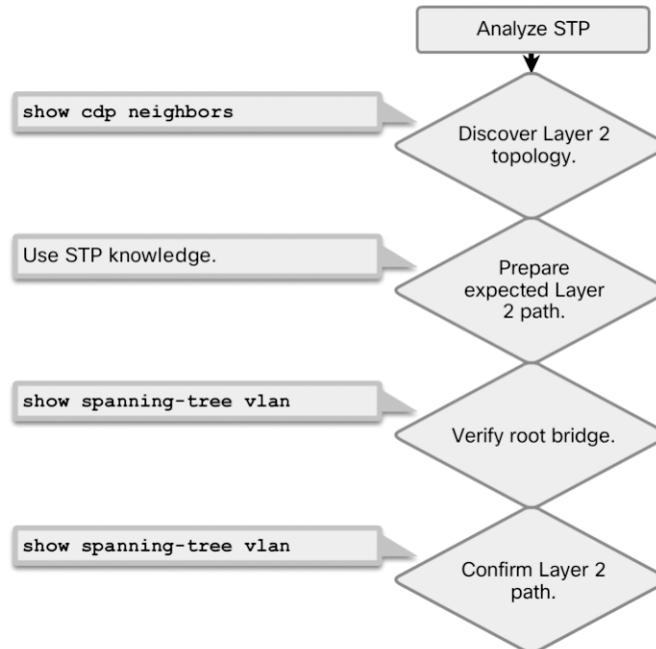


Figure 3-43 Analyzing the STP Topology

- Step 1.** Discover the Layer 2 topology. Use network documentation if it exists or use the **show cdp neighbors** command to discover the Layer 2 topology.
- Step 2.** After discovering the Layer 2 topology, use STP knowledge to determine the expected Layer 2 path. It is necessary to know which switch is the root bridge.
- Step 3.** Use the **show spanning-tree vlan** command to determine which switch is the root bridge.
- Step 4.** Use the **show spanning-tree vlan** command on all switches to find out which ports are in blocking or forwarding state and confirm your expected Layer 2 path.

Expected Topology versus Actual Topology (3.3.3.2)

In many networks, the optimal STP topology is determined as part of the network design and then implemented through manipulation of STP priority and cost values, as shown in Figure 3-44.

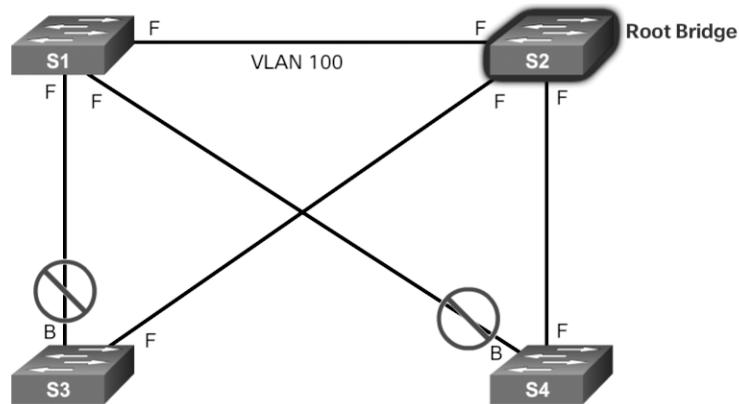


Figure 3-44 Verifying That Actual Topology Matches Expected Topology

Situations may occur in which STP was not considered in the network design and implementation, or in which it was considered or implemented before the network underwent significant growth and change. In such situations, it is important to know how to analyze the STP topology in the operational network.

A big part of troubleshooting consists of comparing the actual state of the network against the expected state of the network and spotting the differences to gather clues about the troubleshooting problem. A network professional should be able to examine the switches and determine the actual topology, as well as understand what the underlying spanning-tree topology should be.

Overview of Spanning Tree Status (3.3.3.3)

Using the **show spanning-tree** command without specifying any additional options provides a quick overview of the status of STP for all VLANs that are defined on a switch.

Use the **show spanning-tree vlan *vlan_id*** command to get STP information for a particular VLAN. Use this command to get information about the role and status of each port on the switch. If you are interested only in a particular VLAN, limit the scope of this command by specifying that VLAN as an option, as shown for VLAN 100 in Figure 3-45.

The output on switch S1 in this example shows all three ports in the forwarding (FWD) state and the roles of the three ports as either designated ports or root ports. Any ports being blocked display the output status as “BLK.”

The output also gives information about the BID of the local switch and the root ID, which is the BID of the root bridge.

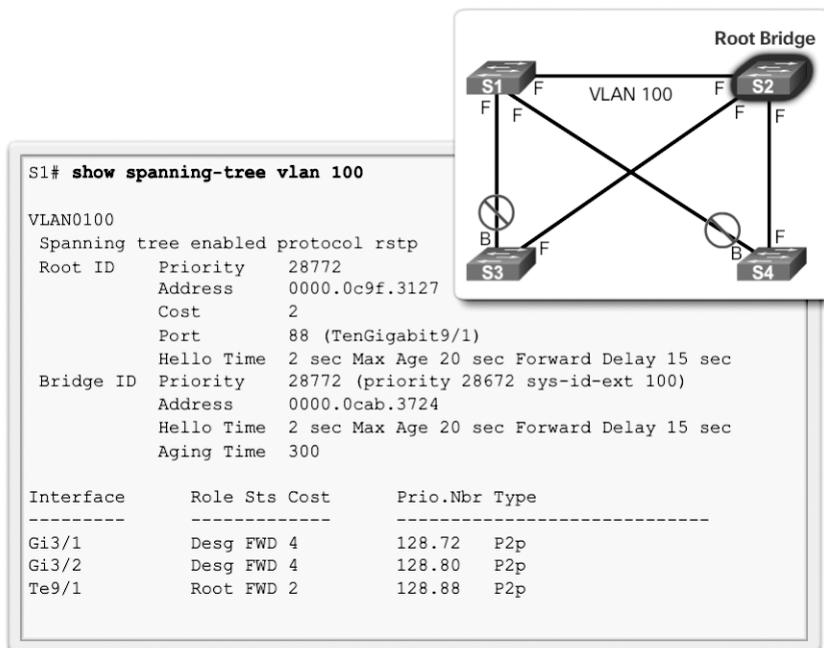


Figure 3-45 Overview of STP Status

Spanning Tree Failure Consequences (3.3.3.4)

Figure 3-46 shows a functional STP network. But what happens when there is an STP failure?

There are two types of STP failure. First, STP might erroneously block ports that should have gone into the forwarding state. Connectivity might be lost for traffic that would normally pass through this switch, but the rest of the network remains unaffected. Second, STP might erroneously move one or more ports into the forwarding state, as shown for S4 in Figure 3-47.

Remember that an Ethernet frame header does not include a TTL field, which means that any frame that enters a bridging loop continues to be forwarded by the switches indefinitely. The only exceptions are frames that have their destination address recorded in the MAC address table of the switches. These frames are simply forwarded to the port that is associated with the MAC address and do not enter a loop. However, any frame that is flooded by a switch enters the loop. This may include broadcasts, multicasts, and unicasts with a globally unknown destination MAC address.

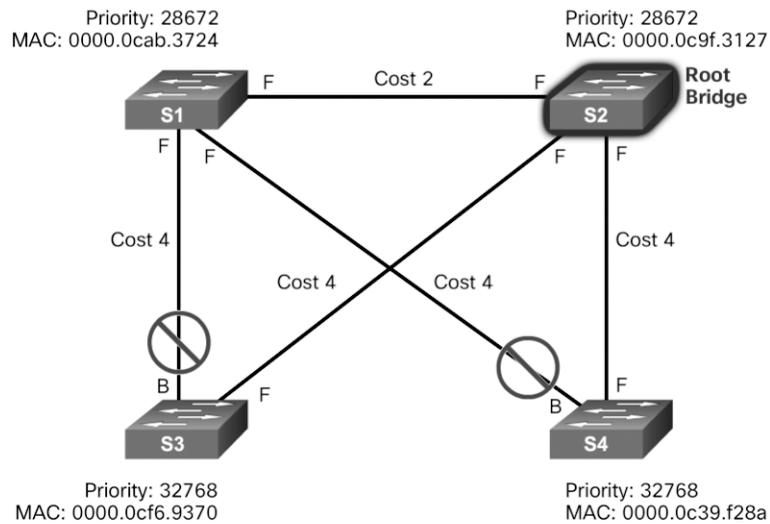


Figure 3-46 STP Switch Topology

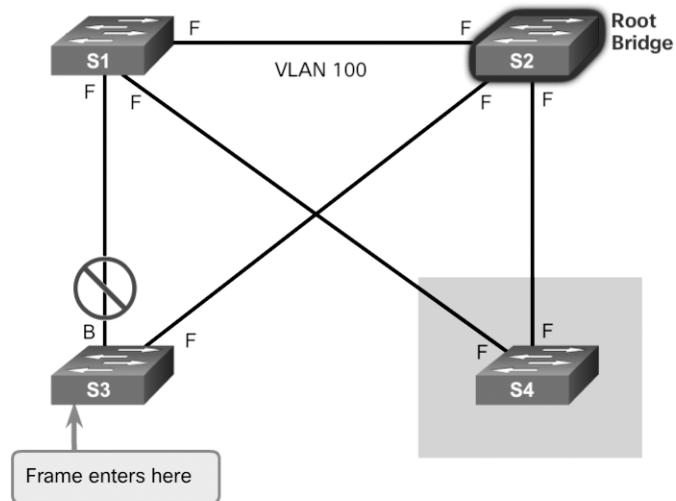


Figure 3-47 Erroneous Transition to Forwarding

Figure 3-48 shows the consequences and corresponding symptoms of STP failure.

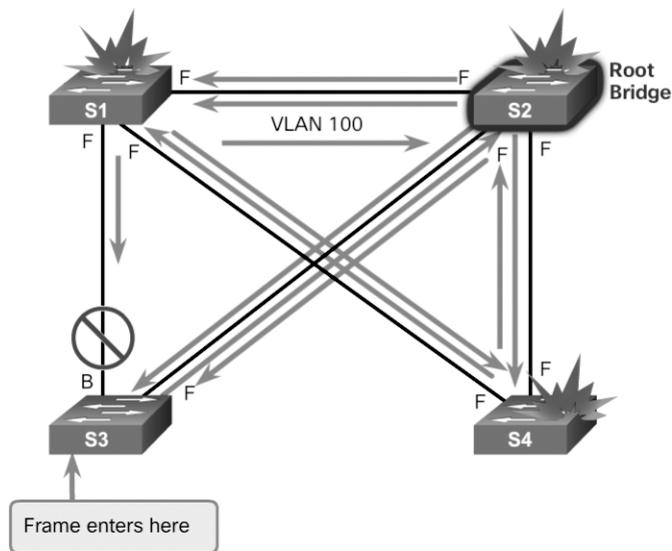


Figure 3-48 Consequences of STP Failure Are Severe

The load on all links in the switched LAN quickly starts increasing as more and more frames enter the loop. This problem is not limited to the links that form the loop but also affects any other links in the switched domain because the frames are flooded on all links. When the spanning-tree failure is limited to a single VLAN only, links in that VLAN are affected. Switches and trunks that do not carry that VLAN operate normally.

If the spanning-tree failure has created a bridging loop, traffic increases exponentially. The switches then flood the broadcasts out multiple ports. This creates copies of the frames every time the switches forward them.

When control plane traffic (for example, routing messages) starts entering the loop, the devices that are running these protocols quickly start getting overloaded. Their CPUs approach 100 percent utilization while they are trying to process an ever-increasing load of control plane traffic. In many cases, the earliest indication of this broadcast storm in progress is that routers or Layer 3 switches report control plane failures and that they are running at a high CPU load.

The switches experience frequent MAC address table changes. If a loop exists, a switch may see a frame with a certain source MAC address coming in on one port and then see another frame with the same source MAC address coming in on a different port a fraction of a second later. This causes the switch to update the MAC address table twice for the same MAC address.

Repairing a Spanning Tree Problem (3.3.3.5)

One way to correct spanning-tree failure is to manually remove redundant links in the switched network, either physically or through configuration, until all loops are eliminated from the topology. When the loops are broken, the traffic and CPU loads should quickly drop to normal levels, and connectivity to devices should be restored.

