

Implementing Cisco IP Switched Networks (SWITCH)

Foundation Learning Guide

CCNP SWITCH 300-115



ciscopress.com

Richard Froom, CCIE No. 5102

Erum Frahim, CCIE No. 7549

FREE SAMPLE CHAPTER



SHARE WITH OTHERS



Cisco
Press

NEW Complete Video Courses for CCNP Routing & Switching 300 Series Exams



These unique products include multiple types of video presentations, including:



- Live instructor whiteboarding
- Real-world demonstrations
- Animations of network activity
- Dynamic KeyNote presentations
- Doodle videos
- Hands-on command-line interface (CLI) demonstrations
- Review quizzes

Complete Video Course
CCNP Routing and Switching v2.0
Kevin Wallace, Elan Beer, and Chris Avants
livelessons®
9780789754493

CCNP Routing and Switching v2.0 — Complete Video Course Library
Specially priced library including ALL THREE Complete Video Courses: *CCNP Routing and Switching ROUTE 300-101*, *CCNP Routing and Switching SWITCH 300-115*, and *CCNP Routing and Switching TSHOOT 300-135*.

Complete Video Course
CCNP Routing and Switching ROUTE 300-101
Kevin Wallace
livelessons®
9780789753731

CCNP Routing and Switching ROUTE 300-101 — Complete Video Course
149 VIDEOS with 12+ HOURS of video instruction from best-selling author, expert instructor, and double CCIE Kevin Wallace walk you through the full range of topics on the CCNP Routing and Switching ROUTE 300-101 exam, including fundamental routing concepts; IGP routing protocols including RIPng, EIGRP, and OSPF; route distribution and selection; BGP; IPv6 Internet connectivity; router security; and routing protocol authentication.

Complete Video Course
CCNP Routing and Switching SWITCH 300-115
Wayne Lewis
livelessons®
9780789754073

CCNP Routing and Switching SWITCH 300-115 — Complete Video Course
10+ HOURS of unique video training walks you through the full range of topics on the CCNP SWITCH 300-115 exam. This complete video course takes you from the design and architecture of switched networks through the key technologies vital to implementing a robust campus network. You will learn, step-by-step, configuration commands for configuring Cisco switches to control and scale complex switched networks.

Complete Video Course
CCNP Routing and Switching TSHOOT 300-135
Elan Beer and Chris Avants
livelessons®
9780789754295

CCNP Routing and Switching TSHOOT 300-135 — Complete Video Course
10+ HOURS of unique video instruction from expert instructors and consultants Elan Beer and Chris Avants walks you through the full range of topics on the CCNP TSHOOT 300-135 exam. This complete video course teaches you the skills you need to plan and perform regular maintenance on complex enterprise routed and switched networks and how to use technology-based practices and a systematic ITIL-compliant approach to perform network troubleshooting commands for configuring Cisco switches to control and scale complex switched networks.

SAVE ON ALL NEW
CCNP R&S 300 Series Products
www.CiscoPress.com/CCNP

Implementing Cisco IP Switched Networks (SWITCH) Foundation Learning Guide

Richard Froom, CCIE No. 5102
Erum Frahim, CCIE No. 7549

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

Implementing Cisco IP Switched Networks (SWITCH) Foundation Learning Guide

Richard Froom, CCIE No. 5102
Erum Frahim, CCIE No. 7549

Copyright© 2015 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 May 2015

Library of Congress Control Number: 2015934731

ISBN-13: 978-1-58720-664-1

ISBN-10: 1-58720-664-1

Warning and Disclaimer

This book is designed to provide information about Cisco CCNP switching. 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.

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 international@pearsoned.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.

Publisher: Paul Boger

Copy Editor: Keith Cline

Associate Publisher: Dave Dusthimer

Technical Editor: Sean Wilkins

Business Operations Manager, Cisco Press:
Jan Cornelssen

Editorial Assistant: Vanessa Evans

Executive Editor: Mary Beth Ray

Designer: Mark Shirar

Managing Editor: Sandra Schroeder

Composition: Bronkella Publishing LLC

Development Editor: Box Twelve Communications

Indexer: Tim Wright

Project Editor: Mandie Frank

Proofreader: The Wordsmithery LLC



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.



CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (08128)

About the Authors

Richard Froom, CCIE No. 5102, is a manager within the Solution Validation Services (SVS) team at Cisco. Richard previously worked as a network engineer in the Cisco TAC and in various customer-facing testing organizations within Cisco. Richard holds CCIEs in Routing and Switching and in Storage Networking. Richard currently focuses on expanding his team's validation coverage to new technologies in the data center, including Application Centric Infrastructure (ACI), OpenStack, Intercloud Fabric, and big data solutions with Hadoop.

Erum Frahim, CCIE No. 7549, is a technical leader working in the Solution Validation Services (SVS) group at Cisco. In her current role, Erum is leading efforts to test data center solutions for several Cisco high-profile customers and leading all the cross-business units interlock. Most recently, she is working on Application Centric Infrastructure (ACI), UCS Director, OpenStack, and big data. Before this, Erum managed the Nexus platform escalation group and served as a team lead for the data center storage-area network (SAN) test lab under the Cisco data center business unit. Erum joined Cisco in 2000 as a technical support engineer. Erum has a Master of Science degree in electrical engineering from Illinois Institute of Technology and also holds a Bachelor of Engineering degree from NED University, Karachi, Pakistan. Erum also authors articles in *Certification Magazine* and on Cisco.com and has participated in many CiscoLive Events. In her spare time, Erum enjoys her time with her husband and child.

About the Technical Reviewer

Sean Wilkins is an accomplished networking consultant for SR-W Consulting (<http://www.sr-wconsulting.com>) and has been in the field of IT since the mid-1990s, working with companies such as Cisco, Lucent, Verizon, and AT&T, in addition to several other private companies. Sean currently holds certifications with Cisco (CCNP/CCDP), Microsoft (MCSE), and CompTIA (A+ and Network+). He also has a Master of Science degree in Information Technology with a focus in network architecture and design, a Master of Science degree in Organizational Management, a Masters Certificate in Network Security degree, a Bachelor of Science degree in Computer Networking, and an Associate of Applied Science in Computer Information Systems degree. In addition to working as a consultant, Sean spends a lot of his time as a technical writer and editor for various companies.

Dedications

From Richard:

This book is dedicated to my wife, Elizabeth, and my son, Nathan. Thank you for your encouragement and patience as I completed this effort.

From Erum:

This book is dedicated to my daughter, my hubby, and my parents, for their love and patience all throughout this process.

Acknowledgments

We want to thank many people for helping to put this book together.

The Cisco Press team: Mary Beth Ray, the executive editor, coordinated the whole project, steered the book through the necessary processes, and understood when the inevitable snags appeared. Sandra Schroeder, the managing editor, brought the book to production. Vanessa Evans was once again wonderful at organizing the logistics and administration. Jeff Riley, the development editor, has been invaluable in coordinating and ensuring that we all focused on producing the best manuscript.

We also want to thank Mandie Frank, the project editor, and Keith Cline, the copy editor, for their excellent work in getting this book through the editorial process.

The Cisco Switch course development team: Many thanks to the members of the team who developed the Switch course. The course was a basis for this book, and without it, we would never have completed the text in short order.

The technical reviewers: We want to thank the technical reviewer of this book, Sean Wilkins, for his thorough review and valuable input.

Our families: Of course, this book would not have been possible without the endless understanding and patience of our families. They have always been there to motivate and inspire us, and we are forever grateful.

Contents at a Glance

	Introduction	xx
Chapter 1	Fundamentals Review	1
Chapter 2	Network Design Fundamentals	9
Chapter 3	Campus Network Architecture	41
Chapter 4	Spanning Tree in Depth	119
Chapter 5	Inter-VLAN Routing	203
Chapter 6	First-Hop Redundancy	247
Chapter 7	Network Management	305
Chapter 8	Switching Features and Technologies for the Campus Network	351
Chapter 9	High Availability	393
Chapter 10	Campus Network Security	409
Appendix A	Answers to Chapter Review Questions	469
	Index	473

Contents

	Introduction	xx
Chapter 1	Fundamentals Review	1
	Switching Introduction	2
	Hubs and Switches	2
	Bridges and Switches	2
	Switches of Today	3
	Broadcast Domains	3
	MAC Addresses	4
	The Basic Ethernet Frame Format	4
	Basic Switching Function	5
	VLANs	6
	The Spanning Tree Protocol	6
	Trunking	7
	Port Channels	7
	Multilayer Switching	8
	Summary	8
Chapter 2	Network Design Fundamentals	9
	Campus Network Structure	9
	Hierarchical Network Design	10
	<i>Access Layer</i>	12
	<i>Distribution Layer</i>	13
	<i>Core Layer (Backbone)</i>	14
	Layer 3 in the Access Layer	17
	The Cisco Enterprise Campus Architecture	19
	The Need for a Core Layer	20
	Types of Cisco Switches	22
	Comparing Layer 2 and Multilayer Switches	24
	<i>MAC Address Forwarding</i>	24
	<i>Layer 2 Switch Operation</i>	25
	<i>Layer 3 (Multilayer) Switch Operation</i>	26
	<i>Useful Commands for Viewing and Editing Catalyst Switch MAC Address Tables</i>	27
	<i>Frame Rewrite</i>	28
	<i>Distributed Hardware Forwarding</i>	28

Cisco Switching Methods	29	
<i>Route Caching</i>	30	
<i>Topology-Based Switching</i>	31	
Hardware Forward Details	33	
Study Tips	34	
Summary	34	
Review Questions	35	
Chapter 3	Campus Network Architecture	41
Implementing VLANs and Trunks in Campus Environment		41
VLAN Overview		42
VLAN Segmentation		44
<i>End-to-End VLANs</i>		44
<i>Local VLANs</i>		45
<i>Comparison of End-to-End VLANs and Local VLANs</i>		46
<i>Mapping VLANs to a Hierarchical Network</i>		47
Implementing a Trunk in a Campus Environment		49
<i>Understanding Native VLAN in 802.1Q Trunking</i>		52
<i>Understanding DTP</i>		53
<i>VLAN Ranges and Mappings</i>		54
Configuring, Verifying, and Troubleshooting VLANs and Trunks		55
<i>Verifying the VLAN Configuration</i>		57
Configuring VLANs and Trunks		61
Best Practices for VLANs and Trunking		65
Voice VLAN Overview		67
Switch Configuration for Wireless Network Support		69
VLAN Trunking Protocol		70
VTP Overview		70
VTP Modes		71
VTP Versions		73
VTP Pruning		74
VTP Authentication		75
VTP Advertisements		75
VTP Messages Types		77
<i>Summary Advertisements</i>		77
<i>Subset Advertisements</i>		77