Although this intervention restores connectivity to the network, it is not the end of the troubleshooting process. All redundancy from the switched network has been removed, and now the redundant links must be restored.

If the underlying cause of the spanning-tree failure has not been fixed, chances are that restoring the redundant links will trigger a new broadcast storm. Before restoring the redundant links, determine and correct the cause of the spanning-tree failure. Carefully monitor the network to ensure that the problem is fixed.

Interactive Graphic

Activity 3.3.3.6: Troubleshoot STP Configuration Issues

Refer to the online course to complete this activity.

Switch Stacking and Chassis Aggregation (3.3.4)

The focus of this topic is to explain the value of switch stacking and chassis aggregation in a small switched LAN.

Switch Stacking Concepts (3.3.4.1)

A switch stack can consist of up to nine Catalyst 3750 switches connected through their StackWise ports. One of the switches controls the operation of the stack and is called the *stack master*. The stack master and the other switches in the stack are stack members.

Figure 3-49 shows the backplane of four Catalyst 3750 switches and how they are connected in a stack.

Every member is uniquely identified by its own stack member number. All members are eligible masters. If the master becomes unavailable, there is an automatic process to elect a new master from the remaining stack members. One of the factors is the stack member priority value. The switch with the highest stack member priority value becomes the master.

Layer 2 and Layer 3 protocols present the entire switch stack as a single entity to the network. One of the primary benefits of switch stacks is that you manage the stack through a single IP address. The IP address is a system-level setting and is not specific to the master or to any other member. You can manage the stack through the same IP address even if you remove the master or any other member from the stack.

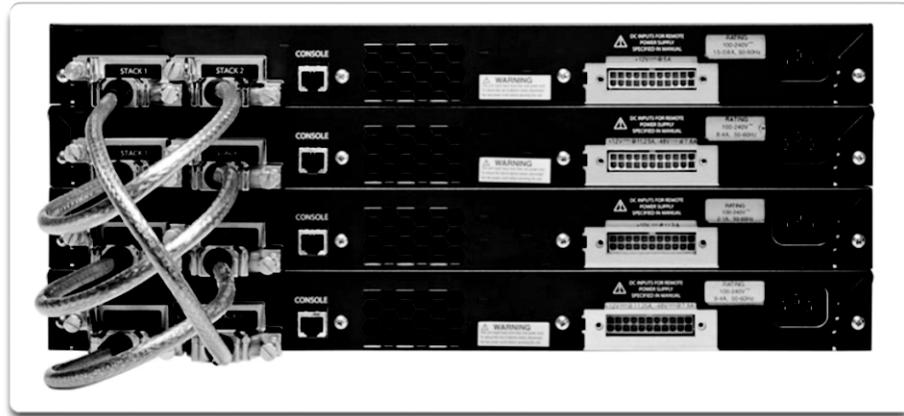


Figure 3-49 Cisco Catalyst 3750 Switch Stack

The master contains the saved and running configuration files for the stack. Therefore, there is only one configuration file to manage and maintain. The configuration files include the system-level settings for the stack and the interface-level settings for each member. Each member has a current copy of these files for backup purposes.

The switch is managed as a single switch, including passwords, VLANs, and interfaces. Example 3-15 shows the interfaces on a switch stack with four 52-port switches. Notice that the first number after the interface type is the stack member number.

Example 3-15 Switch Stack Interfaces

```
Switch# show running-config | begin interface
interface GigabitEthernet1/0/1
!
interface GigabitEthernet1/0/2
!
interface GigabitEthernet1/0/3
!
<output omitted>
!
interface GigabitEthernet1/0/52
!
interface GigabitEthernet2/0/1
!
interface GigabitEthernet2/0/2
!
<output omitted>
!
interface GigabitEthernet2/0/52
!
```

```

interface GigabitEthernet3/0/1
!
interface GigabitEthernet3/0/2
!
<output omitted>
!
interface GigabitEthernet3/0/52
!
interface GigabitEthernet4/0/1
!
interface GigabitEthernet4/0/2
!
<output omitted>
!
interface GigabitEthernet4/0/52
!
Switch#

```

Spanning Tree and Switch Stacks (3.3.4.2)

Another benefit to switch stacking is the ability to add more switches to a single STP instance without increasing the *STP diameter*. The diameter is the maximum number of switches that data must cross to connect any two switches. The IEEE recommends a maximum diameter of seven switches for the default STP timers. For example, in Figure 3-50, the diameter from S1-4 to S3-4 is nine switches. This design violates the IEEE recommendation.

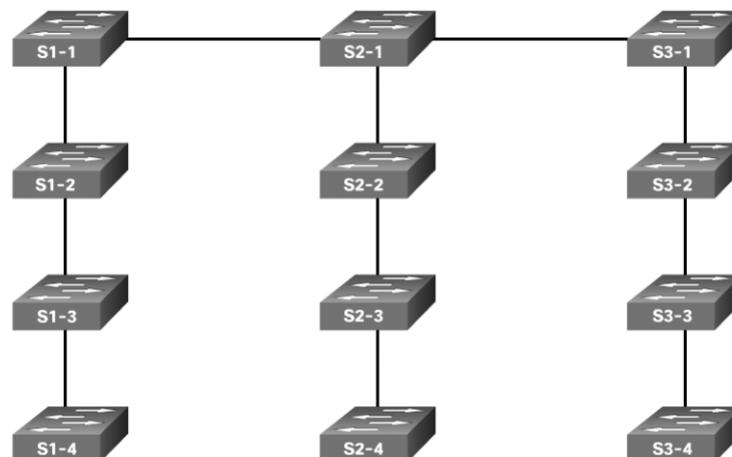


Figure 3-50 Diameter Greater Than 7

The recommended diameter is based on default STP timer values, which are as follows:

- **Hello Timer (2 seconds)**—The interval between BPDU updates.
- **Max Age Timer (20 seconds)**—The maximum length of time a switch saves BPDU information.
- **Forward Delay Timer (15 seconds)**—The time spent in the listening and learning states.

Note

The formulas used to calculate the diameter are beyond the scope of this course. Refer to the following Cisco document for more information: www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/19120-122.html.

Switch stacks help maintain or reduce the impact of diameter on STP reconvergence. In a switch stack, all switches use the same bridge ID for a given spanning-tree instance. This means that, if the switches are stacked, as shown in Figure 3-51, the maximum diameter becomes 3 instead of 9.

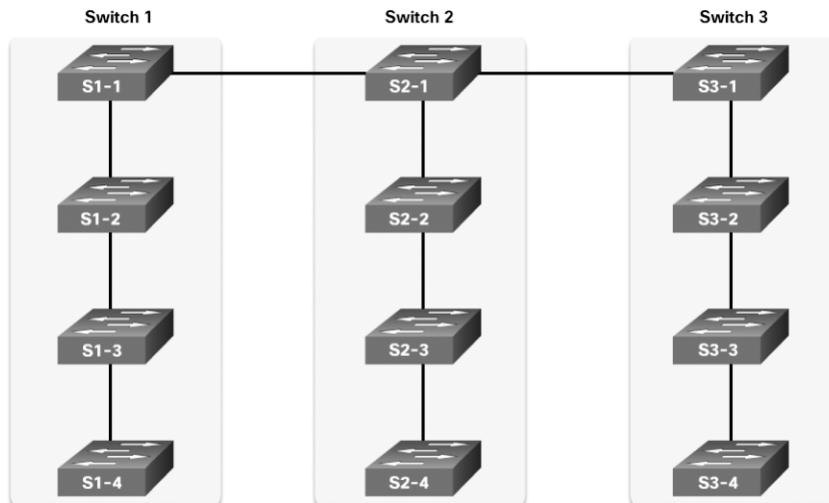


Figure 3-51 Switch Stacking Reduces STP Diameter

Activity 3.3.4.3: Identify Switch Stacking Concepts

Refer to the online course to complete this activity.

Summary (3.4)



Class Activity 3.4.1.1: Documentation Tree

Refer to *Scaling Networks v6 Labs & Study Guide* and the online course to complete this activity.

The employees in your building are having difficulty accessing a web server on the network. You look for the network documentation that the previous network engineer used before he transitioned to a new job; however, you cannot find any network documentation whatsoever.

Therefore, you decide to create your own network record-keeping system. You decide to start at the access layer of your network hierarchy. This is where redundant switches are located, as well as the company servers, printers, and local hosts.

You create a matrix to record your documentation and include access layer switches on the list. You also decide to document switch names, ports in use, cabling connections, root ports, designated ports, and alternate ports.

Problems that can result from a redundant Layer 2 network include broadcast storms, MAC database instability, and duplicate unicast frames. STP is a Layer 2 protocol, which ensures that there is only one logical path between all destinations on the network by intentionally blocking redundant paths that could cause a loop.

STP sends BPDU frames for communication between switches. One switch is elected as the root bridge for each instance of spanning tree. An administrator can control this election by changing the bridge priority. Root bridges can be configured to enable spanning-tree load balancing by a VLAN or by a group of VLANs, depending on the spanning-tree protocol used. STP then assigns a port role to each participating port, using a path cost. The root path cost is equal to the sum of all the port costs along the path to the root bridge. A port cost is automatically assigned to each port; however, it can also be manually configured. Paths with the lowest cost become preferred, and all other redundant paths are blocked.

PVST+ is the default configuration of IEEE 802.1D on Cisco switches. It runs one instance of STP for each VLAN. A newer, faster-converging spanning-tree protocol, RSTP, can be implemented on Cisco switches on a per-VLAN basis in the form of Rapid PVST+. Multiple Spanning Tree (MST) is the Cisco implementation of Multiple Spanning Tree Protocol (MSTP), where one instance of spanning tree runs for a defined group of VLANs. Features such as PortFast and BPDU guard ensure that hosts in the switched environment are provided immediate access to the network without interfering with spanning-tree operation.

Switch stacking allows connection of up to nine Catalyst 3750 switches to be configured and presented to the network as a single entity. STP views the switch stack as a single switch. This additional benefit helps ensure the IEEE recommended maximum diameter of seven switches.

Practice

The following activities provide practice with the topics introduced in this chapter. The Labs and Class Activities are available in the companion *Scaling Networks v6 Labs & Study Guide* (ISBN 9781587134333). The Packet Tracer activity instructions are also in the *Labs & Study Guide*. The PKA files are found in the online course.



Class Activities

Class Activity 3.0.1.2: Stormy Traffic

Class Activity 3.4.1.1: Documentation Tree



Labs

Lab 3.1.2.12: Building a Switched Network with Redundant Links

Lab 3.3.2.3: Configuring Rapid PVST+, PortFast, and BPDU Guard



Packet Tracer Activities

Packet Tracer 3.1.1.5: Examining a Redundant Design

Packet Tracer 3.3.1.5: Configuring PVST+

Packet Tracer 3.3.2.2: Configuring Rapid PVST+

Check Your Understanding Questions

Complete all the review questions listed here to test your understanding of the topics and concepts in this chapter. The appendix “Answers to ‘Check Your Understanding’ Questions” lists the answers.

1. What could be the effect of duplicate unicast frames arriving at a destination device due to multiple active alternative physical paths?
 - A. Application protocols malfunction.
 - B. Frame collisions increase.

- C. The number of broadcast domains increases.
 - D. The number of collision domains increases.
2. What additional information is contained in the 12-bit extended system ID of a BPDU?
- A. IP address
 - B. MAC address
 - C. Port ID
 - D. VLAN ID
3. Which three components are combined to form a bridge ID? (Choose three.)
- A. Bridge priority
 - B. Cost
 - C. Extended system ID
 - D. IP address
 - E. MAC address
 - F. Port ID
4. Which STP port role is adopted by a switch port if there is no other port with a lower cost to the root bridge?
- A. Alternate port
 - B. Designated port
 - C. Disabled port
 - D. Root port
5. Which is the default STP operation mode on Cisco Catalyst switches?
- A. MST
 - B. MSTP
 - C. PVST+
 - D. Rapid PVST+
 - E. RSTP
6. What is an advantage of PVST+?
- A. PVST+ optimizes performance on the network through autoselection of the root bridge.
 - B. PVST+ optimizes performance on the network through load sharing.

- C. PVST+ reduces bandwidth consumption compared to traditional implementations of STP that use CST.
 - D. PVST+ requires fewer CPU cycles for all the switches in the network.
7. In which two port states does a switch learn MAC addresses and process BPDUs in a PVST network? (Choose two.)
- A. Blocking
 - B. Disabled
 - C. Forwarding
 - D. Learning
 - E. Listening
8. Which STP priority configuration would ensure that a switch would always be the root switch?
- A. `spanning-tree vlan 10 priority 0`
 - B. `spanning-tree vlan 10 priority 4096`
 - C. `spanning-tree vlan 10 priority 61440`
 - D. `spanning-tree vlan 10 root primary`
9. To obtain an overview of the spanning-tree status of a switched network, a network engineer issues the **show spanning-tree** command on a switch. Which two items of information does this command display? (Choose two.)
- A. The IP address of the management VLAN interface
 - B. The number of broadcasts received on each root port
 - C. The role of the ports in all VLANs
 - D. The root bridge BID
 - E. The status of native VLAN ports
10. Which two network design features require Spanning Tree Protocol (STP) to ensure correct network operation? (Choose two.)
- A. Implementing VLANs to contain broadcasts
 - B. Link-state dynamic routing that provides redundant routes
 - C. Redundant links between Layer 2 switches
 - D. Removing single points of failure with multiple Layer 2 switches
 - E. Static default routes

11. What value determines the root bridge when all switches connected by trunk links have default STP configurations?
 - A. Bridge priority
 - B. Extended system ID
 - C. MAC address
 - D. VLAN ID

12. Which two concepts relate to a switch port that is intended to have only end devices attached and intended never to be used to connect to another switch? (Choose two.)
 - A. Bridge ID
 - B. Edge port
 - C. Extended system ID
 - D. PortFast
 - E. PVST+

13. Which Cisco switch feature ensures that configured switch edge ports do not cause Layer 2 loops if a port is mistakenly connected to another switch?
 - A. BPDU guard
 - B. Extended system ID
 - C. PortFast
 - D. PVST+

Symbols

? (question mark), 292
 0.0.0.0 static route, 380
 2-WAY/DROTHER state, 539
 802.1D BPDU
 frame format, 128–131
 propagation and process, 131–136
 802.1D-2004, 148
 802.3ad, 180

A

ABRs (area border routers), 425, 500
 accumulated costs (OSPF), 455–456
 Acknowledgement packets, 279, 281–282
 active mode (LACP), 187
 active state (HSRP), 206
 AD (administrative distance), 246, 418–419
 Address Families (AF) feature, 417
 addresses. *See also* IPv4; IPv6
 global unicast, 474
 link-local, 472–473, 475–476
 assigning, 476–477
 IPv6, 344–345, 347–348
 verifying, 477
 multicast, 242
 adjacencies, 433–435
 adjacency database, 419
 DRs (designated routers), 538–540
 neighbor adjacency, 310–311
 adjacency database, 419
 administrative distance (AD), 246, 418–419
 advertisements
 LSAs (link-state advertisements), 501–502
 flooding, 436–438, 531
 type 1 LSAs, 502–503
 type 2 LSAs, 503–504
 type 3 LSAs, 504–505
 type 4 LSAs, 505
 type 5 LSAs, 506
 VTP (VLAN Trunking Protocol), 52
 AF (Address Families) feature, 417
 aggregation. *See* link aggregation
 algorithms
 Bellman-Ford, 245
 Dijkstra's, 248–249, 420
 distance vector, 242–245
 DUAL (Diffusing Update Algorithm), 323–324
 FC (*feasibility condition*), 326–327
 FD (*feasible distance*), 324–325
 feasible successors, 335–337
 FS (*feasible successors*), 326–327
 FSM (*Finite State Machine*), 334
 no feasible successors, 338–340
 RD (*reported distance*), 326–327
 successor distance, 324–325
 STA (Spanning Tree Algorithm), 114–117
 port roles, 117–119
 root bridges, 119–120
 root path cost, 121–124
 RSTP (*Rapid Spanning Tree Protocol*), 124–126
 alternate ports, 127–128
 application-specific integrated circuits (ASIC), 24
 area 0 (backbone area), 424–425
 area border routers (ABRs), 425, 500
 Area ID field (Hello packets), 429
 ASBRs (Autonomous System Boundary Routers),
 501, 547
 ASIC (application-specific integrated circuits), 24
 assigning
 link-local addresses, 476–477
 ports to VLANs, 66
 router IDs
 OSPFv2, 445
 OSPFv3, 479
 authentication (EIGRP), 279
 auto mode (PAGP), 185
 auto-cost reference-bandwidth command, 456
 automatic summarization (EIGRP)
 configuration, 371–372
 disabling, 408

- enabling, 408
- how it works, 369–370
- network topology, 367–369
- routing table, 376–378
- summary route, 378–380
- topology table, 375–376
- troubleshooting, 405–408
- verifying
 - routing table*, 376–378
 - show ip protocols command*, 372–375
 - topology table*, 375–376

autonomous system (AS), 224

Autonomous System Boundary Routers (ASBRs), 501, 547

autonomous system numbers, 291–292

autosummarization. *See* automatic summarization (EIGRP)

auto-summary command, 372, 408

B

- backbone area, 424–425
- backbone routers, 500
- backup designated routers. *See* BDRs (backup designated routers)
- backup ports, 118
- backup routers, 202
- balancing load. *See* load balancing
- in-band management, 30–31
- bandwidth
 - EIGRP (Enhanced Interior Gateway Routing Protocol)
 - bandwidth metric*, 316–318
 - utilization*, 385–386
 - increasing, 13–14
 - interface bandwidth
 - adjusting*, 462–463
 - default interface bandwidths*, 460–462
 - reference bandwidth, 456–460
- bandwidth command, 316–318, 369, 462–463
- BDRs (backup designated routers), 435–438
 - adjacencies, 538–540
 - election process, 540–543
 - roles, 535–538
- Bellman-Ford algorithm, 245
- best path, 221

- BGP (Border Gateway Protocol), 223, 292
- BID (bridge ID), 117, 154–155
- bit bucket, 377
- blocking state (ports), 145
- Border Gateway Protocol (BGP), 223, 292
- bounded updates, 247, 282
- BPDU (bridge protocol data units)
 - 802.1D BPDU
 - frame format*, 128–131
 - propagation and process*, 131–136
 - BPDU Guard, 156–158
 - Version 2 BPDUs, 149–150
- branch routers, 27
- bridge ID (BID), 117, 154–155
- Bridge ID field (BPDU), 129
- bridge protocol data units. *See* BPDUs (bridge protocol data units)
- bridges, root, 119–120
- broadcast multiaccess networks, 529
- broadcast storms, 111–112
- building switch block, 13

C

- campus wired LAN (local area network) design, 4
 - Cisco validated designs
 - hierarchical design model*, 6–8
 - need for network scaling*, 4–6
 - network expansion
 - access layer*, 14–15
 - bandwidth*, 13–14
 - design for scalability*, 8–10
 - failure domains*, 11–13
 - redundancy planning*, 10–11
 - routing protocols, fine-tuning*, 15–17
 - switches, 17
- Can Submarines Swim? class activity, 416
- Catalyst switches. *See* switches
- CEF (Cisco Express Forwarding), 91, 390
- channel-group command, 190
- chassis aggregation, 169–172
- child routes, 232
- CIDR (classless interdomain routing), 228
- Cisco Express Forwarding (CEF), 91, 390
- Cisco IOS files and licensing, 30
- Cisco validated designs

- hierarchical design model, 6–8
- need for network scaling, 4–6
- class activities**
 - Can Submarines Swim?416
 - Classless EIGRP, 274
 - Digital Trolleys, 522
 - Documentation Tree, 173
 - DR and BDR Election, 528
 - How Much Does This Cost?, 221–222
 - Imagine This, 180
 - Layered Network Design Simulation, 43
 - Leaving on a Jet Plane, 494
 - Linking Up, 214
 - Network by Design, 3
 - Portfolio RIP and EIGRP, 358
 - SPF Troubleshooting Mastery, 585
 - Stepping Through OSPFv3, 486
 - Stormy Traffic, 107
 - Tuning EIGRP, 410
- classful routing protocols, 228–231**
- classification of routing protocols, 222–224**
- Classless EIGRP class activity, 274**
- classless interdomain routing (CIDR), 228**
- classless routing protocols, 231–233**
- clear ip ospf command, 566**
- clear ip ospf process command, 447, 545**
- clear ipv6 ospf command, 580**
- clear ipv6 ospf process command, 480**
- clearing OSPF process, 447, 480**
- clients (FTP)**
 - configuration, 60
 - verification, 62–63
- cloud-managed switches, 18**
- cold start, 237–238**
- collapsed core design, 7**
- Coltun, Rob, 417**
- Common Spanning Tree. *See* STP (Spanning Tree Protocol)**
- composite metric (EIGRP), 313–315**
- configuration**
 - DTP (Dynamic Trunking Protocol)
 - initial configuration, 71*
 - negotiated interface modes, 72–73*
 - verification, 72–73*
 - EIGRP for IPv4
 - automatic summarization, 371–372*
 - autonomous system numbers, 291–292*
 - compared to EIGRP for IPv6, 342–343*
 - network command, 296–300*
 - network topology, 289–290*
 - passive interfaces, 300–302*
 - router eigrp command, 292–293*
 - router IDs, 293–296*
 - verification, 296*
 - EIGRP for IPv6, 341
 - automatic summarization, 408*
 - compared to EIGRP for IPv4, 342–343*
 - ipv6 eigrp command, 350–352*
 - link-local addresses, 344–345, 347–348*
 - network topology, 345–347*
 - routing process, 349–350*
 - EtherChannel, 183–184
 - guidelines, 188–189*
 - interfaces, 189–191*
 - global unicast addresses, 474
 - HSRP (Hot Standby Router Protocol), 206–208
 - intervals
 - OSPFv2, 555–557*
 - OSPFv3, 557–559*
 - multiarea OSPF
 - multiarea OSPFv2, 511–513*
 - multiarea OSPFv3, 513–514*
 - multiarea OSPF (Open Shortest Path First), 510–511
 - OSPF in multiaccess networks
 - challenges, 531–533*
 - DR/BDR adjacencies, 538–540*
 - DR/BDR election process, 540–543*
 - DR/BDR roles, 535–538*
 - network types, 528–531*
 - OSPF DRs, 533–534*
 - OSPF priority, 544–546*
 - PVST+
 - BPDU Guard, 156–158*
 - bridge IDs, 154–155*
 - Catalyst 2960 default configuration, 153–154*
 - PortFast, 156–158*
 - Rapid PVST+161–163
 - routers
 - enabling, 32*
 - OSPFv2, 445–447*
 - OSPFv3, 479–480*

- running configuration*, 33–34
- troubleshooting*, 82–83
- single-area OSPF
 - OSPFv2*, 448–453
 - OSPFv3*, 473–481
 - passive interfaces*, 450–453
 - reference bandwidth*, 456–460
 - router IDs*, 441–447, 477–480
- switches
 - enabling*, 39
 - running configuration*, 39–40
 - verification*, 79–81
- VLANs (virtual local area networks)
 - assigning ports to*, 66
 - Layer 3 switching*, 95–96
 - verification*, 67–69
 - VLAN creation*, 65–66
- VTP (VLAN Trunking Protocol), 57
 - cautions*, 55–56
 - clients*, 60
 - default configuration*, 53–55
 - domain name and password*, 59
 - verification*, 62–63
 - VLANs*, 60–61
 - VTP server*, 58–59
- convergence, 241–242, 247, 312–313
- copy running-config startup-config command, 65
- core layer, 7
- costs
 - load balancing
 - equal-cost load balancing*, 388
 - unequal-cost load balancing*, 391
 - path cost, 234
 - root path cost, 121–124
 - single-area OSPF (Open Shortest Path First), 453–464
 - accumulated costs*, 455–456
 - calculating*, 454–455
 - default interface bandwidths*, 460–462
 - interface bandwidth*, 462–463
 - reference bandwidth*, 456–460
 - setting manually*, 463–464
- CST (Common Spanning Tree). *See* STP (Spanning Tree Protocol)