Configuring and Verifying VTP	78
Overwriting VTP Configuration (Very Common Issue with VTP)	87
Best Practices for VTP Implementation	93
Implementing EtherChannel in a Switched Network	94
The Need for EtherChannel	94
EtherChannel Mode Interactions	97
<i>LACP</i>	97
<i>PAgP</i>	98
Layer 2 EtherChannel Configuration Guidelines	99
EtherChannel Load-Balancing Options	100
Configuring EtherChannel in a Switched Network	102
<i>EtherChannel Configuration and Load Balancing</i>	103
<i>EtherChannel Guard</i>	108
Study Tips	109
Summary	110
Review Questions	110
Chapter 4 Spanning Tree in Depth	119
Spanning Tree Protocol Overview	120
STP Need	120
STP Standards	121
STP Operations	122
Bridge Protocol Data Units	124
Root Bridge Election	124
Root Port Election	126
Designated Port Election	128
STP Port States	129
Per-VLAN STP Plus (PVST+)	130
STP Topology Changes	131
Rapid Spanning Tree Protocol	133
RSTP Port Roles	134
Comparison of RSTP and STP Port States	135
RSTP Topology Changes	136
RSTP Link Types	138

Configuring and Modifying STP Behavior	140
<i>Changing STP Priority</i>	143
<i>STP Path Manipulation</i>	145
<i>STP Timers</i>	148
Implementing STP Stability Mechanisms	151
Use UplinkFast	153
Use BackboneFast	154
Use PortFast	156
Securing PortFast Interface with BPDU Guard	158
Disabling STP with BPDU Filter	159
Use Root Guard	161
Loop Guard Overview	164
Use UDLD	166
<i>UDLD Recommended Practices</i>	170
Use FlexLinks	171
STP Stability Mechanisms Recommendations	175
Configuring Multiple Spanning Tree Protocol	179
Introducing MST	179
MST Regions	182
STP Instances with MST	183
Extended System ID for MST	185
Configuring and Verifying MST	185
Configuring MST Path Cost	192
Configuring MST Port Priority	193
MST Protocol Migration	194
MST Recommended Practices	194
Troubleshooting STP	196
Potential STP Problems	196
<i>Duplex Mismatch</i>	196
<i>Unidirectional Link Failure</i>	197
<i>Frame Corruption</i>	197
<i>Resource Errors</i>	198
<i>PortFast Configuration Errors</i>	198
Study Tips	198
Summary	199
Review Questions	200

Chapter 5	Inter-VLAN Routing	203
	Describing Inter-VLAN Routing	204
	Introduction to Inter-VLAN Routing	204
	Inter-VLAN Routing Using an External Router	206
	<i>Configuring Inter-VLAN Routing Using an External Router</i>	207
	<i>Routing with an External Router</i>	208
	<i>External Routers: Advantages Disadvantages</i>	211
	Inter-VLAN Routing Using Switch Virtual Interfaces	212
	<i>SVI: Advantages and Disadvantages</i>	214
	Routing with Routed Ports	214
	<i>Routed Ports: Advantages</i>	215
	Configuring Inter-VLAN Routing Using SVI and Routed Ports	216
	<i>Routing on a Multilayer Switch</i>	217
	<i>Using the SVI autostate exclude Command</i>	220
	<i>SVI Configuration Checklist</i>	221
	Troubleshooting Inter-VLAN Problems	222
	<i>Example of a Troubleshooting Plan</i>	223
	Layer 2 Versus Layer 3 EtherChannel	225
	Layer 3 EtherChannel Configuration	226
	Verifying Routing Protocols	229
	Implementing DHCP	231
	DHCP Overview	231
	Configuring DHCP in Multilayer Switched Network	233
	<i>Configuring a DHCP Relay</i>	239
	<i>Configuring DHCP Options</i>	239
	Study Tips	240
	Summary	241
	Review Questions	242
Chapter 6	First-Hop Redundancy	247
	Overview of FHRP and HSRP	247
	The Need for First-Hop Redundancy	248
	HSRP Overview	250
	HSRP State Transition	253
	Aligning HSRP with STP Topology	254

Configuring and Tuning HSRP	255
<i>Forwarding Through the Active Router</i>	257
Load Sharing with HSRP	263
The Need for Interface Tracking with HSRP	265
HSRP Interface Tracking	266
HSRP and Object Tracking	268
Configuring HSRP Authentication	271
Tuning HSRP Timers	272
HSRP Versions	274
Configuring Layer 3 Redundancy with VRRP	274
About VRRP	275
Configuring VRRP and Spotting the Differences from HSRP	276
<i>VRRP and Authentication</i>	279
Tracking and VRRP	280
Configuring Layer 3 Redundancy with GLBP	282
Introducing GLBP	282
Comparing GLPB to HSRP	283
GLBP States	284
Configuring and Verifying GLBP	285
GLBP Load-Balancing Options	294
GLBP Authentication	295
GLBP and STP	295
Tracking and GLBP	296
Study Tips	300
Summary	301
References	301
Review Questions	302

Chapter 7 Network Management 305

AAA	305
Authentication Options	307
RADIUS and TACACS+ Overview	308
<i>RADIUS Authentication Process</i>	309
<i>TACACS+ Authentication Process</i>	310
Configuring AAA	311
Configuring RADIUS for Console and vty Access	311
Configuring TACACS+ for Console and vty Access	312

AAA Authorization	313
AAA Accounting	314
Limitations of TACACS+ and RADIUS	315
Identity-Based Networking	316
IEEE 802.1X Port-Based Authentication Overview	316
IEEE 802.1X Configuration Checklist	318
Network Time Protocols	319
The Need for Accurate Time	320
Configuring the System Clock Manually	320
Network Time Protocol Overview	323
NTP Modes	324
Other NTP Configuration Options	326
NTP Example	326
NTP Design Principles	329
Securing NTP	331
NTP Source Address	333
NTP Versions	333
SNTP	335
PTP/IEEE-1588	336
SNMP	336
SNMP Overview	337
SNMP Versions	339
SNMP Best Practices	339
SNMPv3 Configuration Example	340
<i>Verifying SNMP Version 3 Configuration</i>	342
Study Tips	344
Summary	345
Review Questions	345
Chapter 8 Switching Features and Technologies for the Campus Network	351
Discovery Protocols	352
Introduction to LLDP	352
Basic Configuration of LLDP	353
Discovering Neighbors Using LLDP	355

Unidirectional Link Detection	357
UDLD Mechanisms and Specifics	358
UDLD Configuration	358
Leveraging UDLD and STP Loop Guard Together	360
Power over Ethernet	360
PoE Components	362
PoE Standards	362
<i>PoE Negotiation</i>	362
Configuring and Verifying PoE	363
SDM Templates	364
SDM Template Types	365
Choosing the Right SDM Template	367
System Resource Configuration on Other Platforms	367
Monitoring Features	368
SPAN and RSPAN Overview	368
SPAN Configuration	371
RSPAN Configuration	372
IP SLA	374
Introduction to IP SLA	375
IP SLA Source and Responder	377
IP SLA Configuration	377
IP SLA Operation with Responder	379
IP SLA Time Stamps	381
Configuring Authentication for IP SLA	382
IP SLA Example for UDP Jitter	383
Study Tips	384
Summary	385
Review Questions	385
Chapter 9 High Availability	393
The Need for Logical Switching Architectures	394
What Is StackWise?	395
StackWise Benefits	396
Verifying StackWise	396
What Is VSS?	397
VSS Benefits	398
Verifying VSS	399

Redundant Switch Supervisors	401
Supervisor Redundancy Modes	402
<i>Stateful Switchover</i>	403
<i>Nonstop Forwarding</i>	404
Study Tips	405
Summary	405
Review Questions	406
References	406
Chapter 10 Campus Network Security	409
Overview of Switch Security Issues	410
Cisco Switch Security Configuration Best Practices	411
Campus Network Vulnerabilities	414
Rogue Access	414
Switch Vulnerabilities	415
MAC Flooding Attacks	417
Introducing Port Security	419
Port Security Configuration	420
Port Error Conditions	422
<i>Err-Disabled Automatic Recovery</i>	423
Port Access Lists	424
Storm Control	425
Introduction to Storm Control	426
Configuring and Verifying Storm Control on an Interface	427
Mitigating Spoofing Attacks	430
DHCP Spoofing Attacks	430
DHCP Snooping	432
<i>DHCP Option 82</i>	433
<i>DHCP Snooping Example Configuration</i>	433
IP Source Guard	436
IPSG Configuration	438
ARP Spoofing	439
Dynamic ARP Inspection	440
<i>DAI Configuration</i>	441
Securing VLAN Trunks	443
Switch Spoofing	444
VLAN Hopping	446
<i>Protecting Against VLAN Hopping</i>	447

VLAN Access Lists	448
<i>VACL Interaction with ACLs and PACLs</i>	449
<i>Configuring VACLs</i>	450
Private VLANs	451
Introduction to PVLANS	452
<i>PVLAN Port Types</i>	453
<i>PVLAN Configuration</i>	454
<i>PVLAN Verification</i>	456
<i>PVLANS Across Multiple Switches</i>	457
<i>Using the Protected Port Feature</i>	458
Study Tips	458
Summary	459
Review Questions	460
Appendix A Answers to Chapter Review Questions	469
Index	473

Icons Used in This Book



Router




Switch



Multilayer
Switch



Cisco IOS
Firewall



Route/Switch
Processor



Access Server



PIX Firewall



Laptop



Server



PC



Authentication
Server



Camera
PC/Video



Ethernet
Connection



Serial Line
Connection



Network
Cloud



IP Phone



Analog
Phone

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

This book starts you down the path toward attaining your CCNP or CCDP certification, providing in-depth information to help you prepare for the SWITCH exam (300-115).