D

- data center switches, 18
- data structures (OSPF), 419–420, 583
- Database Description (DBD) packet, 428
- databases
 - LSDB (link-state database), 259–260
 - OSPF (Open Shortest Path First), 419–420
 - vlan.dat, 50
- DBD (Database Description) packet, 428
- dead interval field (Hello packets), 430
- dead intervals
 - modifying
 - OSPFv2*, 555–557
 - OSPFv3*, 557–559
 - verifying, 554–555
- debug eigrp fsm command, 339
- debug standby command, 210–211
- debugging HSRP (Hot Standby Router Protocol), 23, 210–213. *See also* troubleshooting
- DEC (Digital Equipment Corporation), 267
- default configuration (VTP), 53–55
- default election process (DR/BDR), 540–543
- default gateway limitations, 198
- default interface bandwidths, 460–462
- default route propagation
 - EIGRP (Enhanced Interior Gateway Routing Protocol)
 - IPv4*, 380–382
 - IPv6*, 383–384
 - verification*, 382–383
 - OSPF (Open Shortest Path First), 547
 - OSPFv2*, 547–548
 - OSPFv3*, 551–552
 - propagated IPv4 route verification*, 549–550
 - propagated IPv6 route verification*, 552–554
 - static route propagation
 - IPv4*, 380–382
 - IPv6*, 383–384
 - verification*, 382–383
- default-information originate command, 548
- delay (DLY) metric, 319–320
- delay values, 319
- deleting VLANs, 75–77
- departmental switch block, 13

- design. *See* LAN (local area network) design
- designated ports, 118, 127–128
- designated routers. *See* DRs (designated routers)
- desirable mode (PAGP), 185
- destination IPv6 addresses, 344, 472
- device management. *See* network device management
- device selection. *See* network device selection
- Diffusing Update Algorithm. *See* DUAL (Diffusing Update Algorithm)
- Digital Equipment Corporation (DEC), 267
- Digital Trolleys class activity, 522
- Dijkstra, Edsger Wybe, 416
- Dijkstra's algorithm, 248–249, 420
- directly connected networks, detecting, 237
- disabled ports, 118, 145
- disabling
 - EIGRP automatic summarization, 408
 - passive interfaces, 572
- discontiguous networks, 228
- discovery
 - initial route discovery
 - convergence*, 312–313
 - neighbor adjacency*, 310–311
 - topology table*, 311–312
 - network discovery, 238–239
- distance, 226
- distance vector algorithms, 242–245
- distance vector dynamic routing, 226. *See also* EIGRP (Enhanced Interior Gateway Routing Protocol)
- convergence, 241–242
 - distance vector algorithms, 242–245
 - network discovery, 238–239
 - operation, 236–238
 - protocols, 16
 - routing information exchange, 239–241
 - technologies, 242
- distribution layer, 7
- DLY (delay) metric, 319–320
- Documentation Tree class activity, 173
- domains
 - failure domains, 11–13
 - VTP (VLAN Trunking Protocol), 59
- Down state (OSPF), 432, 561
- DR and BDR Election class activity, 528
- DROTHERs, 437, 533
- DRs (designated routers), 435–438, 533–534
 - adjacencies, 538–540
 - BDRs (backup designated routers)
 - adjacencies*, 538–540
 - election process*, 540–543
 - roles*, 535–538
 - election process, 540–543
 - roles, 535–538
- DTP (Dynamic Trunking Protocol), 48
 - initial configuration, 71
 - negotiated interface modes, 72–73
 - troubleshooting, 89
 - verification, 72–73
- DUAL (Diffusing Update Algorithm), 245, 323–324
 - FC (feasibility condition), 326–327
 - FD (feasible distance), 324–325
 - feasible successors, 335–337
 - FS (feasible successors), 326–327
 - FSM (Finite State Machine), 334
 - no feasible successors, 338–340
 - RD (reported distance), 326–327
 - successor distance, 324–325
- duplicate unicast frames, 113
- dynamic routing, 221. *See also* EIGRP (Enhanced Interior Gateway Routing Protocol); OSPF (Open Shortest Path First)
- BGP (Border Gateway Protocol), 223
 - classful routing protocols, 228–231
 - classless routing protocols, 231–233
 - distance vector dynamic routing, 226
 - convergence*, 241–242
 - distance vector algorithms*, 242–245
 - network discovery*, 238–239
 - operation*, 236–238
 - routing information exchange*, 239–241
 - technologies*, 242
 - EGP (Exterior Gateway Protocols), 224–225
 - IGP (Interior Gateway Protocols), 224–225
 - IGRP (Interior Gateway Routing Protocol), 246–247
 - IS-IS (Intermediate System-to-Intermediate System), 267
 - link-state dynamic routing
 - advantages of*, 264–265
 - Dijkstra's algorithm*, 248–249
 - disadvantages of*, 265–266

- protocols*, 226–228
- SPF (Shortest Path First)*, 248, 249–251
- link-state updates
 - flooding LSPs*, 258–259
 - Hello packets*, 256–257
 - link-state routing process*, 251–253
 - LSDB (link-state database)*, 259–260
 - LSP (link-state packets)*, 257
 - OSPF routes*, 264
 - SPF (Shortest Path First) tree*, 260–263
- protocol classification, 222–224
- RIP (Routing Information Protocol), 245–246
- routing protocol characteristics, 233–234
- routing protocol metrics, 234–236
- Dynamic Trunking Protocol. See DTP (Dynamic Trunking Protocol)**

E

- edge ports, 150–151
- edge routers, 547
- EGP (Exterior Gateway Protocols), 224–225
- EIGRP (Enhanced Interior Gateway Routing Protocol), 15, 246–247, 274
 - authentication, 279
 - automatic summarization
 - configuration*, 371–372
 - disabling*, 408
 - enabling*, 408
 - how it works*, 369–370
 - network topology*, 367–369
 - routing table*, 376–378
 - summary route*, 378–380
 - topology table*, 375–376
 - verification*, 372–378
 - bandwidth utilization, 385–386
 - characteristics of, 223
 - configuration for IPv4
 - autonomous system numbers*, 291–292
 - compared to EIGRP for IPv6*, 342–343
 - network command*, 296–300
 - network topology*, 289–290
 - passive interfaces*, 300–302
 - router eigrp command*, 292–293
 - router IDs*, 293–296
 - verifying EIGRP process*, 296
 - configuration for IPv6, 341
 - compared to EIGRP for IPv4*, 342–343
 - ipv6 eigrp command*, 350–352
 - link-local addresses*, 344–345, 347–348
 - network topology*, 345–347
 - routing process*, 349–350
 - default route propagation
 - IPv4*, 380–382
 - IPv6*, 383–384
 - verification*, 382–383
 - DUAL (Diffusing Update Algorithm), 323–324
 - FC (feasibility condition)*, 326–327
 - FD (feasible distance)*, 324–325
 - feasible successors*, 335–337
 - FS (feasible successors)*, 326–327
 - FSM (Finite State Machine)*, 334
 - no feasible successors*, 338–340
 - RD (reported distance)*, 326–327
 - successor distance*, 324–325
 - features of, 274–276
 - Hello and Hold timers, 367–386
 - initial route discovery
 - convergence*, 312–313
 - neighbor adjacency*, 310–311
 - topology table*, 311–312
 - load balancing
 - IPv4*, 388–390
 - IPv6*, 390–392
 - metrics
 - bandwidth metric*, 316–318
 - calculating*, 320–323
 - composite metric*, 313–315
 - delay metric*, 319–320
 - interface metric values*, 315–316
 - named EIGRP, 275
 - packets
 - Acknowledgement packets*, 281–282
 - encapsulating*, 284–285
 - Hello packets*, 280–281
 - packet headers and TLV*, 285–288
 - Query packets*, 283–284
 - Reply packets*, 283–284
 - table of*, 279–280
 - Update packets*, 281–282
 - PDMs (protocol dependence modules), 276–277
 - RTP (Reliable Transport Protocol), 278

topology table
 no feasible successor, 332–334
 show ip eigrp topology command, 328–332

troubleshooting

- automatic summarization*, 405–408
- basic commands*, 392–394
- components*, 394–395
- EIGRP parameters*, 398–399
- interfaces*, 399–401
- Layer 3 connectivity*, 397–398
- missing network statement*, 403–405
- neighbor issues*, 397–401
- passive interfaces*, 401–403
- routing table issues*, 401–408

tuning, 366

- automatic summarization*, 366–380
- bandwidth utilization*, 385–386
- default route propagation*, 380–384
- Hello and Hold timers*, 367–386
- IPv4 load balancing*, 388–390
- IPv6 load balancing*, 390–392

verification with IPv4

- neighbors*, 302–304
- routing table*, 306–309
- show ip protocols command*, 304–306

verification with IPv6

- neighbor table*, 352–354
- routing table*, 355–356
- show ipv6 protocols command*, 354–355

eigrp log-neighbor-changes command, 298

eigrp router-id command, 295, 349

election process (DR/BDR), 540–543

EM (Extended Maintenance), 30

enabling. *See also* configuration

- routers, 32
- switches, 39

encapsulating messages, 284–285, 426–427

Enhanced Interior Gateway Routing Protocol.
See EIGRP (Enhanced Interior Gateway Routing Protocol)

enterprise networks, 4–6

entrance routers, 547

equal-cost load balancing, 388

EtherChannel, 13–14, 180–181

- advantages of, 182–183
- configuration

- guidelines*, 188–189
- interfaces*, 189–191
- implementation restrictions, 183–184
- LACP (Link Aggregation Control Protocol), 186–187
- PAgP (Port Aggregation Protocol), 185–186
- troubleshooting, 194–197
- verifying, 191–194

Ethernet, PoE (Power over Ethernet), 23–24

Exchange state (OSPF), 433, 561

exchanging routing information, 239–241

expanding networks

- access layer, 14–15
- bandwidth, 13–14
- design for scalability, 8–10
- failure domains, 11–13
- redundancy planning, 10–11
- routing protocols, fine-tuning, 15–17

expected versus actual topology, 164–165

ExStart state (OSPF), 433, 561

Extended Maintenance (EM), 30

extended system ID, 136–140, 145–147

extended VLANs, 63

- definition of, 64
- VLAN ranges on Catalyst switches, 63–64

Exterior Gateway Protocols (EGP), 224–225

external LSA entries, 506

F

failover, routers, 200–201

failure

- failure domains, 11–13
- HSRP (Hot Standby Router Protocol), 209–210
- STP (Spanning Tree Protocol), 166–168

feasibility condition (FC), 326–327

feasible distance (FD), 324–325, 376

feasible successors (FS), 326–327, 376

Ferguson, Dennis, 417

FHRPs (First Hop Redundancy Protocols)

- default gateway limitations, 198
- GLBP (Gateway Load Balancing Protocol), 202
- HSRP (Hot Standby Router Protocol)
 - configuration*, 206–208
 - definition of*, 202
 - operation of*, 203–204
 - preemption*, 205

priority, 204–205
states and timers, 205–206
troubleshooting, 209–213
verification, 208–209
versions, 204
 IRDP (ICMP Router Discovery Protocol), 202
 router failover, 200–201
 router redundancy, 199–200
 VRRP (Virtual Router Redundancy Protocol), 202
fields
 BPDU (bridge protocol data units), 129–130
 Hello packets, 429–430
files (IOS), 30
fine-tuning routing protocols, 15–17
Finite State Machine (FSM), 324, 334
fixed configuration switches, 19–20
Flags field (BPDU), 129
flooding
 LSAs (link-state advertisements), 436–438, 531
 LSP (link-state packets), 258–259
form factors (router), 28–29
forwarding database, 420
forwarding rates, 22
forwarding state (ports), 145
frame format, 802.1D BPDU, 128–131
FS (feasible successors), 326–327, 335–340, 376
FSM (Finite State Machine), 324, 334
Full state (OSPF), 433, 562
FULL/BDR state, 539
FULL/DR state, 538
FULL/DROTHER state, 539

G

Garcia-Luna-Aceves, J. J. 245
GATED, 417
Gateway Load Balancing Protocol (GLBP), 202
gateways, 198, 547
GLBP (Gateway Load Balancing Protocol), 202
global unicast addresses, 474

H

headers, 285–288
 EIGRP (Enhanced Interior Gateway Routing Protocol), 285–288
 OSPF (Open Shortest Path First), 427

hello intervals
 EIGRP (Enhanced Interior Gateway Routing Protocol), 367–386
 OSPF (Open Shortest Path First), 430
 OSPFv2, 555–557
 OSPFv3, 557–559
 verifying, 554–555
Hello keepalive mechanism, 247
Hello packets
 EIGRP (Enhanced Interior Gateway Routing Protocol), 279, 280–281
 OSPF (Open Shortest Path First), 256–257, 428–430
Hello Time field (BPDU), 130
hierarchical design model, 6–8
hold times, tuning, 367–386
Hot Standby Router Protocol. *See* HSRP (Hot Standby Router Protocol)
How Much Does This Cost? class activity, 221–222
HSRP (Hot Standby Router Protocol)
 configuration, 206–208
 definition of, 202
 operation of, 203–204
 preemption, 205
 priority, 204–205
 states and timers, 205–206
 troubleshooting
 common configuration issues, 213
 debug commands, 210–213
 failure, 209–210
 verification, 208–209
 versions, 204
hybrid routing protocols, 276

I

IANA (Internet Assigned Numbers Authority), 291
ICMP Router Discovery Protocol (IRDP), 202
IDs. *See* router IDs
IEEE 802.1D-2004, 148
IEEE 802.1w. *See* RSTP (Rapid Spanning Tree Protocol)
IEEE 802.3ad, 180
IETF (Internet Engineering Task Force), 417
IGP (Interior Gateway Protocols), 224–225
IGRP (Interior Gateway Routing Protocol), 223, 246–247
Imagine This class activity, 180

- implementation. *See* configuration
- Init state (OSPF), 432, 561
- initial route discovery
 - convergence, 312–313
 - neighbor adjacency, 310–311
 - topology table, 311–312
- initial state (HSRP), 206
- interarea routing, 425
- interface bandwidth
 - adjusting, 462–463
 - default interface bandwidths, 460–462
- interface metric values, 315–316
- interface port-channel command, 190–191
- interface range command, 66, 189–190
- interface table (OSPF), 583
- interfaces
 - EIGRP (Enhanced Interior Gateway Routing Protocol)
 - bandwidth utilization*, 385–386
 - Hello and Hold timers*, 367–386
 - IPv4 load balancing*, 388–390
 - IPv6 load balancing*, 390–392
 - troubleshooting*, 399–401
 - EtherChannel, 189–191
 - IPv6-enabled interfaces, verifying, 475
 - loopback interfaces, 447
 - Null0, 377–378
 - OSPF (Open Shortest Path First)
 - interface table*, 583
 - OSPFv2*, 448
 - OSPFv3*, 481
 - verification*, 468–469
 - passive interfaces, 450–453
 - configuration*, 300–302
 - disabling*, 572
 - verification*, 302
 - port channel interfaces, 182
 - SVIs (switch virtual interfaces)
 - definition of*, 91
 - inter-VLAN routing with*, 91–94
 - switch virtual interfaces, inter-VLAN routing with, 91–94
 - troubleshooting, 81
- Interior Gateway Protocols (IGP), 224–225
- Interior Gateway Routing Protocol (IGRP), 223, 246–247
- Intermediate System-to-Intermediate System (IS-IS), 223, 267, 417
- internal routers, 499
- International Organization for Standardization (ISO), 267, 417
- Internet Assigned Numbers Authority (IANA), 291
- Internet Engineering Task Force (IETF), 417
- Internet Protocol Security (IPsec), 418
- intervals
 - dead intervals, 554–555
 - hello intervals, 554–555
 - modifying
 - OSPFv2*, 555–557
 - OSPFv3*, 557–559
- inter-VLAN configuration
 - deleting VLANs, 75–77
 - interface issues, 81
 - inter-VLAN routing
 - with routed ports*, 94–95
 - with switch virtual interfaces*, 91–94
 - legacy inter-VLAN routing solution, 77
 - router-on-a-stick inter-VLAN routing, 78
 - routing configuration, 82–83
 - switch configuration, 79–81
 - switch port issues, 77–79
- IOS (Internetwork Operating System) files and licensing, 30
- ip address command, 84
- ip bandwidth-percent eigrp command, 385
- ip hello-interval eigrp command, 387
- ip hold-time eigrp command, 367–387
- ip mtu command, 573
- ip ospf cost command, 463–464
- ip ospf dead-interval command, 555–556
- ip ospf hello-interval command, 555–556
- ip ospf priority command, 544
- IPsec (Internet Protocol Security), 418
- IPv4
 - default routes
 - propagating in OSPFv2*, 547–548
 - verifying in OSPFv2*, 549–550
 - EIGRP (Enhanced Interior Gateway Routing Protocol)
 - autonomous system numbers*, 291–292
 - bandwidth utilization*, 385–386
 - compared to EIGRP for IPv6*, 342–343

- default route propagation*, 380–382
- load balancing*, 388–390
- neighbors*, 302–304
- network command*, 296–300
- network topology*, 289–290
- passive interfaces*, 300–302
- router eigrp command*, 292–293
- router IDs*, 293–296
- routing table*, 306–309
- show ip protocols command*, 304–306
- verifying EIGRP process*, 296

IP addressing issues

- errors with IP addresses and subnet masks*, 83–85
- verifying configuration*, 85–87

routing table, examining, 306–309

troubleshooting, 83–87

IPv6

- default routes
 - propagating in OSPFv3*, 551–552
 - verifying in OSPFv2*, 552–554
- EIGRP (Enhanced Interior Gateway Routing Protocol), 341
 - bandwidth utilization*, 385–386
 - compared to EIGRP for IPv4*, 342–343
 - default route propagation*, 383–384
 - ipv6 eigrp command*, 350–352
 - link-local addresses*, 344–345, 347–348
 - load balancing*, 390–392
 - neighbor table*, 352–354
 - network topology*, 345–347
 - routing process*, 349–350
 - routing table*, 355–356
 - show ipv6 protocols command*, 354–355
- GLBP (Gateway Load Balancing Protocol), 202
- HSRP (Hot Standby Router Protocol) for IPv6, 202
- IPv6-enabled interfaces, verifying, 475
- link-local addresses, 472–473
- VRRP (Virtual Router Redundancy Protocol), 202

ipv6 address command, 347

ipv6 bandwidth-percent eigrp command, 386

ipv6 eigrp command, 350–352, 405

ipv6 hello-interval eigrp command, 367–387

ipv6 hold-time eigrp command, 367–387

ipv6 ospf command, 479–480

ipv6 ospf dead-interval command, 557–558

ipv6 ospf hello-interval command, 557–558

ipv6 ospf priority command, 544

ipv6 router ospf command, 479, 480

ipv6 unicast-routing command, 349

IRDP (ICMP Router Discovery Protocol), 202

IS-IS (Intermediate System-to-Intermediate System), 223, 267, 417

ISO (International Organization for Standardization), 267, 417

J-K-L

k values, verifying, 313–315

LACP (Link Aggregation Control Protocol), 186–187

LAN (local area network) design, 3

- campus wired LANs, 4
 - Cisco validated designs*, 4–8
 - network expansion*, 8–17
- network device management, 29
 - in-band versus out-of-band management*, 30–31
 - basic router CLI commands*, 30–31
 - basic switch CLI commands*, 38–40
 - IOS files and licensing*, 30
 - router show commands*, 34–38
 - switch show commands*, 40–42
- network device selection, 17
 - router hardware*, 26–29
 - switch hardware*, 17–25
- network expansion
 - access layer*, 14–15
 - bandwidth*, 13–14
 - design for scalability*, 8–10
 - failure domains*, 11–13
 - redundancy planning*, 10–11
 - routing protocols, fine-tuning*, 15–17

Layer 3 switching, 89–91, 397–398

- configuration, 95–96
- inter-VLAN routing with routed ports, 94–95
- inter-VLAN routing with switch virtual interfaces, 91–94
- troubleshooting, 95–98
- verification, 570

Layered Network Design Simulation class activity, 43

- learn state (HSRP), 206
- learning state (ports), 145
- Leaving on a Jet Plane class activity, 494
- legacy inter-VLAN routing solutions, 77
- licensing, 30
- link aggregation, 9
 - definition of, 181
 - EtherChannel, 180–181
 - advantages of*, 182–183
 - configuration*, 188–191
 - implementation restrictions*, 183–184
 - LACP (Link Aggregation Control Protocol), 186–187
 - PAgP (Port Aggregation Protocol), 185–186
 - troubleshooting*, 194–197
 - verification*, 191–194
- FHRPs (First Hop Redundancy Protocols)
 - default gateway limitations*, 198
 - GLBP (Gateway Load Balancing Protocol), 202
 - HSRP (Hot Standby Router Protocol). *See* HSRP (Hot Standby Router Protocol)
 - IRDP (ICMP Router Discovery Protocol), 202
 - router failover*, 200–201
 - router redundancy*, 199–200
 - VRRP (Virtual Router Redundancy Protocol), 202
- HSRP (Hot Standby Router Protocol)
 - configuration*, 206–208
 - operation of*, 203–204
 - preemption*, 205
 - priority*, 204–205
 - troubleshooting*, 209–213
 - verification*, 208–209
 - versions*, 204
- LACP (Link Aggregation Control Protocol), 186–187
- link types, 152–153
- Linking Up class activity, 214
- link-local addresses, 344–345, 347–348, 472–473, 475–476
 - assigning, 476–477
 - verifying, 477
- Link-State Acknowledgement (LSAck) packet, 428
- link-state advertisements. *See* LSAs (link-state advertisements)
- link-state database (LSDB), 259–260, 419, 583
- link-state packets (LSP), 257, 258–259, 428
- Link-State Request (LSR) packet, 428
- link-state routing, 420–424, 430–431. *See also* OSPF (Open Shortest Path First)
 - advantages of, 264–265
 - Dijkstra's algorithm, 248–249
 - disadvantages of, 265–266
 - Hello packets, 256–257
 - link-state updates
 - flooding LSPs*, 258–259
 - Hello packets*, 256–257
 - link-state routing process*, 251–253
 - LSDB (link-state database)*, 259–260
 - LSP (link-state packets)*, 257
 - OSPF routes*, 264
 - SPF (Shortest Path First) tree*, 260–263
 - LSDB (link-state database), 259–260, 419, 583
 - LSP (link-state packets), 257, 258–259, 428
 - OSPF routes, 264
 - routing process, 251–253
 - routing protocols, 15, 226–228
 - SPF (Shortest Path First), 248, 249–251
 - SPF (Shortest Path First) tree, 260–263
- Link-State Update (LSU) packet, 428
- listen state (HSRP), 206
- listening state (ports), 145
- load balancing, 14, 183
 - equal-cost load balancing, 388
 - IPv4, 388–390
 - IPv6, 390–392
 - PVST+158–160
 - unequal-cost load balancing, 391
- Loading state (OSPF), 433, 562
- loopback interfaces, 447
- LSAck (Link-State Acknowledgement) packet, 428
- LSAs (link-state advertisements), 501–502
 - flooding*, 436–438, 531
 - type 1 LSAs, 502–503
 - type 2 LSAs, 503–504
 - type 3 LSAs, 504–505
 - type 4 LSAs, 505
 - type 5 LSAs, 506
- LSDB (link-state database), 259–260, 419, 583
- LSP (link-state packets), 257, 258–259, 428, 428
- LSR (Link-State Request) packet, 428
- LSU (Link-State Update) packet, 428

M

MAC database instability, 109–111

managing network devices, 29

in-band versus out-of-band management, 30–31

basic router CLI commands, 30–31

basic switch CLI commands, 38–40

IOS files and licensing, 30

router show commands, 34–38

switch show commands, 40–42

master routers, 202

Max Age field (BPDU), 130

maximum-paths command, 390, 411

MD5 (Message Digest 5), 418

Message Age field (BPDU), 130

Message Type field (BPDU), 129

messages

EIGRP (Enhanced Interior Gateway Routing Protocol)