The commands and configuration examples presented in this book are based on Cisco Catalyst IOS for the Catalyst 3750 and 6500.

In terms of content, campus networks continue to evolve, scale, and require minimal convergence and downtime. As these campus networks grow to need these parameters, Cisco has created new switching features to support growth of the networks. Features found in spanning-tree enhancements, port channeling, and trunking all drive the evolving campus networks and are discussed in this book, among other features.

Moreover, as with Internet security, security within the campus network is paramount. Most enterprises focus heavily on security at the Internet edge, but focus is also needed on internal security. Rogue access by hackers to either create a denial-of-service attack or steal data is an example where internal security is needed. This book covers the basic building blocks of campus networks, with a new and heavy emphasis placed on campus network security.

In terms of the structure, configuration examples and sample verification outputs throughout this book demonstrate troubleshooting techniques and illustrate critical issues surrounding network operation. Chapter-ending review questions illustrate and will help solidify the concepts presented in this book.

Who Should Read This Book?

This book is intended for network architects, network designers, systems engineers, network managers, and network administrators who are responsible for implementing and troubleshooting campus networks.

If you are planning to take the SWITCH exam toward your CCNP or CCDP certification, this book provides you with in-depth study material. To fully benefit from this book, you should have your CCNA Routing and Switching certification or possess the same level of knowledge, including an understanding of the following topics:

- A working knowledge of the OSI reference model and networking fundamentals
- The ability to operate and configure a Cisco router/switch, including the following:
 - Displaying and interpreting a router's or switch's routing table
 - Configuring management IP address
 - Configuring static and default routes
 - Enabling a switch interface
 - Configuring IP standard and extended access lists
 - Managing network device security

- Configuring network management protocols and managing device configurations and Cisco Catalyst IOS images and licenses
- Verifying router and switch configurations with available tools, such as **show** and **debug** commands
- Working knowledge of the TCP/IP stack and IPv6
- The ability to configure, verify, and troubleshoot basic IP connectivity and switching problems

If you lack this knowledge and these skills, you can gain them by completing the Interconnecting Cisco Network Devices Part 1 (ICND1) and Interconnecting Cisco Network Devices Part 2 (ICND2) courses or by reading the related Cisco Press books.

Switch Exam Topic Coverage

The Cisco website has the following information on the exam topics page for the SWITCH exam (300-115) (available at <https://learningnetwork.cisco.com/docs/DOC-24499>):

“The following topics are general guidelines for the content that is likely to be included on the practical exam. However, other related topics may also appear on any specific delivery of the exam. In order to better reflect the contents of the exam and for clarity purposes, the following guidelines may change at any time without notice.”

The referenced list of exam topics available at the time of this writing is provided in Table I-1.

The Cisco SWITCH course does not cover all the listed exam topics and may not cover other topics to the extent needed by the exam, because of classroom time constraints. The Cisco SWITCH course is not created by the same group that created the exam.

This book does provide information on each of these exam topics (except when the topic is covered by prerequisite material as noted), as identified in the “Where Topic Is Covered” column in Table I-1. This book provides information related to all the exam topics to a depth that should be adequate for the exam. Note, however, that because the wording of the topics is quite general in nature and the exam itself is Cisco proprietary and subject to change, the authors of this book cannot guarantee that all the details on the exam are covered.

As mentioned, some of the listed SWITCH exam topics are actually covered by the prerequisite material. You may already be familiar with this material, and so this book provides pointers to the relevant chapters of the *ICND1* and *ICND2 Foundation Learning Guide* (ISBN: 978-1587143762 and 978-1587143779) Cisco Press books for these topics.

Table I-1 SWITCH Exam Topic Coverage

Topic #	Topic	Where Topic Is Covered
1.0	Layer 2 Technologies	
1.1	Configure and Verify Switch Administration	
	SDM Templates	Chapter 8
	Managing MAC Address Table	Chapters 1–10
	Troubleshoot Err-Disable Recovery	Chapter 10
1.2	Configure and Verify Layer 2 Protocols	
	CDP, LLDP	Chapter 8
	UDLD	Chapter 8
1.3	Configure and Verify VLANs	
	Access Ports	Chapter 3
	VLAN Database	Chapter 3
	Normal, Extended VLAN, Voice VLAN	Chapter 3
1.4	Configure and Verify Trunking	
	VTPv1, VTPv2, VTPv3, VTP Pruning	Chapter 3
	Dot1Q	Chapter 3
	Native VLAN	Chapter 3
	Manual Pruning	Chapter 3
1.5	Configure and Verify EtherChannels	
	LACP, PAgP, Manual	Chapter 3
	Load Balancing	Chapter 3
	EtherChannel Misconfiguration Guard	Chapter 3
1.6	Configure and Verify Spanning Tree	
	PVST+, RPVST+, MST	Chapter 4
	Switch Priority, Port Priority, Path Cost, STP Timers	Chapter 4
	PortFast, BPDU Guard, BPDU Filter	Chapter 4
	Loop Guard and Root Guard	Chapter 4
1.7	Configure and Verify Other LAN Switching Technologies	
	SPAN, RSPAN	Chapter 8

Topic #	Topic	Where Topic Is Covered
1.8	Describe Chassis Virtualization and Aggregation Technologies	
	StackWise	Chapter 9
2.0	Infrastructure Security	
2.1	Configure and Verify Switch Security Features	
	DHCP Snooping	Chapter 10
	IP Source Guard	Chapter 10
	Dynamic ARP Inspection	Chapter 10
	Port Security	Chapter 10
	Private VLAN	Chapter 10
	Storm Control	Chapter 10
2.2	Describe Device Security Using Cisco IOS AAA with TACACS+ and RADIUS	
	AAA with TACACS+ and RADIUS	Chapter 7
	Local Privilege Authorization Fallback	Chapter 7
3.0	Infrastructure Services	
3.1	Configure and Verify First-Hop Redundancy Protocols	
	HSRP	Chapter 6
	VRRP	Chapter 6
	GLBP	Chapter 6

How This Book Is Organized

The chapters and appendix in this book are as follows:

Chapter 1, “Fundamentals Review,” begins with a review of basic switching terminology and previews a couple of terms used in later chapters. The chapter attempts to prevent excessive cross-referencing, because many switching technologies are applicable to all chapters.

Chapter 2, “Network Design Fundamentals,” covers campus network design fundamentals, including campus network structure, Cisco Catalyst switches, and Layer 2 versus multilayer switches. A brief on Catalyst switching hardware functions is also included.

Chapter 3, “Campus Network Architecture,” introduces VLANs, VTP, trunking, and port channeling.

Chapter 4, “Spanning Tree in Depth,” goes into detail about spanning tree and its enhancements that are useful in today’s network.

Chapter 5, “Inter-VLAN Routing,” discusses the fundamentals of routing between VLANs and associated network designs and best practices. In addition, it also discusses Dynamic Host Configuration Protocol (DHCP) services and layer 3 Portchannels.

Chapter 6, “First-Hop Redundancy,” covers the protocols leveraged by Cisco Catalyst switches to support first-hop redundancy, including Hot Standby Router Protocol (HSRP), Gateway Load Balancing Protocol (GLBP), and Virtual Router Redundancy Protocol (VRRP).

Chapter 7, “Network Management,” covers AAA (authentication, authorization, and accounting), Network Time Protocol (NTP), 802.1X, and Simple Network Management Protocol (SNMP) to present a holistic view of network management and Cisco Catalyst device security.

Chapter 8, “Switching Features and Technologies for the Campus Network,” describes how campus networks use advanced features to add resiliency and availability. Network monitoring using Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) is also covered, in addition to the Cisco IOS IP SLA (Service Level Agreement) feature.

Chapter 9, “High Availability,” discusses switch physical redundancy using StackWise, Virtual Switching System (VSS), or redundant supervisors.

Chapter 10, “Campus Network Security,” delves into a plethora of network security features, such as Dynamic Host Configuration Protocol (DHCP) snooping, IP Source Guard, dynamic ARP inspection (DAI), port security, private VLANs, and storm control.

Appendix A, “Answers to Chapter Review Questions,” contains the answers to the review questions that appear at the end of each chapter.

Fundamentals Review

Before journeying into Cisco campus networks and detail technology readouts to prepare for CCNP: Switch, this chapter quickly reviews several topics covered in CCNA and briefly introduces a few topics to ease comprehension of this book. Because each technology covered, such as spanning tree or virtual LANs (VLANs), can exist by itself, the short technology highlights in the chapter reduce cross-referencing of chapters.

If you have a very good understanding of switching terminology and a basic understanding of switching technology, you may want to skip this chapter and begin with Chapter 2, “Network Design Fundamentals.”

This chapter covers the following basic switching topics as a review to CCNA and serves as a teaser for topics covered later in chapter:

- Hubs and switches
- Bridges and switches
- Switches of today
- Broadcast domains
- MAC addresses
- The basic Ethernet frame format
- Basic switching function
- VLANs
- The Spanning Tree Protocol
- Trunking
- Port channels
- Multilayer switching (MLS)

Switching Introduction

The term *LAN switching* is becoming legacy. LAN switching was a popular term to describe LANs built on Cisco Catalyst switches in the 1990s to mid-2000s. In today's networks, LANs have been segmented into distinct functional areas: data centers and campus networks.

This book focuses on campus networks. Campus networks generally take a more conservative approach to architectures, using Cisco Catalyst switches and leveraging traditional Layer 2 and Layer 3 hierarchical designs. Data centers are in a state of evolution, with the focus on applications, dev/ops, and software programmability. These architectures use bleeding-edge technologies such as FabricPath, Dynamic Fabric Allocation (DFA), Application Centric Infrastructure (ACI), and so on.

The remainder of this chapter focuses on a couple of key switching concepts in relation to campus networks that are found throughout this text. Many of these concepts are discussed in more detail in later chapters, but a quick review and definition will help you understand the following chapters. Moreover, because all campus network features are heavily intertwined, it is difficult to present topics in a serial fashion. Definitions in this chapter will ease reading in that manner as well.

Hubs and Switches

Hubs are archaic, and the terminology should be avoided. Even the simplest multiport Ethernet devices for the home are switches.