Acknowledgement packets, 281–282

encapsulating, 284–285

Hello packets, 280–281

packet headers and TLV, 285–288

Query packets, 283–284

Reply packets, 283–284

types of, 279–280

Update packets, 281–282

OSPF (Open Shortest Path First), 426–431

DBD (Database Description) packet, 428

encapsulating, 426–427

Hello intervals, 430

Hello packets, 428–430

link-state updates, 430–431

LSAck (Link-State Acknowledgement) packet, 428

LSR (Link-State Request) packet, 428

LSU (Link-State Update) packet, 428

metric weights command, 314

metrics

EIGRP (Enhanced Interior Gateway Routing Protocol)

bandwidth metric, 316–318

calculating, 320–323

composite metric, 313–315

delay metric, 319–320

interface metric values, 315–316

routing protocols, 234–236

missing network statement, troubleshooting, 403–405

mission-critical services, 4

modes (VTP), 50–51

modifying. *See* configuration

modular configuration switches, 19–20

Moy, John, 417

MSTP (Multiple Spanning Tree Protocol)

characteristics of, 142

definition of, 141

MTU size, 573

multiaccess networks, OSPF (Open Shortest Path First) in

challenges, 531–533

DRs (designated routers), 533–534

adjacencies, 538–540

election process, 540–543

roles, 535–538

network types, 528–531

OSPF priority, 544–546

multiarea OSPF (Open Shortest Path First), 15–16, 494. *See also* single-area OSPF (Open Shortest Path First)

advantages of, 424–426, 495–497

configuration

multiarea OSPFv2, 511–513

multiarea OSPFv3, 513–514

implementation, 510–511

LSAs (link-state advertisements), 501–502

type 1 LSAs, 502–503

type 2 LSAs, 503–504

type 3 LSAs, 504–505

type 4 LSAs, 505

type 5 LSAs, 506

messages, 426–431

DBD (Database Description) packet, 428

encapsulating, 426–427

Hello intervals, 430

Hello packets, 428–430

link-state updates, 430–431

LSAck (Link-State Acknowledgement) packet, 428

LSR (Link-State Request) packet, 428

LSU (Link-State Update) packet, 428

route calculation, 508–509

routers, 499–501

routing table entries, 506–508

troubleshooting
data structures, 583
overview, 582

two-layer area hierarchy, 498–499

verification
multiarea OSPFv2, 515–518
multiarea OSPFv3, 518–521

multicast addresses, 242

multihomed, 225

multilayer switching, 24–25

Multiple Spanning Tree Protocol. *See* MSTP (Multiple Spanning Tree Protocol)

multi-VLAN issues
 DTP (Dynamic Trunking Protocol) issues, 89
 inter-VLAN configuration
deleting VLANs, 75–77
interface issues, 81
routing configuration, 82–83
switch configuration, 79–81
switch port issues, 77–79
 IP addressing issues
errors with IP addresses and subnet masks, 83–85
verifying configuration, 85–87
 VTP (VLAN Trunking Protocol) issues, 88

N

named EIGRP (Enhanced Interior Gateway Routing Protocol), 275

naming VLANs, 65

NBMA (nonbroadcast multiaccess), 281, 529

negotiated interface modes (DTP), 72–73

neighbor tables, 352–354, 583

neighbors
 EIGRP (Enhanced Interior Gateway Routing Protocol)
adjacencies, 310–311
examining, 302–304
troubleshooting, 397–401

OSPF (Open Shortest Path First)
adjacency, 433–435
list of, 430
neighbor table, 583
OSPFv2 troubleshooting, 569–573
OSPFv2 verification, 465–466
OSPFv3 verification, 482–483
troubleshooting flowcharts, 566

Network by Design class activity, 3

network command, 296–300, 449–450, 574

network device management, 29
 in-band versus out-of-band management, 30–31
 IOS files and licensing, 30
 routers
basic router CLI commands, 30–31
router show commands, 34–38
 switches
basic switch CLI commands, 38–40
switch show commands, 40–42

network device selection, 17
 router hardware, 26
Cisco routers, 27–28
form factors, 28–29
router requirements, 26
 switch hardware, 17
forwarding rates, 22
multilayer switching, 24–25
PoE (Power over Ethernet), 23–24
port density, 21–22
switch platforms, 17–21

network discovery, 238–239

network edge routers, 28

network expansion
 access layer, 14–15
 bandwidth, 13–14
 design for scalability, 8–10
 failure domains, 11–13
 redundancy planning, 10–11
 routing protocols, fine-tuning, 15–17

network link entries, 503–504

network mask field (Hello packets), 429

network operations center (NOC), 5

network redundancy. *See* redundancy

network topology
 EIGRP (Enhanced Interior Gateway Routing Protocol), 367–369
EIGRP for IPv4, 289–290
EIGRP for IPv6, 345–347
 OSPF (Open Shortest Path First)
OSPFv2, 441–442
OSPFv3, 473–475

no auto-summary command, 380, 408

- no bandwidth command, 317, 462
- no ip ospf dead-interval command, 556
- no ip ospf hello-interval command, 556
- no ipv6 ospf dead-interval command, 557
- no ipv6 ospf hello-interval command, 557
- no passive-interface command, 301, 452, 572
- no router eigrp command, 293
- no spanning-tree cost command, 122
- no switchport command, 91, 95
- no vlan command, 76
- NOC (network operations center), 5
- nonbroadcast multiaccess (NBMA), 281, 529
- noncontiguous networks, 228
- Nonstop Forwarding (NSF), 399
- normal-range VLANs, 64
- NSF (Nonstop Forwarding), 399
- Null0 interface, 377–378
- numbers, autonomous system numbers, 291–292

O

- on mode (EtherChannel), 185
- Open Shortest Path First. *See* OSPF (Open Shortest Path First)
- optimization. *See* tuning
- order of precedence (router IDs), 443
- OSI layers, redundancy at, 108–109
- OSPF (Open Shortest Path First), 15, 416. *See also* multiarea OSPF (Open Shortest Path First); single-area OSPF (Open Shortest Path First)
 - characteristics of, 223
 - components of, 419–420
 - default routes, 547
 - propagating in OSPFv2*, 547–548
 - propagating in OSPFv3*, 551–552
 - verifying in OSPFv2*, 549–550
 - evolution of, 417–418
 - features of, 418
 - interface bandwidth
 - adjusting*, 462–463
 - default interface bandwidths*, 460–462
 - intervals
 - dead intervals*, 554–555
 - modifying in OSPFv2*, 555–557
 - modifying in OSPFv3*, 557–559
 - link-local addresses, 472–473
 - link-state operation, 420–424
 - messages, 426–431
 - DBD (Database Description) packet*, 428
 - encapsulating*, 426–427
 - Hello intervals*, 430
 - Hello packets*, 428–430
 - link-state updates*, 430–431
 - LSAck (Link-State Acknowledgement) packet*, 428
 - LSR (Link-State Request) packet*, 428
 - LSU (Link-State Update) packet*, 428
 - in multiaccess networks
 - challenges*, 531–533
 - DR/BDR adjacencies*, 538–540
 - DR/BDR election process*, 540–543
 - DR/BDR roles*, 535–538
 - network types*, 528–531
 - OSPF DRs*, 533–534
 - OSPF priority*, 544–546
 - multilayer switching, 42
 - network topology, 441–442
 - operation
 - BDRs (backup designated routers)*, 435–438
 - database synchronization*, 438–440
 - DRs (designated routers)*, 435–438
 - neighbor adjacencies*, 433–435
 - states*, 432–433
 - reference bandwidth, 456–460
 - routes, adding to routing table, 264
 - states, 560–562
 - troubleshooting
 - data structures*, 583
 - flowcharts*, 566–568
 - OSPFv2 neighbor issues*, 569–573
 - OSPFv2 routing table issues*, 573–575
 - OSPFv2 troubleshooting commands*, 562–566
 - OSPFv3 routing tables*, 580–582
 - OSPFv3 troubleshooting commands*, 576–580
 - overview*, 560
 - states*, 560–562
- out-of-band management, 30–31

P

- packets
 - EIGRP (Enhanced Interior Gateway Routing Protocol)

Acknowledgement packets, 281–282
encapsulating, 284–285
Hello packets, 279, 280–281
packet headers and TLV, 285–288
Query packets, 283–284
Reply packets, 283–284
table of, 279–280
Update packets, 281–282
 LSP (link-state packets), 257, 258–259
 OSPF Hello packets, 256–257, 428–430
PAgP (Port Aggregation Protocol), 185–186
parameters (EIGRP), 398–399
parent routes, 232
partial updates, 282
passive interfaces, 401–403, 450–453
 configuration, 300–302
 disabling, 572
 verification, 302
passive mode (LACP), 187
passive-interface command, 300–302, 401–403, 451, 572
passwords (VTP), 59
path cost, 121–124, 234
path selection, troubleshooting, 568
path-vector routing protocol, 223
PDMs (protocol-dependent modules), 247, 276–277
performance tuning. *See* tuning
periodic updates, 227, 242
Perlman, Radia, 117
Per-VLAN Spanning Tree. *See* PVST+
PoE (Power over Ethernet), 23–24
point-to-multipoint access, 530
point-to-point links, 152
point-to-point networks, 529
Port Aggregation Protocol (PAgP), 185–186
port channel interfaces, 14, 182
Port ID field (BPDU), 129
PortFast, 156–158
Portfolio RIP and EIGRP class activity, 358
ports
 alternate ports, 127–128
 assigning to VLANs, 66
 backup ports, 118
 density, 21–22
 designated ports, 118, 127–128
 disabled ports, 118

 edge ports, 150–151
 port channel interfaces, 14, 182
 roles, 117–119, 124–126
 root ports, 118
 routed ports, inter-VLAN routing with, 94–95
 states, 144–146
 troubleshooting, 77–79
Power over Ethernet (PoE), 23–24
preemption (HSRP), 205
priority
 HSRP (Hot Standby Router Protocol), 204–205
 OSPF (Open Shortest Path First), 544–546
process information (OSPF), 466–468
Protocol ID field (BPDU), 129
protocol-dependent modules (PDMs), 247, 276–277
PuTTY, 31
PVST+. *See also* Rapid PVST+
 BPDU Guard, 156–158
 bridge IDs, 154–155
 Catalyst 2960 default configuration, 153–154
 characteristics of, 142
 definition of, 141
 extended system ID, 145–147
 load balancing, 158–160
 overview, 143–144
 port states, 144–146
 PortFast, 156–158

Q

quad zero default static route, 380
Query packets, 280, 283–284
question mark (?), 292

R

rack units, 20
Rapid PVST+ 161–163
 BPDUs (bridge protocol data units), 149–150
 characteristics of, 142
 definition of, 141
 edge ports, 150–151
 link types, 152–153
 overview, 148–149
Rapid Spanning Tree Protocol. *See* RSTP (Rapid Spanning Tree Protocol)