In review, hubs died off as a product because they are shared-bandwidth devices. Switches introduced dedicated bandwidth. A hub allows multiple devices to be connected to the same network segment. The devices on that segment share the bandwidth with each other. As an example with a 100-Mbps hub, and there are six devices connected to six different ports on the hub, all six devices share the 100 Mbps of bandwidth with each other. A 100-Mbps hub shares 100 Mbps of bandwidth among the connected devices. In terms of the OSI reference model, a hub is considered a Layer 1 (physical layer) device. It hears an electrical signal on the wire and passes it along to the other ports.

A switch allows multiple devices to be connected to the same network, just like a hub does, but this is where the similarity ends. A switch allows each connected device to have dedicated bandwidth instead of shared bandwidth. The bandwidth between the switch and the device is reserved for communication to and from that device alone. Six devices connected to six different ports on a 1-Gbps switch each have 1 Gbps of bandwidth to work with, instead of shared bandwidth with the other devices. A switch can greatly increase the available bandwidth in your network, which can lead to improved network performance. Switches also support additional capabilities beyond what hubs support. Later sub-sections describe some of these features.

Bridges and Switches

A basic switch is considered a Layer 2 device. When we use the word *layer*, we are referring to the seven-layer OSI reference model. A switch does not just pass electrical

signals along, like a hub does; instead, it assembles the signals into a frame (Layer 2), and then decides what to do with the frame. A switch determines what to do with a frame by borrowing an algorithm from a previously common networking device: a transparent bridge. Logically, a switch acts just like a transparent bridge would, but it can handle frames much faster than a transparent bridge could (because of special hardware and architecture). Once a switch decides where the frame should be sent, it passes the frame out the appropriate port (or ports). You can think of a switch as a device creating instantaneous connections between various ports, on a frame-by-frame basis.

Switches of Today

Today's switches have evolved beyond just switching frames. Most modern switches can actually route traffic. In addition, switches can prioritize traffic, support no downtime through redundancy, and provide convergence services around IP telephony and wireless networks.

In summary, to meet evolving network needs of today, Cisco Catalyst switch designs include support for the following industry-leading features beyond the legacy features found in all switches:

- **Application intelligence:** This helps networks recognize many types of applications and secure and prioritize those applications to provide the best user experience.
- **Unified network services:** Combining the best elements of wireless and wired networking allows you to consistently connect to any resource or person with any device. 10 Gigabit Ethernet technology and Power over Ethernet (PoE) technology support new applications and devices.
- **Nonstop communications:** Features such as redundant hardware, and nonstop forwarding and stateful switchover (NSF/SSO) technology support more-reliable connections.
- **Integrated security:** LAN switches provide the first line of defense against internal network attacks and prevent unauthorized intrusion.
- **Operational manageability:** To more easily manage the network, IT staff must be able to remotely configure and monitor network devices from a central location.

Broadcast Domains

In a review from CCNA material, a broadcast domain is a set of network devices that receive broadcast frames originating from any device within the group. Routers typically bound broadcast domains because routers do not forward broadcast frames. VLANs are an example of broadcast domain. Broadcast domains are generally limited to a specific Layer 2 segment that contains a single IP subnet. The next section discusses the addresses used within broadcast domains.

MAC Addresses

MAC addresses are standardized data link layer addresses that are required for every port or device that connects to a LAN. Other devices in the network use these addresses to locate specific ports in the network and to create and update routing tables and data structures. MAC addresses are 6 bytes long and are controlled by the IEEE. MAC addresses are also known as a hardware address, MAC layer address, and physical address.

A MAC address is also applied to virtual devices. Virtual machines on a server may all contain individual MAC addresses. Moreover, most devices have more than one MAC address. A simple example is your laptop; it has both a LAN MAC address and a wireless MAC address. The next section covers the basic frame structure used in Ethernet.

The Basic Ethernet Frame Format

The IEEE 802.3 standard defines a basic data frame format that is required for all MAC implementations, plus several additional optional formats that are used to extend the protocol's basic capability. The basic data frame format contains the following seven fields, as shown in Figure 1-1.

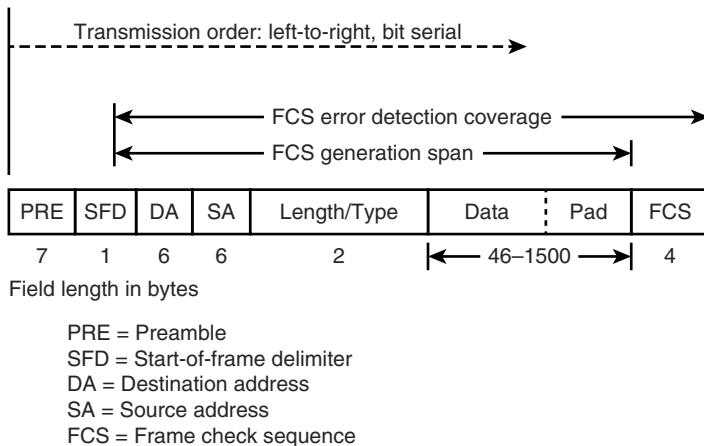


Figure 1-1 *The Basic IEEE 802.3 MAC Data Frame Format*

- **Preamble (PRE):** Consists of 7 bytes. The PRE is an alternating pattern of 1s and 0s that tells receiving stations that a frame is coming, and that provides a means to synchronize the frame-reception portions of receiving physical layers with the incoming bit stream.
- **Start-of-frame delimiter (SOF):** Consists of 1 byte. The SOF is an alternating pattern of 1s and 0s, ending with two consecutive 1 bits, indicating that the next bit is the leftmost bit in the leftmost byte of the destination address.

- **Destination address (DA):** Consists of 6 bytes. The DA field identifies which station(s) should receive the frame. In the first byte of the DA, the 2 least significant bits are used to indicate whether the destination is an individual address or group address (that is, multicast). The first of these 2 bits indicates whether the address is an individual address (indicated by a 0) or a group address (indicated by a 1). The second bit indicates whether the DA is globally administered (indicated by a 0) or locally administered (indicated by a 1). The remaining bits are a uniquely assigned value that identifies a single station, a defined group of stations, or all stations on the network.
- **Source addresses (SA):** Consists of 6 bytes. The SA field identifies the sending station. The SA is always an individual address, and the leftmost bit in the SA field is always 0.
- **Length/Type:** Consists of 2 bytes. This field indicates either the number of MAC-client data bytes that are contained in the data field of the frame, or the frame type ID if the frame is assembled using an optional format. If the Length/Type field value is less than or equal to 1500, the number of LLC bytes in the Data field is equal to the Length/Type field value. If the Length/Type field value is greater than 1536, the frame is an optional type frame, and the Length/Type field value identifies the particular type of frame being sent or received.
- **Data:** Is a sequence of n bytes of any value, where n is less than or equal to 1500. If the length of the Data field is less than 46, the Data field must be extended by adding a filler (a pad) sufficient to bring the Data field length to 46 bytes.

Note that jumbo frames up to 9000 bytes are supported on the current-generation Cisco Catalyst switches.
- **Frame check sequence (FCS):** Consists of 4 bytes. This sequence contains a 32-bit cyclic redundancy check (CRC) value, which is created by the sending MAC and is recalculated by the receiving MAC to check for damaged frames. The FCS is generated over the DA, SA, Length/Type, and Data fields.

Basic Switching Function

When a switch receives a frame, it must decide what to do with that frame. It could ignore the frame, it could pass the frame out one other port, or it could pass the frame out many other ports.

To know what to do with the frame, the switch learns the location of all devices on the segment. This location information is placed in a content addressable memory table (CAM, named for the type of memory used to store these tables). The CAM table shows, for each device, the MAC address of the device, out which port that MAC address can be found, and with which VLAN this port is associated. The switch continually performs this learning process as frames are received into the switch. The CAM table of the switch is continually updated. The next chapter discusses the CAM table in more detail.

This information in the CAM table is used to decide how a received frame is handled. To decide where to send a frame, the switch looks at the destination MAC address in a received frame and looks up that destination MAC address in the CAM table. The CAM table shows the port that the frame must be sent out for that frame to reach the specified destination MAC address. In brief, the basic switching function at Layer 2 adheres to these rules for determining forwarding responsibility:

- If the destination MAC address is found in the CAM table, the switch sends the frame out the port that is associated with that destination MAC address in the CAM table. This process is called *forwarding*.
- If the associated port to send the frame out is the same port that the frame originally came in on, there is no need to send the frame back out that same port, and the frame is ignored. This process is called *filtering*.
- If the destination MAC address is not in the CAM table (that is, unknown unicast), the switch sends the frame out all other ports that are in the same VLAN as the received frame. This is called *flooding*. It does not flood the frame out the same port on which the frame was received.
- If the destination MAC address of the received frame is the broadcast address (FFFF.FFFF.FFFF), the frame is sent out all ports that are in the same VLAN as the received frame. This is also called *flooding*. The only exception is the frame is not sent out the same port on which the frame was received.

The next section introduces a widely popular feature leveraged by Cisco Catalyst switches and Nexus switches to segment groups of ports into their own LAN segments.

VLANs

Because the switch decides on a frame-by-frame basis which ports exchange data, it is a natural extension to put logic inside the switch to allow it to choose ports for special groupings. This grouping of ports is called a *virtual local-area network* (VLAN). The switch makes sure that traffic from one group of ports never gets sent to other groups of ports (which would be routing). These port groups (VLANs) can each be considered an individual LAN segment.

VLANs are also described as broadcast domains. This is because of the transparent bridging algorithm, which says that broadcast packets (packets destined for the *all devices* address) be sent out all ports that are in the same group (that is, in the same VLAN). All ports that are in the same VLAN are also in the same broadcast domain.

The next section introduces the legacy spanning tree technology used to build Layer 2 domains.

The Spanning Tree Protocol

As discussed previously, the switch forwarding algorithm floods unknown and broadcast frames out of all the ports that are in the same VLAN as the received frame. This causes

a potential problem. If the network devices that run this algorithm are connected together in a physical loop, flooded frames (like broadcasts) are passed from switch to switch, around and around the loop, forever. Depending on the physical connections involved, the frames can actually multiply exponentially because of the flooding algorithm, which can cause serious network problems.

There is a benefit to a physical loop in your network: It can provide redundancy. If one link fails, there is still another way for the traffic to reach its destination. To allow the benefits derived from redundancy, without breaking the network because of flooding, a protocol called the *Spanning Tree Protocol* (STP) was created. Spanning tree was standardized in the IEEE 802.1D specification.

The purpose of STP is to identify and temporarily block the loops in a network segment or VLAN. The switches run STP, which involves electing a root bridge or switch. The other switches measure their distance from the root switch. If there is more than one way to get to the root switch, there is a loop. The switches follow the algorithm to determine which ports must be blocked to break the loop. STP is dynamic; if a link in the segment fails, ports that were originally blocking can possibly be changed to forwarding mode.

Spanning tree is covered in more detail later in this book. The next section covers how to pass multiple VLANs on a single port.

Trunking

Trunking is a mechanism that is most often used to allow multiple VLANs to function independently across multiple switches. Routers and servers can use trunking, as well, which allows them to live simultaneously on multiple VLANs. If your network only has one VLAN in it, you might never need trunking; but if your network has more than one VLAN, you probably want to take advantage of the benefits of trunking.

A port on a switch normally belongs to only one VLAN; any traffic received or sent on this port is assumed to belong to the configured VLAN. A trunk port, however, is a port that can be configured to send and receive traffic for many VLANs. It accomplishes this when it attaches VLAN information to each frame, a process called *tagging* the frame. Also, trunking must be active on both sides of the link; the other side must expect frames that include VLAN information for proper communication to occur. As with all the section briefs in this chapter, more information is found later in this book.

Port Channels

Utilizing port channels (EtherChannels) is a technique that is used when you have multiple connections to the same device. Rather than each link functioning independently, port channels group the ports together to work as one unit. Port channels distribute traffic across all the links and provide redundancy if one or more links fail. Port channel settings must be the same on both sides of the links involved in the channel. Normally, spanning tree would block all of these parallel connections between devices because

they are loops, but port channels run *underneath* spanning tree, so that spanning tree thinks all the ports within a given port channel are only a single port. Later chapters discuss port channels in more detail.

Multilayer Switching

Multilayer switching (MLS) is the ability of a switch to forward frames based on information in the Layer 3 and sometimes Layer 4 header. Almost all Cisco Catalyst switches model 3500 or later support MLS. MLS is becoming a legacy term due to the wide support. The most important aspect to MLS is recognizing that switches can route or switch frames at wire-rate speeds using specialized hardware. This effectively bundles the routing function into the switch and is specifically useful for routing between VLANs in the core of the network. The next chapter discusses this capability in more detail.

Summary

This chapter briefly reviewed several common technology topics pertaining to switching. The remaining chapters of this book cover these topics and other (newer) switching technology related to security.

Network Design Fundamentals

Every time you go to an office to work or go to class at school, college, or university, you will use a campus network to access critical applications, tools, the Internet, and so on over wired or wireless connections. Often, you may even gain access by using a portable device such as an Apple iPhone connected on a corporate Wi-Fi to reach applications such as e-mail, calendaring, or instant messaging over a campus network. Therefore, the persons responsible for building this network need to deploy sound fundamentals and design principles for the campus networks to function adequately and provide the necessary stability, scalability, and resiliency necessary to sustain interconnectivity with a 100 percent uptime.

This chapter begins the journey of exploring campus network design fundamentals by focusing on a few core concepts around network design and structure and a few details about the architecture of Cisco switches. This is useful knowledge when designing and building campus networks. Specifically, this chapter focuses on the following two high-level topics:

- Campus network structure
- Introduction to Cisco switches and their associated architecture

Campus Network Structure

A campus network describes the portion of an enterprise infrastructure that interconnects end devices such as computers, laptops, and wireless access points to services such as intranet resources or the Internet. Intranet resources may be company web pages, call center applications, file and print services, and almost anything end users connect to from their computer.

In different terms, the campus network provides for connectivity to company applications and tools that reside in a data center for end users. Originally, prior to around 2005, the term *campus network* and its architectures were relevant for application server farms and computing infrastructure as well. Today, the infrastructure that interconnects

server farms, application servers, and computing nodes are clearly distinguished from campus networks and referred to as *data centers*.

Over the past few years, data center architectures have become more complex and require sophistication not required in the campus network due to high-availability, low-latency, and high-performance requirements. Therefore, data centers may use bleeding-edge technologies that are not found in the campus network, such as FabricPath, VXLAN, and Application Centric Infrastructure (ACI). For the purpose of CCNP Switch at the time of this writing, these technologies, as well as data center architectures, are out of scope. Nevertheless, we will point out some of the differences as to avoid any confusion with campus network fundamentals.

The next subsection describes the hierarchical network design with the following subsections breaking down the components of the hierarchical design in detail.

Hierarchical Network Design

A flat enterprise campus network is where all PCs, servers, and printers are connected to each other using Layer 2 switches. A flat network does not use subnets for any design purposes. In addition, all devices on this subnet are in the same broadcast domain, and broadcasts will be flooded to all attached network devices. Because a broadcast packet received by an end device, such as tablet or PC, uses compute and I/O resources, broadcasts will waste available bandwidth and resources. In a network size of ten devices on the same flat network, this is not a significant issue; however, in a network of thousands of devices, this is a significant waste of resources and bandwidth (see Figure 2-1).




Figure 2-1 Flat Versus Hierarchical Network Design

As a result of these broadcast issues and many other limitations, flat networks do not scale to meet the needs of most enterprise networks or of many small and medium-size businesses. To address the sizing needs of most campus networks, a hierarchical model is used. Figure 2-2 illustrates, at a high level, a hierarchical view of campus network design versus a flat network.




Figure 2-2 *The Hierarchical Model*

Hierarchical models for network design allow you to design any networks in layers. To understand the importance of layering, consider the OSI reference model, which is a layered model for understanding and implementing computer communications. By using layers, the OSI model simplifies the task that is required for two computers to communicate. Leveraging the hierarchical model also simplifies campus network design by allowing focus at different layers that build on each other.

Referring to Figure 2-2, the layers of the hierarchical model are divided into specific functions categorized as core, distribution, and access layers. This categorization provides for modular and flexible design, with the ability to grow and scale the design without major modifications or reworks.

For example, adding a new wing to your office building may be as simple as adding a new distribution layer with an access layer while adding capacity to the core layer. The existing design will stay intact, and only the additions are needed. Aside from the simple physical additions, configuration of the switches and routes is relatively simple because most of the configuration principles around hierarchy were in place during the original design.

By definition, the access, distribution, and core layer adhere to the following characteristics:

- **Access layer:** The access layer is used to grant the user access to network applications and functions. In a campus network, the access layer generally incorporates

switched LAN devices with ports that provide connectivity to workstations, IP phones, access points, and printers. In a WAN environment, the access layer for teleworkers or remote sites may provide access to the corporate network across WAN technologies.

- **Distribution layer:** The distribution layer aggregates the access layer switches wiring closets, floors, or other physical domain by leveraging module or Layer 3 switches. Similarly, a distribution layer may aggregate the WAN connections at the edge of the campus and provides policy-based connectivity.
- **Core layer (also referred to as the backbone):** The core layer is a high-speed backbone, which is designed to switch packets as fast as possible. In most campus networks, the core layer has routing capabilities, which are discussed in later chapters of this book. Because the core is critical for connectivity, it must provide a high level of availability and adapt to changes quickly. It also provides for dynamic scalability to accommodate growth and fast convergence in the event of a failure.

The next subsections of this chapter describe the access layer, distribution layer, and core layer in more detail.

Access Layer

The access layer, as illustrated in Figure 2-3, describes the logical grouping of the switches that interconnect end devices such as PCs, printers, cameras, and so on. It is also the place where devices that extend the network out one more level are attached. Two such prime examples are IP phones and wireless APs, both of which extend the connectivity out one more layer from the actual campus access switch.




Figure 2-3 *Access Layer*

The wide variety of possible types of devices that can connect and the various services and dynamic configuration mechanisms that are necessary make the access layer one of the most capable parts of the campus network. These capabilities are as follows:

- **High availability:** The access layer supports high availability via default gateway redundancy using dual connections from access switches to redundant distribution layer switches when there is no routing in the access layer. This mechanism

behind default gateway redundancy is referred to as *first-hop redundancy protocol* (FHRP). FHRP is discussed in more detail in later chapters of this book.

- **Convergence:** The access layer generally supports inline Power over Ethernet (PoE) for IP telephony, thin clients, and wireless access points (APs). PoE allows customers to easily place IP phones and wireless APs in strategic locations without the need to run power. In addition, the access layers allow support for converged features that enable optimal software configuration of IP phones and wireless APs, as well. These features are discussed in later chapters.
- **Security:** The access layer also provides services for additional security against unauthorized access to the network by using tools such as port security, quality of service (QoS), Dynamic Host Configuration Protocol (DHCP) snooping, dynamic ARP inspection (DAI), and IP Source Guard. These security features are discussed in more detail in later chapters of this book.

The next subsection discusses the upstream layer from the access layer, the distribution layer.

Distribution Layer

The distribution layer in the campus design has a unique role in which it acts as a services and control boundary between the access layer and the core. Both the access layer and the core are essentially dedicated special-purpose layers. The access layer is dedicated to meeting the functions of end-device connectivity, and the core layer is dedicated to providing nonstop connectivity across the entire campus network. The distribution layer, in contrast, serves multiple purposes. Figure 2-4 references the distribution layer.




Figure 2-4 *Distribution Layer*

Availability, fast path recovery, load balancing, and QoS are all important considerations at the distribution layer. Generally, high availability is provided through Layer 3 redundant paths from the distribution layer to the core, and either Layer 2 or Layer 3 redundant paths from the access layer to the distribution layer. Keep in mind that Layer 3 equal-cost load sharing allows both uplinks from the distribution to the core layer to be used for traffic in a variety of load-balancing methods discussed later in this chapter.

Note Equal-cost multipathing (ECMP) is another term used to describe equal-cost load sharing. However, the term ECMP is typically used with respect to data center architectures and not campus architectures. This book uses both terms, equal-cost load sharing and ECMP, interchangeably.