- RD (reported distance), 326–327, 376
- redistribute static command, 381
- redistribution, route, 501
- redundancy. *See also* STP (Spanning Tree Protocol)
 - NSF (Nonstop Forwarding), 399
 - planning for, 10–11
 - routers, 199–200
- redundant switched networks. *See* STP (Spanning Tree Protocol)
- reference bandwidth, 456–460
- regional Internet registry (RIR), 291
- Reliable Transport Protocol (RTP), 278
- repairing STP (Spanning Tree Protocol), 169
- Reply packets, 280, 283–284
- reported distance (RD), 326–327, 376
- requests, advertisement, 52
- Retransmission Timeout (RTO), 303
- RIP (Routing Information Protocol), 223, 245–246
- RIPng, 246
- RIR (regional Internet registry), 291
- roles
 - DRs (designated routers), 535–538
 - port roles, 117–119, 124–126
- root bridges, 119–120
- Root ID field (BPDU), 129
- root path cost, 121–124
- Root Path Cost field (BPDU), 129
- root ports, 118
- route discovery (EIGRP)
 - convergence, 312–313
 - neighbor adjacency, 310–311
 - topology table, 311–312
- route propagation
 - EIGRP (Enhanced Interior Gateway Routing Protocol)
 - IPv4, 380–382
 - IPv6, 383–384
 - verification, 382–383
 - OSPF (Open Shortest Path First), 547
 - OSPFv2, 547–548
 - OSPFv3, 551–552
 - propagated IPv4 route verification, 549–550
 - propagated IPv6 route verification, 552–554
- route redistribution, 501
- routed ports, inter-VLAN routing with, 94–95
- router eigrp command, 292–293, 372
- Router ID field (Hello packets), 429
- router IDs
 - EIGRP, 293–296
 - OSPFv2, 441–447
 - assigning, 445
 - configuration, 442–445
 - loopback interfaces as, 447
 - modifying, 445–447
 - network topology, 441–442
 - order of precedence, 443
 - router OSPF configuration mode, 442
 - verifying, 446
 - OSPFv3
 - assigning, 479
 - configuration, 477–479
 - modifying, 479–480
 - verifying, 480
- router link entries, 502–503
- router ospf command, 445
- router priority field (Hello packets), 430
- router-id command, 444, 479, 541
- router-on-a-stick inter-VLAN routing, 78
- routers. *See also* router IDs
 - ASBRs (Autonomous System Boundary Routers), 547
 - backup routers, 202
 - BDRs (backup designated routers), 435–438
 - adjacencies, 538–540
 - election process, 540–543
 - roles, 535–538
 - configuration, troubleshooting, 82–83
 - DROTHERs, 437
 - DRs (designated routers), 435–438, 533–534
 - adjacencies, 538–540
 - election process, 540–543
 - roles, 535–538
 - edge routers, 547
 - enabling, 32
 - failover, 200–201
 - gateways, 547
 - HSRP (Hot Standby Router Protocol)
 - configuration, 206–208
 - debug commands, 210–213
 - definition of, 202
 - failure, 209–210

operation of, 203–204
preemption, 205
priority, 204–205
states and timers, 205–206
verification, 208–209
versions, 204
 IRDP (ICMP Router Discovery Protocol), 202
 master routers, 202
 OSPF (Open Shortest Path First), 499–501
 redundancy, 199–200
 router CLI commands
 basic commands, 30–31
 show commands, 34–38
 router hardware, 26
 Cisco routers, 27–28
 form factors, 28–29
 router requirements, 26
 running configuration, 33–34
 VRRP (Virtual Router Redundancy Protocol), 202
routing configuration, verifying, 82–83
routing information exchange, 239–241
Routing Information Protocol (RIP), 223, 245–246
routing tables
 EIGRP (Enhanced Interior Gateway Routing Protocol), 376–378
 EIGRP for IPv4, 306–309
 EIGRP for IPv6, 355–356
 troubleshooting, 401–408
 OSPF (Open Shortest Path First), 583
 OSPF routes, adding, 264
 OSPFv3, 580–582
 route calculation, 508–509
 routing table entries, 506–508
 troubleshooting, 566–568, 573–575
RSTP (Rapid Spanning Tree Protocol)
 BPDUs (bridge protocol data units), 149–150
 characteristics of, 142
 definition of, 141
 edge ports, 150–151
 port role decisions, 124–126
RTO (Retransmission Timeout), 303
RTP (Reliable Transport Protocol), 278
running configuration
 routers, 33–34
 switches, 39–40

S

scalability, design for, 8–10
 Secure Hash Algorithm (SHA), 418
 security (OSPF), 418
 selecting network devices. *See* network device selection
 servers (VTP), 58–59
 service provider routers, 28
 service provider switches, 18
 SFP (small form-factor pluggable) devices, 22
 SHA (Secure Hash Algorithm), 418
 shared links, 152
 Shortest Path First. *See* SPF (Shortest Path First)
 show cdp neighbors command, 38, 164
 show dtp interface command, 73–74
 show etherchannel port-channel command, 192–193
 show etherchannel summary command, 192–197
 show interfaces command, 35–36, 41–42, 79–80, 315–316, 460–461
 show interfaces etherchannel command, 193–194
 show interfaces port-channel command, 191–192
 show interfaces vlan command, 67–69
 show ip eigrp neighbors command, 302–303, 392–393
 show ip eigrp topology all-links command, 333, 375
 show ip eigrp topology command, 328–332
 show ip interface brief command, 37, 397, 566
 show ip interface command, 36–37
 show ip ospf command, 466–467, 564–565
 show ip ospf database command, 517–518
 show ip ospf interface brief command, 468, 516
 show ip ospf interface command, 459, 468–469, 536, 554–555, 557, 563–564, 566, 571
 show ip ospf neighbor command, 465–466, 538, 555, 556, 563, 566
 show ip protocols command, 34
 EIGRP (Enhanced Interior Gateway Routing Protocol), 296, 304–306, 371, 372–375, 382, 388–389, 393–394
 OSPF (Open Shortest Path First), 446, 447, 451–452, 466, 515–516, 562–563, 571–572
 show ip route command, 35, 306–309, 382, 516–517, 549–550
 show ip route eigrp command, 371, 393

- show ip route ospf command, 565–566, 568
- show ipv6 eigrp neighbors command, 352–354
- show ipv6 interface brief command, 348, 354, 476–477
- show ipv6 ospf command, 483, 578–579
- show ipv6 ospf database command, 520–521
- show ipv6 ospf interface brief command, 483, 519
- show ipv6 ospf interface command, 483–484, 536, 559, 578
- show ipv6 ospf neighbor command, 482, 558, 577
- show ipv6 protocols command, 354–355, 480, 483, 518, 577
- show ipv6 route command, 355–356, 384, 508
- show ipv6 route ospf command, 484, 519, 551, 579
- show ipv6 route static command, 552–553
- show mac address-table command, 42
- show port-security address command, 41
- show protocols command, 38
- show running-config command, 80–81, 158, 367–369
- show spanning-tree command, 122, 155, 159, 165
- show spanning-tree vlan command, 164, 165
- show standby brief command, 208–209
- show standby command, 208–209
- show vlan brief command, 62, 66
- show vlan command, 67–68, 70
- show vtp password command, 59
- show vtp status command, 53–55, 58, 62–63
- shutdown command, 339
- single-area OSPF (Open Shortest Path First), 15, 424
 - costs, 453–464
 - accumulated costs, 455–456
 - calculating, 454–455
 - reference bandwidth, 456–460
 - setting manually, 463–464
 - default routes, 547
 - propagating in OSPFv2, 547–548
 - propagating in OSPFv3, 551–552
 - verifying in OSPFv2, 549–550
 - verifying in OSPFv3, 552–554
 - interface bandwidth
 - adjusting, 462–463
 - default interface bandwidths, 460–462
 - intervals
 - dead intervals, 554–555
 - modifying in OSPFv2, 555–557
 - modifying in OSPFv3, 557–559
 - limitations of, 494
 - messages, 426–431
 - DBD (Database Description) packet, 428
 - encapsulating, 426–427
 - Hello intervals, 430
 - Hello packets, 428–430
 - link-state updates, 430–431
 - LSAck (Link-State Acknowledgement) packet, 428
 - LSR (Link-State Request) packet, 428
 - LSU (Link-State Update) packet, 428
 - in multiaccess networks
 - challenges, 531–533
 - DR/BDR adjacencies, 538–540
 - DR/BDR election process, 540–543
 - DR/BDR roles, 535–538
 - network types, 528–531
 - OSPF DRs, 533–534
 - OSPF priority, 544–546
 - network topology, 441–442
 - operation, 431–440
 - BDRs (backup designated routers), 435–438
 - database synchronization, 438–440
 - DRs (designated routers), 435–438
 - neighbor adjacencies, 433–435
 - states, 432–433
- OSPFv2 configuration, 448–453
 - enabling OSPF on interfaces, 448
 - network command, 449–450
 - passive interfaces, 450–453
 - wildcard masks, 448–449
- OSPFv2 router IDs, 441–447
 - assigning, 445
 - configuration, 442–445
 - loopback interfaces as, 447
 - modifying, 445–447
 - network topology, 441–442
 - order of precedence, 443
 - router OSPF configuration mode, 442
 - verifying, 446
- OSPFv2 verification, 464–469
 - interface settings, 468–469
 - neighbors, 465–466
 - process information, 466–468
 - protocol settings, 466
- OSPFv2 versus OSPFv3, 469–473
 - differences, 471–472
 - link-local addresses, 472–473

- similarities*, 471
- OSPFv3 configuration, 473–481
 - enabling OSPFv3 on interfaces*, 481
 - link-local addresses*, 475–477
 - network topology*, 473–475
 - router IDs*, 477–480
- OSPFv3 router IDs
 - assigning*, 479
 - configuration*, 477–479
 - modifying*, 479–480
 - verifying*, 480
- OSPFv3 verification, 481–485
- reference bandwidth, 456–460
- troubleshooting
 - flowcharts*, 566–568
 - OSPFv2 neighbor issues*, 569–573
 - OSPFv2 routing table issues*, 573–575
 - OSPFv2 troubleshooting commands*, 562–566
 - OSPFv3*, 576–582
 - overview*, 560
 - states*, 560–562
- single-homed, 225
- small form-factor pluggable (SFP) devices, 22
- Smooth Round Trip Timer (SRTT), 303
- source IP to destination IP load balancing, 183
- source IPv6 addresses, 344, 472
- source MAC to destination MAC load balancing, 183
- Spanning Tree Algorithm. *See* STA (Spanning Tree Algorithm)
- spanning tree mode (Rapid PVST+), 161–163
- Spanning Tree Protocol. *See* STP (Spanning Tree Protocol)
- spanning-tree bpduguard enable command, 157
- spanning-tree cost command, 122
- spanning-tree link-type command, 152
- spanning-tree mode rapid-pvst command, 161
- spanning-tree portfast bpduguard default command, 158
- spanning-tree portfast command, 157
- spanning-tree vlan command, 154–155, 159
- speak state (HSRP), 206
- SPF (Shortest Path First), 248
 - example of, 249–251
 - tree, building, 260–263
- SPF Troubleshooting Mastery class activity, 585
- split horizon, 241
- SRTT (Smooth Round Trip Timer), 303
- STA (Spanning Tree Algorithm), 114–117
 - port roles, 117–119
 - root bridges, 119–120
 - root path cost, 121–124
- stack master, 169
- stackable configuration switches, 19–20
- standby ip-address command, 207
- standby preempt command, 207, 212
- standby priority command, 207
- standby routers, 200. *See also* HSRP (Hot Standby Router Protocol)
- standby state (HSRP), 206
- standby version 2 command, 207
- states
 - HSRP (Hot Standby Router Protocol), 205–206
 - multiaccess networks, 538–539
 - OSPF (Open Shortest Path First), 432–433, 560–562
 - port states, 144–146
- static route propagation
 - EIGRP (Enhanced Interior Gateway Routing Protocol)
 - IPv4*, 380–382
 - IPv6*, 383–384
 - verification*, 382–383
 - OSPF (Open Shortest Path First), 547
 - OSPFv2*, 547–548
 - OSPFv3*, 551–552
 - propagated IPv4 route verification*, 549–550
 - propagated IPv6 route verification*, 552–554
- status (STP), 165
- Stepping Through OSPFv3 class activity, 486
- Stormy Traffic class activity, 107
- STP (Spanning Tree Protocol), 10, 107
 - 802.1D BPDU (bridge protocol data unit)
 - frame format*, 128–131
 - propagation and process*, 131–136
 - alternate ports, 127–128
 - BDPUs, 118
 - broadcast storms, 111–112
 - characteristics of, 141
 - comparison of protocols, 142–143
 - designated ports, 127–128
 - duplicate unicast frames, 113
 - extended system ID, 136–140