With a Layer 2 design in the access layer, the distribution layer generally serves as a routing boundary between the access and core layer by terminating VLANs. The distribution layer often represents a redistribution point between routing domains or the demarcation between static and dynamic routing protocols. The distribution layer may perform tasks such as controlled routing decision making and filtering to implement policy-based connectivity, security, and QoS. These features allow for tighter control of traffic through the campus network.

To improve routing protocol performance further, the distribution layer is generally designed to summarize routes from the access layer. If Layer 3 routing is extended to the access layer, the distribution layer generally offers a default route to access layer switching while leveraging dynamic routing protocols when communicating with core routers.

In addition, the distribution layer optionally provides default gateway redundancy by using a first-hop routing protocol (FHRP) such as Host Standby Routing Protocol (HSRP), Gateway Load Balancing Protocol (GLBP), or Virtual Router Redundancy Protocol (VRRP). FHRPs provide redundancy and high availability for the first-hop default gateway of devices connected downstream on the access layer. In designs that leverage Layer 3 routing in the access layer, FHRP might not be applicable or may require a different design.

In summary, the distribution layer performs the following functions when Layer 3 routing is not configured in the access layer:

- Provides high availability and equal-cost load sharing by interconnecting the core and access layer via at least dual paths
- Generally terminates a Layer 2 domain of a VLAN
- Routes traffic from terminated VLANs to other VLANs and to the core
- Summarizes access layer routes
- Implements policy-based connectivity such as traffic filtering, QoS, and security
- Provides for an FHRP

Core Layer (Backbone)

The core layer, as illustrated in Figure 2-5, is the backbone for campus connectivity, and is the aggregation point for the other layers and modules of an enterprise network. The core must provide a high level of redundancy and adapt to changes quickly.




Figure 2-5 Core Layer

From a design point-of-view, the campus core is in some ways the simplest yet most critical part of the campus. It provides a limited set of services and is designed to be highly available and requires 100 percent uptime. In large enterprises, the core of the network must operate as a nonstop, always-available service. The key design objectives for the campus core are based on providing the appropriate level of redundancy to allow for near-immediate data-flow recovery in the event of the failure of any component (switch, supervisor, line card, or fiber interconnect, power, and so on). The network design must also permit the occasional, but necessary, hardware and software upgrade or change to be made without disrupting any network applications. The core of the network should not implement any complex policy services, nor should it have any directly attached user or server connections. The core should also have the minimal control plane configuration that is combined with highly available devices that are configured with the correct amount of physical redundancy to provide for this nonstop service capability. Figure 2-6 illustrates a large campus network interconnected by the core layer (campus backbone) to the data center.

From an enterprise architecture point-of-view, the campus core is the backbone that binds together all the elements of the campus architecture to include the WAN, the data center, and so on. In other words, the core layer is the part of the network that provides for connectivity between end devices, computing, and data storage services that are located within the data center, in addition to other areas and services within the network.

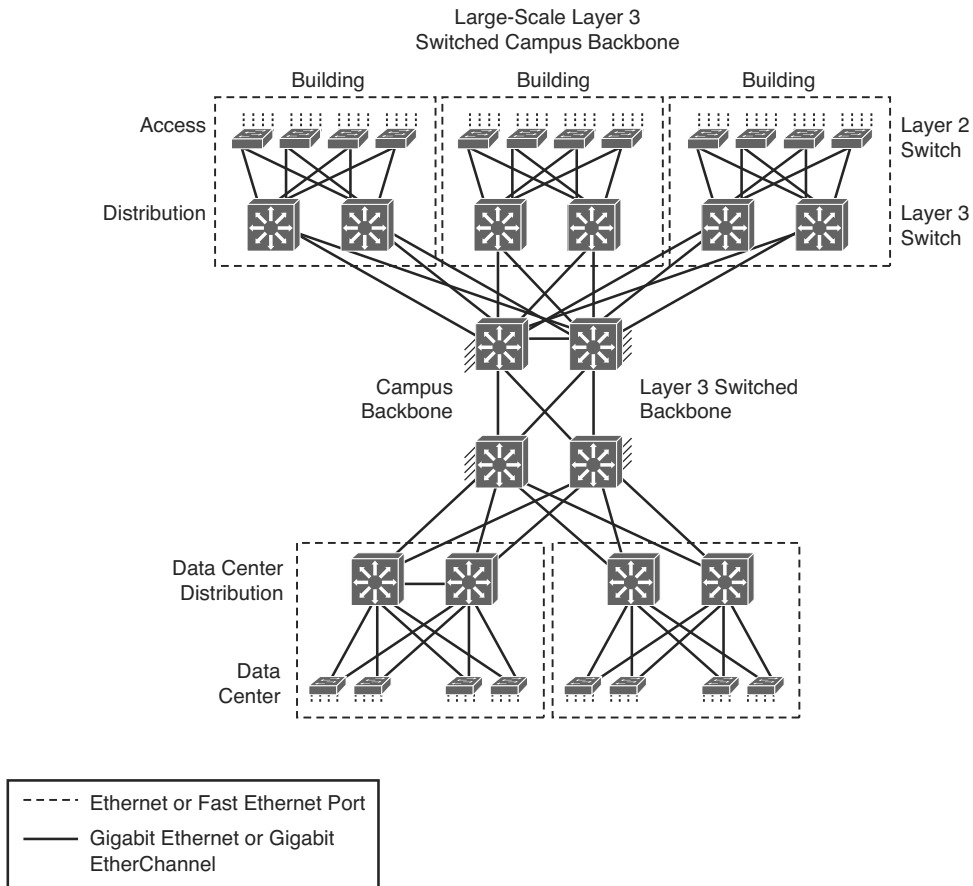


Figure 2-6 *Large Campus Network*

Figure 2-7 illustrates an example of the core layer interconnected with other parts of the enterprise network. In this example, the core layer interconnects with a data center and edge distribution module to interconnect WAN, remote access, and the Internet. The network module operates out of band from the network but is still a critical component.

In summary, the core layer is described as follows:

- Aggregates the campus networks and provides interconnectivity to the data center, the WAN, and other remote networks
- Requires high availability, resiliency, and the ability to make software and hardware upgrades without interruption
- Designed without direct connectivity to servers, PCs, access points, and so on
- Requires core routing capability

- Architected for future growth and scalability
- Leverages Cisco platforms that support hardware redundancy such as the Catalyst 4500 and the Catalyst 6800

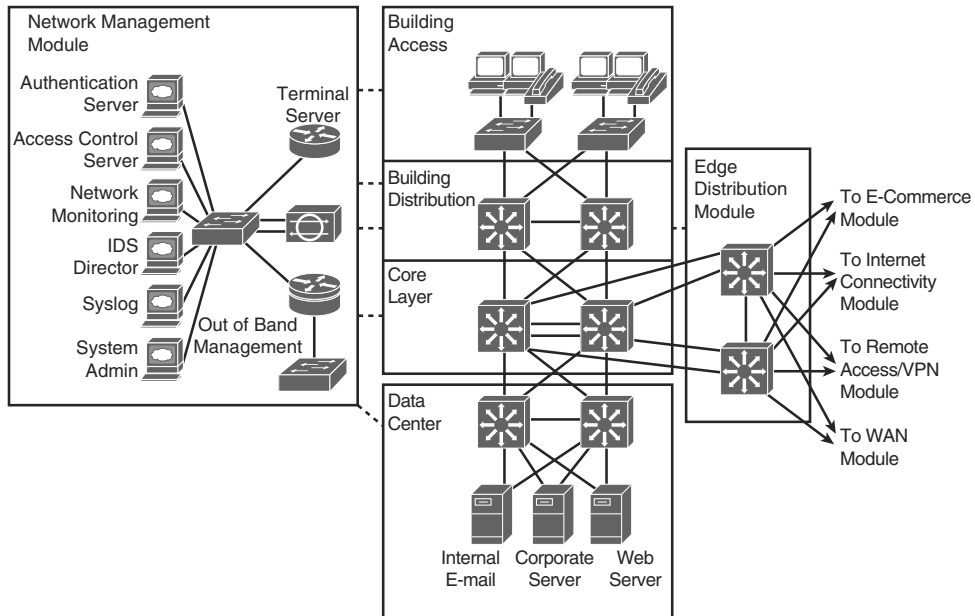


Figure 2-7 Core Layer Interconnecting with the Enterprise Network

Layer 3 in the Access Layer

As switch products become more commoditized, the cost of Layer 3 switches has diminished significantly. Because of the reduced cost and a few inherent benefits, Layer 3 switching in the access layer has become more common over typical Layer 2 switching in the access layer. Using Layer 3 switching or traditional Layer 2 switching in the access layer has benefits and drawbacks. Figure 2-8 illustrates the comparison of Layer 2 from the access layer to the distribution layer with Layer 3 from the access layer to the distribution layer.

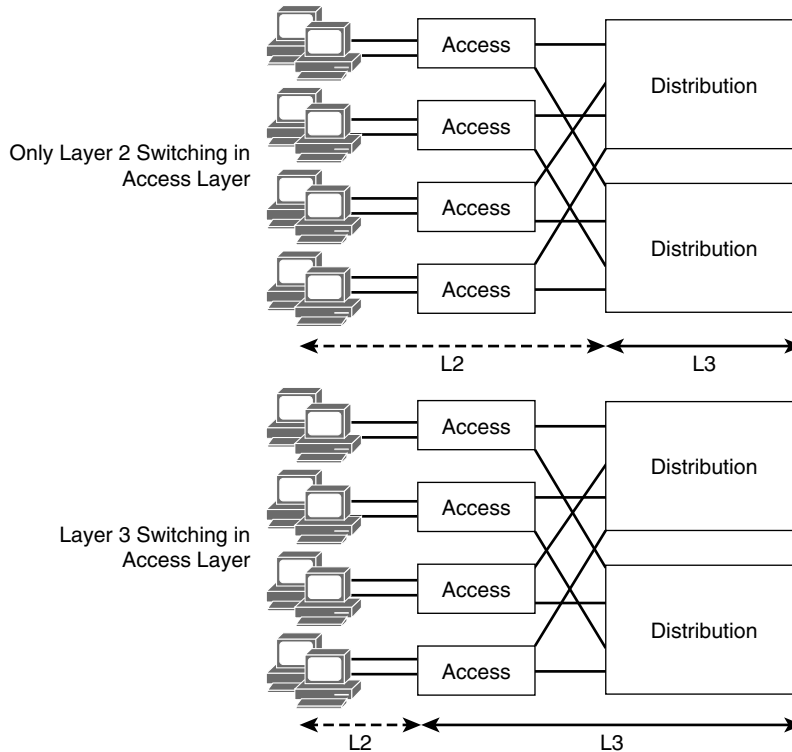


Figure 2-8 *Layer 3 in the Access Layer*

As discussed in later chapters, deploying a Layer 2 switching design in the access layer may result in suboptimal usage of links between the access and distribution layer. In addition, this method does not scale as well in very large numbers because of the size of the Layer 2 domain.