- MAC database instability, 109–111
- MSTP (Multiple Spanning Tree Protocol)
 - characteristics of*, 142
 - definition of*, 141
- PVST+
 - BPDUs Guard*, 156–158
 - bridge IDs*, 154–155
 - Catalyst 2960 default configuration*, 153–154
 - characteristics of*, 142
 - definition of*, 141
 - extended system ID*, 145–147
 - load balancing*, 158–160
 - overview*, 143–144
 - port states*, 144–146
 - PortFast*, 156–158
- Rapid PVST+
 - 161–163
 - BPDUs (bridge protocol data units)*, 149–150
 - characteristics of*, 142
 - definition of*, 141
 - edge ports*, 150–151
 - link types*, 152–153
 - overview*, 148–149
- redundancy at OSI layers 1 and 2, 108–109
- repairing, 169
- RSTP (Rapid Spanning Tree Protocol)
 - BPDUs (bridge protocol data units)*, 149–150
 - characteristics of*, 142
 - definition of*, 141
 - edge ports*, 150–151
 - port role decisions*, 124–126
- STA (Spanning Tree Algorithm), 114–117
 - port roles*, 117–119
 - root bridges*, 119–120
 - root path cost*, 121–124
- switch stacking, 169–172
- troubleshooting, 163–169
 - expected versus actual topology*, 164–165
 - failure*, 166–168
 - status*, 165
 - STP repair*, 169
- stub networks, 268
- subnet masks
 - troubleshooting, 83–87
 - VLSM (variable-length subnet mask), 228
- subset advertisements, 52
- successor distance, 324–325
- summarization. *See* automatic summarization (EIGRP)
- summary advertisements, 52
- summary route (EIGRP), 378–380
- SVIs (switch virtual interfaces)
 - definition of*, 91
 - inter-VLAN routing with*, 91–94
- switch platforms, 17–21
- switch virtual interfaces. *See* SVIs (switch virtual interfaces)
- switches. *See also* ports
 - enabling, 39
 - Layer 3 switching, 89–91, 397–398
 - configuration*, 95–96
 - inter-VLAN routing with routed ports*, 94–95
 - inter-VLAN routing with switch virtual interfaces*, 91–94
 - troubleshooting*, 95–98
 - verification*, 570
 - running configuration, 39–40
 - stacking, 169–172
 - SVIs (switch virtual interfaces)
 - definition of*, 91
 - inter-VLAN routing with*, 91–94
 - switch CLI commands
 - basic commands*, 38–40
 - show commands*, 40–42
 - switch hardware, 17
 - forwarding rates*, 22
 - multilayer switching*, 24–25
 - PoE (Power over Ethernet)*, 23–24
 - port density*, 21–22
 - switch platforms*, 17–21
 - verification*, 79–81
 - VLAN ranges on, 63–64
- switchport access vlan command, 66
- switchport command, 91
- switchport mode access command, 66, 73
- switchport mode dynamic auto command, 73
- switchport mode dynamic desirable command, 73
- switchport mode trunk command, 73, 78–79
- switchport nonegotiate command, 73
- synchronizing OSPF (Open Shortest Path First) databases, 438–440

T

tables. *See also* routing tables

- neighbor tables, 352–354, 583
- topology table (EIGRP), 247, 311–312
 - EIGRP for IPv6*, 345–347
 - show ip eigrp topology command*, 328–334

TeraTerm, 31

timers (HSRP), 205–206

TLV (type, length, value) field, 285–288

topology

- EIGRP (Enhanced Interior Gateway Routing Protocol), 311–312, 367–369
 - EIGRP for IPv4*, 289–290
 - EIGRP for IPv6*, 345–347
 - no feasible successor*, 332–334
 - show ip eigrp topology command*, 328–332

OSPF (Open Shortest Path First)

- OSPFv2*, 441–442
- OSPFv3*, 473–475

STP (Spanning Tree Protocol), 164–165

traditional inter-VLAN routing, 77

traffic-share balanced command, 391

transparent VTP mode, 48

trees (SPF), 260–263

troubleshooting

- EIGRP (Enhanced Interior Gateway Routing Protocol)
 - automatic summarization*, 405–408
 - basic commands*, 392–394
 - components*, 394–395
 - EIGRP parameters*, 398–399
 - interfaces*, 399–401
 - Layer 3 connectivity*, 397–398
 - missing network statement*, 403–405
 - neighbor issues*, 397–401
 - passive interfaces*, 401–403
 - routing table issues*, 401–408
- EtherChannel, 194–197
- HSRP (Hot Standby Router Protocol)
 - common configuration issues*, 213
 - debug commands*, 210–213
 - failure*, 209–210
- multiarea OSPF (Open Shortest Path First)
 - data structures*, 583
 - overview*, 582

single-area OSPF (Open Shortest Path First)

- flowcharts*, 566–568
- OSPFv2 neighbor issues*, 569–573
- OSPFv2 routing table issues*, 573–575
- OSPFv2 troubleshooting commands*, 562–566
- OSPFv3*, 580–582
- OSPFv3 troubleshooting commands*, 576–580
- overview*, 560
- states*, 560–562

STP (Spanning Tree Protocol), 163–169

- expected versus actual topology*, 164–165
- failure*, 166–168
- status*, 165
- STP repair*, 169

VLANs (virtual local area networks)

- deleting VLANs*, 75–77
- DTP (Dynamic Trunking Protocol) issues*, 89
- interface issues*, 81
- IP addressing issues*, 83–87
- Layer 3 switching*, 95–98
- routing configuration*, 82–83
- switch configuration*, 79–81
- switch port issues*, 77–79

VTP (VLAN Trunking Protocol), 88

trunking

- DTP (Dynamic Trunking Protocol), 48
 - initial configuration*, 71
 - negotiated interface modes*, 72–73
 - troubleshooting*, 89
 - verification*, 72–73
- VTP (VLAN Trunking Protocol), 48–49
 - advertisements*, 52
 - cautions*, 55–56
 - components*, 50
 - configuration*, 53–55, 57–63
 - modes*, 50–51
 - transparent VTP mode*, 48
 - troubleshooting*, 88
 - verification*, 62–63
 - versions*, 53

tuning

- EIGRP (Enhanced Interior Gateway Routing Protocol), 366
 - automatic summarization*, 366–380
 - bandwidth utilization*, 385–386
 - default route propagation*, 380–384

Hello and Hold timers, 367–386
IPv4 load balancing, 388–390
IPv6 load balancing, 390–392
 multiaccess networks, OSPF (Open Shortest Path First) in
challenges, 531–533
DR/BDR adjacencies, 538–540
DR/BDR election process, 540–543
DR/BDR roles, 535–538
network types, 528–531
OSPF DRs, 533–534
OSPF priority, 544–546
 OSPF (Open Shortest Path First)
default route propagation, 547–554
interfaces, 554–559
in multiaccess networks, 528–546
 routing protocols, 15–17
 Tuning EIGRP class activity, 410
 two-layer area hierarchy, 498–499
 Two-Way state (OSPF), 433, 561
 type, length, value (TLV) field, 285–288
 type 1 LSAs (link-state advertisements), 502–503
 type 2 LSAs (link-state advertisements), 503–504
 type 3 LSAs (link-state advertisements), 504–505
 type 4 LSAs (link-state advertisements), 505
 type 5 LSAs (link-state advertisements), 506
 Type field (Hello packets), 429

U

ultimate routes, 232
 unequal-cost load balancing, 391
 unicast frames, duplicate, 113
 Update packets, 279, 281–282
 updates, link-state
 flooding LSPs, 258–259
 Hello packets, 256–257
 link-state routing process, 251–253
 LSDB (link-state database), 259–260
 LSP (link-state packets), 257
 OSPF routes, 264
 SPF (Shortest Path First) tree, 260–263

V

variable-length subnet mask (VLSM), 228
 variance command, 391

verification. *See also* troubleshooting
 DRs (designated routers)
adjacencies, 538–540
roles, 535–538
 DTP (Dynamic Trunking Protocol), 73–74
 EIGRP for IPv4
automatic summarization, 372–378
EIGRP processes, 296
neighbors, 302–304
passive interfaces, 302
propagated default routes, 382–383
routing table, 306–309
show ip protocols command, 304–306
 EIGRP for IPv6
neighbor table, 352–354
routing table, 355–356
show ipv6 protocols command, 354–355
 EtherChannel, 191–194
 HSRP (Hot Standby Router Protocol), 208–209
 IP addresses, 85–87
 IPv6-enabled interfaces, 475
 k values, 313–315
 link-local addresses, 477
 multiarea OSPF (Open Shortest Path First)
multiarea OSPFv2, 515–518
multiarea OSPFv3, 518–521
 propagated default routes
IPv4, 549–550
IPv6, 552–554
 router IDs, 480
 routing configuration, 82–83
 single-area OSPF (Open Shortest Path First)
costs, 459
default bandwidth, 460–461
interface settings, 468–469
Layer 3 connectivity, 570
OSPFv2, 464–469
OSPFv3, 481–485
passive interfaces, 451–452
propagated IPv4 routes, 549–550
propagated IPv6 routes, 552–554
route metric, 459–460
router IDs, 446
 switch configuration, 79–81
 VLANs (virtual local area networks), 67–69
 VTP (VLAN Trunking Protocol), 62–63
 Version 0 BPDUs, 149

Version 2 BPDUs, 149–150
 Version field (BPDU), 129
 versions
 HSRP (Hot Standby Router Protocol), 204
 VTP (VLAN Trunking Protocol), 53
 virtual links, 530
 virtual local area networks. *See* VLANs (virtual local area networks)
 virtual networking switches, 18
 Virtual Router Redundancy Protocol (VRRP), 202
 vlan command, 66
 VLAN Trunking Protocol. *See* VTP (VLAN Trunking Protocol)
 vlan.dat database, 50
 VLANs (virtual local area networks), 48
 assigning ports to, 66
 configuration on VTP server, 60–61
 creating, 65–66
 deleting, 75–77
 DTP (Dynamic Trunking Protocol)
 initial configuration, 71
 negotiated interface modes, 72–73
 troubleshooting, 89
 verification, 72–73
 extended VLANs, 63
 definition of, 64
 VLAN ranges on Catalyst switches, 63–64
 Layer 3 switching, 89–91
 configuration, 95–96
 inter-VLAN routing with routed ports, 94–95
 inter-VLAN routing with switch virtual interfaces, 91–94
 troubleshooting, 95–98
 naming, 65
 normal-range VLANs, 64
 troubleshooting
 deleting VLANs, 75–77
 DTP (Dynamic Trunking Protocol) issues, 89
 interface issues, 81
 IP addressing issues, 83–87
 routing configuration, 82–83
 switch configuration, 79–81
 switch port issues, 77–79
 VTP (VLAN Trunking Protocol) issues, 88
 verification, 67–69
 VTP (VLAN Trunking Protocol), 48–49

advertisements, 52
 cautions, 55–56
 components, 50
 configuration, 53–55, 57–63
 modes, 50–51
 transparent VTP mode, 48
 troubleshooting, 88
 verification, 62–63
 versions, 53
 VLSM (variable-length subnet mask), 228
 VRRP (Virtual Router Redundancy Protocol), 202
 VTP (VLAN Trunking Protocol), 48–49
 advertisements, 52
 cautions, 55–56
 components, 50
 configuration, 57
 clients, 60
 default configuration, 53–55
 domain name and password, 59
 verification, 62–63
 VLANs, 60–61
 VTP server, 58–59
 modes, 50–51
 transparent VTP mode, 48
 troubleshooting, 88
 versions, 53
 vtp domain command, 59
 vtp mode server command, 58
 vtp password command, 59

W-X-Y-Z

wildcard masks, 296–300, 448–449
 wire speed, 22
 wired LAN (local area network) design, 4
 Cisco validated designs
 hierarchical design model, 6–8
 need for network scaling, 4–6
 network expansion
 access layer, 14–15
 bandwidth, 13–14
 design for scalability, 8–10
 failure domains, 11–13
 redundancy planning, 10–11
 routing protocols, fine-tuning, 15–17
 wireless access points, 23



The Cisco Learning Network

The IT Community that helps you get Cisco Certified.



Be a Part of the
Community



Prepare for
Success



Interact with
Professionals



Mentor, Share,
Achieve

Join over 1 Million Members on the Cisco Learning Network, featuring powerful study resources like IT Training Videos, Study Groups and Certification Exam Topics.

Connect with us on social media at:
cs.co/LearningatCisco-About



ciscolearningnetwork.com