Using a design that leverages Layer 3 switching to the access layer VLANs scales better than Layer 2 switching designs because VLANs get terminated on the access layer devices. Specifically, the links between the distribution and access layer switches are routed links; all access and distribution devices would participate in the routing scheme.

The Layer 2-only access design is a traditional, slightly cheaper solution, but it suffers from optimal use of links between access and distribution due to spanning tree. Layer 3 designs introduce the challenge of how to separate traffic. (For example, guest traffic should stay separated from intranet traffic.) Layer 3 designs also require careful planning with respect to IP addressing. A VLAN on one Layer 3 access device cannot be on another access layer switch in a different part of your network because each VLAN is globally significant. Traditionally, mobility of devices is limited in the campus network of the enterprise in Layer 3 access layer networks, without using an advanced mobility networking features.

Note Modern technologies such as Dynamic Fabric Allocation (DFA) and ACI enable simplified mobility of devices while maintaining a scalable and resilient architecture. At the time of this writing, DFA and ACI were data center only technologies and beyond the scope of CCNP.

In summary, campus networks with Layer 3 in the access layer are becoming more popular. Moreover, next-generation architectures will alleviate the biggest problem with Layer 3 routing in the access layer: mobility.

The next subsection of this chapter applies the hierarchical model to an enterprise architecture.

The Cisco Enterprise Campus Architecture

The Cisco enterprise campus architecture refers to the traditional hierarchical campus network applied to the network design, as illustrated in Figure 2-9.

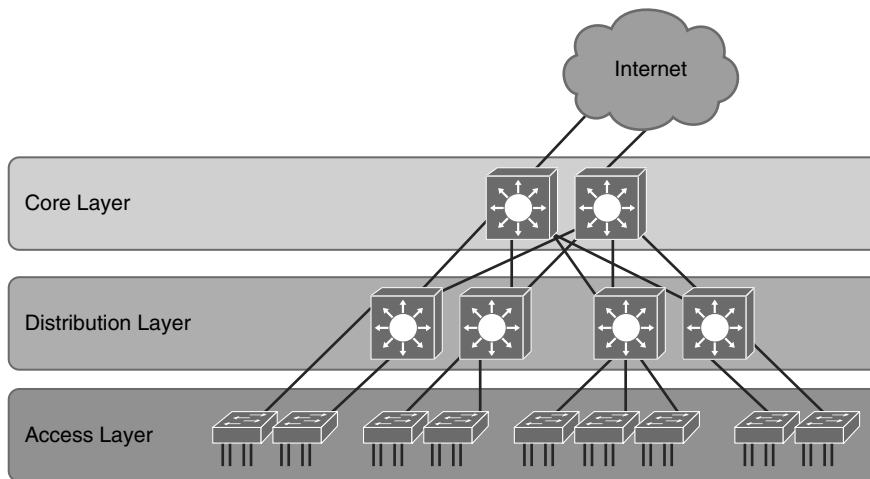


Figure 2-9 *Cisco Enterprise Campus Network*

The Cisco enterprise campus architecture divides the enterprise network into physical, logical, and functional areas while leveraging the hierarchical design. These areas allow network designers and engineers to associate specific network functionality on equipment that is based on its placement and function in the model.

Note that although the tiers do have specific roles in the design, no absolute rules apply to how a campus network is physically built. Although it is true that many campus networks are constructed by three physical tiers of switches, this is not a strict requirement. In a smaller campus, the network might have two tiers of switches in which the core and distribution elements are combined in one physical switch: a collapsed distribution and core. However, a network may have four or more physical tiers of switches because the scale, wiring plant, or physical geography of the network might require that the core be extended.

The hierarchy of the network often defines the physical topology of the switches, but they are not the same thing. The key principle of the hierarchical design is that each element in the hierarchy has a specific set of functions and services that it offers and a specific role to play in the design.

In reference to CCNP Switch, the access layer, the distribution layer, and core layer may be referred to as the *building access layer*, the *building distribution layer*, and the *building core layer*. The term *building* implies but does not limit the context of layers as physical buildings. As mentioned previously, the physical demarcation does not have to be a building; it can be a floor, group of floors, wiring closets, and so on. This book will solely use the terms *access layer*, *distribution layer*, and *core layer* for simplicity.

In summary, network architects build Cisco enterprise campus networks by leveraging the hierarchical model and dividing the layers by some physical or logical barrier. Although campus network designs go much further beyond the basic structure, the key takeaway of this section is that the access, distribution, and core layers are applied to either physical or logical barriers.

The Need for a Core Layer

When first studying campus network design, persons often question the need for a core layer. In a campus network contained with a few buildings or a similar physical infrastructure, collapsing the core into the distribution layer switches may save on initial cost because an entire layer of switches is not needed. Figure 2-10 shows a network design example where the core layer has been collapsed into the distribution layer by fully meshing the four distinct physical buildings.

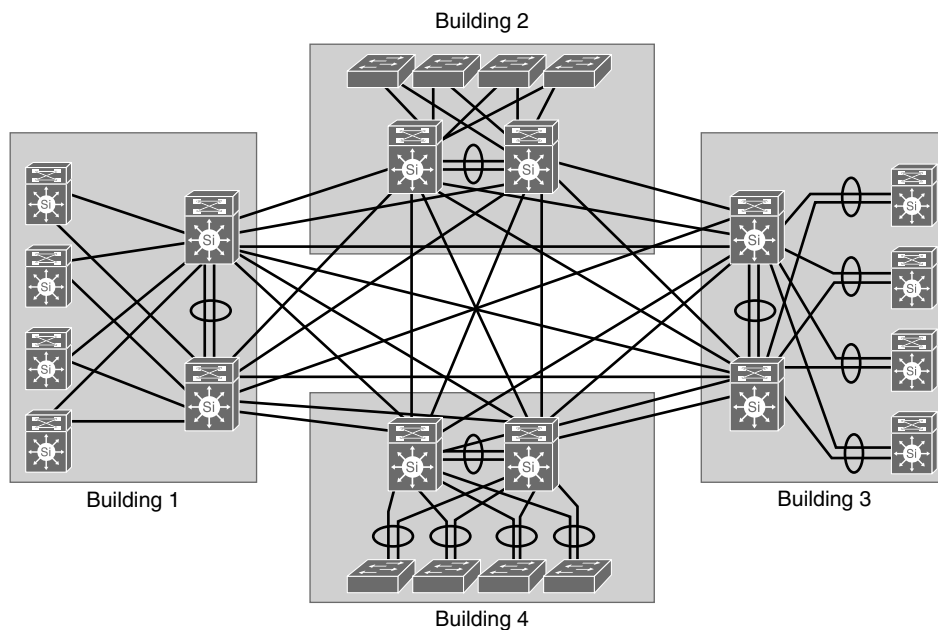


Figure 2-10 *Collapsed Core Design*

Despite a possible lower cost to the initial build, this design is difficult to scale. In addition, cabling requirements increase dramatically with each new building because of the need for full-mesh connectivity to all the distribution switches. The routing complexity also increases as new buildings are added because additional routing peers are needed.

With regard to Figure 2-10, the distribution module in the second building of two interconnected switches requires four additional links for full-mesh connectivity to the first module. A third distribution module to support the third building would require 8 additional links to support the connections to all the distribution switches, or a total of 12 links. A fourth module supporting the fourth building would require 12 new links for a total of 24 links between the distribution switches.

As illustrated in Figure 2-11, having a dedicated core layer allows the campus to accommodate growth without requiring full-mesh connectivity between the distribution layers. This is particularly important as the size of the campus grows either in number of distribution blocks, geographical area, or complexity. In a larger, more complex campus, the core provides the capacity and scaling capability for the campus as a whole and may house additional services such as security features.




Figure 2-11 *Scaling with a Core Layer*

The question of when a separate physical core is necessary depends on multiple factors. The ability of a distinct core to allow the campus network to solve physical design challenges is important. However, remember that a key purpose of having a distinct campus core is to provide scalability and to minimize the risk from (and simplify) moves, adds,

and changes in the campus network. In general, a network that requires routine configuration changes to the core devices does not yet have the appropriate degree of design modularization. As the network increases in size or complexity and changes begin to affect the core devices, it often points out design reasons for physically separating the core and distribution functions into different physical devices.

In brief, although design networks without a core layer may work at small scale, medium-sized to enterprise-sized networks, they require a core layer for design modularization and scalability.

In conclusion of the hierarchical model presented in this section, despite its age, the hierarchical model is still relevant to campus network designs. For review, the layers are described as follows:

- The access layer connects end devices such as PCs, access points, printers, and so on to the network.
- The distribution layer has multiple roles, but primarily aggregates the multiple access layers. The distribution may terminate VLANs in Layer 2 to the access layer designs or provide routing downstream to the access layer with Layer 3 to the access layer designs.

The next section delves into a major building block of the campus network: the Cisco switch itself.

Types of Cisco Switches

Switches are the fundamental interconnect component of the campus network. Cisco offers a variety of switches specifically designed for different functions. At the time of this writing, Cisco designs the Catalyst switches for campus networks and Nexus switches for data centers. In the context of CCNP, this book focuses mostly on Catalyst switches.

Figure 2-12 illustrates the current recommended Catalyst switches. However, in the competitive campus switch marketplace, Cisco continuously updates the Catalyst switches with new capabilities, higher performance, higher density, and lower cost.

Interesting enough, the Catalyst 6500 was not detailed in Figure 2-12. Despite its extremely long life cycle, Cisco marketing has finally shifted focus to the Catalyst 6800. For a large number of you reading this book, you have likely come across the Catalyst 6500 at some point in your career.

Cisco offers two types of network switches: fixed configuration and modular switches. With fixed configuration switches, you cannot swap or add another module, like you can with a modular switch. In enterprise access layers, you will find fixed configuration switches, like the Cisco Catalyst, 2960-X series. It offers a wide range of deployments.

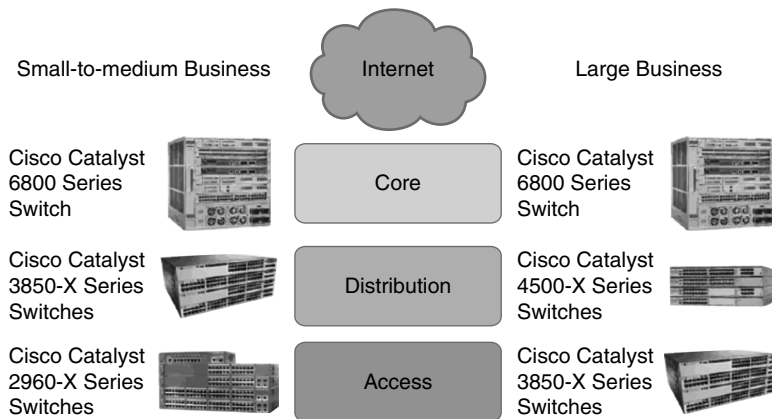


Figure 2-12 Cisco Catalyst Switches

In the enterprise distribution layer, you will find either fixed or modular switches depending on campus network requirements. An example of a modular switch that can be found in the distribution layer is the Cisco Catalyst 3850-X series. This series of switches allows you to select different network modules (Ethernet or fiber optic) and redundant power supply modules. In small businesses without a distribution layer, the 3850-X can be found in the core layer. In large enterprise networks, you might find 3850-X in the access layer in cases where high redundancy and full Layer 3 functionality at the access layer are requirements.

In the enterprise core layer, you will often find modular switches such as the Cisco Catalyst 6500 or the Catalyst 6800 series. With the 6800 switch, nearly every component, including the route processing/supervisor module and Ethernet models to power supplies) is individually installed in a chassis. This individualization allows for customization and high-availability options when necessary.

If you have a network where there is a lot of traffic, you have the option to leverage the Cisco Catalyst 4500-X series switches into the distribution layer. The Catalyst 4500-X supports supervisor/route process redundancy and supports 10 Gigabit Ethernet.

All switches within the 2960-X, 3850-X, 4500-X, and 6800-X series are managed. This means that you can configure an IP address on the device. By having a management IP address, you can connect to the device using Secure Shell (SSH) or Telnet and change device settings. An unmanaged switch is only appropriate for a home or very small business environment. It is highly recommended not to use an unmanaged switch in any campus network.

This section just described a few examples of Cisco switches and their placement in the network. For more information, go to <http://www.cisco.com/c/en/us/products/switches/index.html>.

The next section compares Layer 2 and Layer 3 (multilayer switches).

Comparing Layer 2 and Multilayer Switches

A Layer 2 Ethernet switch operates at the Data Link Layer of the OSI model. These types of switches make decisions about forwarding frames based on the destination MAC addresses found within the frame.

Recalling basic networking: A switch collision domain is only port to port because each switch port and its associated end device is its own collision domain. Because there is no contention on the media, all hosts can operate in full-duplex mode, which means that they can receive and transmit data at the same time. The concept of half duplex is legacy and applies only to hubs and older 10/100-Mbps switches, because 1 Gbps operates by default at full duplex.

When a switch receives in store-n-forward mode, the frame is checked for errors, and frames with a valid cyclic redundancy check (CRC) are regenerated and transmitted. Some models of switches, mostly Nexus switches, opt to switch frames based only on reading the Layer 2 information and bypassing the CRC check. This bypass, referred to as cut-through switching, lowers the latency of the frame transmission as the entire frame is not stored before transmission to another port. Lower switching latency is beneficial for low-latency applications such as algorithm trading programs found in the data center. The assumption is that the end device network interface card (NIC) or an upper-level protocol will eventually discard the bad frame. Most Catalyst switches are store-n-forward.

MAC Address Forwarding

To figure out where a frame must be sent, the switch will look up its MAC address table. This information can be told to the switch or it can learn it automatically. The switch listens to incoming frames and checks the source MAC addresses. If the address is not in the table already, the MAC address, switch port, and VLAN will then get recorded in the forwarding table. The forwarding table is also called the *CAM table*.

What happens if the destination MAC address of the frame is unknown to the switch? The switch then forwards the frame through all ports within a VLAN except the port the frame was received on. This is known as *unknown unicast flooding*. Broadcast and multicast traffic is destined for multiple destinations, so it will get flooded by default.

Referring to Figure 2-13, in the first example, the switch receives a frame on port 1. The destination MAC address for the frame is 0000.0000.5555. The switch will look up its forwarding table and figure out that MAC address 0000.0000.5555 is recorded on port 5. The switch will then forward the frame through port 5.

In the second example, the switch receives a broadcast frame on port 1. The switch will forward the frame through all ports that are within the same VLAN except port 1. The frame was received on port 1, which is in VLAN 1; therefore, the frame is forwarded through all ports on the switch that belong to VLAN 1 (all ports except port 3).

The next subsection discusses Layer 2 switch operation from a mechanics point of view.

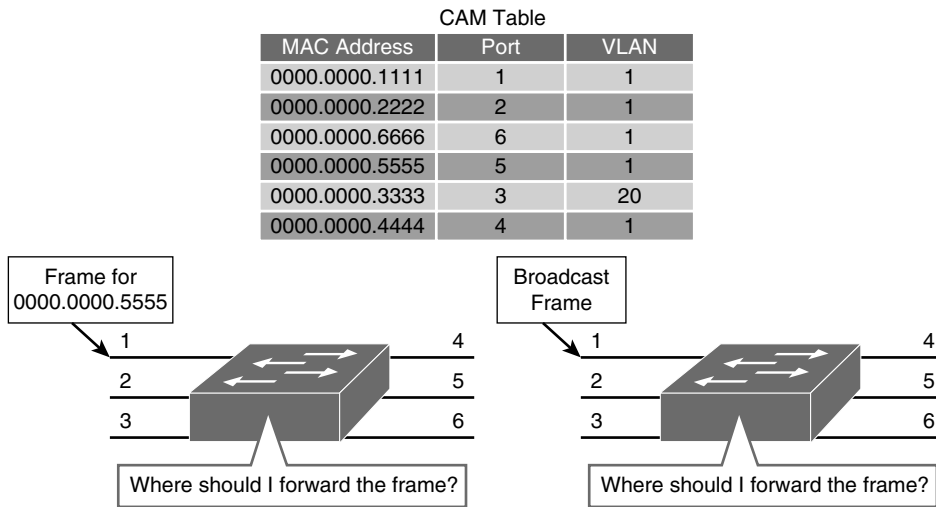


Figure 2-13 Layer 2 Switching Operation: MAC Address Forwarding

Layer 2 Switch Operation

When a switch receives a frame, it places the frame into an ingress queue. A port can have multiple ingress queues, and typically these queues are used to service frames differently (for example, apply quality of service [QoS]). From a simplified viewpoint, when the switch selects a frame from a queue to transmit, the switches need to answer a few questions:

- Where should the frame be forwarded?
- Are there restrictions preventing the forwarding of the frame?
- Is there any prioritization or marking that needs to be applied to the frame?

Decisions about these three questions are answered, respectively, as illustrated in Figure 2-14 and described in the list that follows.

- **Layer 2 forwarding table:** The Layer 2 forwarding table, also called the *MAC table*, contains information about where to forward the frame. Specifically, it contains MAC addresses and destination ports. The switches reference the destination MAC address of the incoming frame in the MAC table and forward the frames to the destination ports specified in the table. If the MAC address is not found, the frame is flooded through all ports in the same VLAN.
- **ACLs:** Access control lists (ACLs) do not only apply to routers. Switches can also apply ACLs based on MAC and IP addresses. Generally only higher-end switches support ACLs based on both MAC and IP addresses, whereas Layer 2 switches support ACLs only with MAC addresses.
- **QoS:** Incoming frames can be classified according to QoS parameters. Traffic can then be marked, prioritized, or rate-limited.

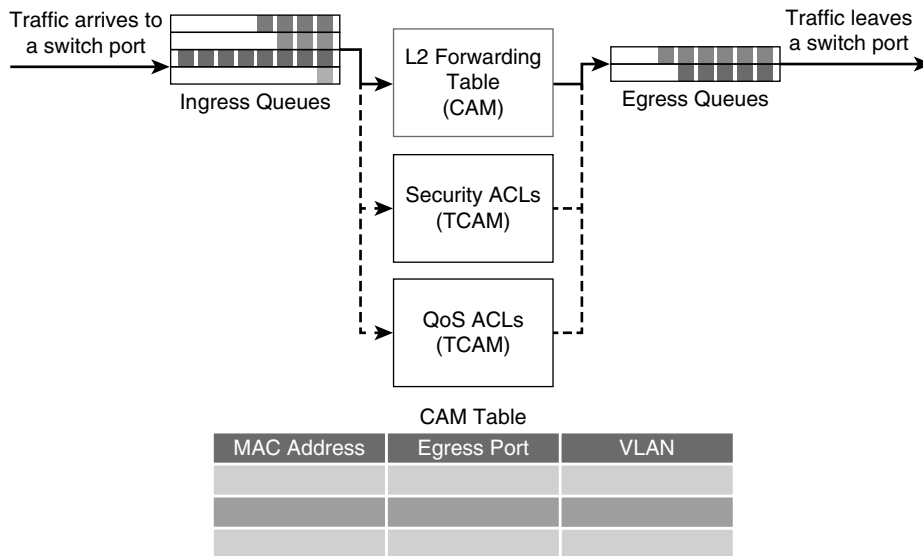


Figure 2-14 *Layer 2 Switch Operation: Mechanics*

Switches use specialized hardware to house the MAC table, ACL lookup data, and QoS lookup data. For the MAC table, switches use content-addressable memory (CAM), whereas the ACL and QoS tables are housed in ternary content-addressable memory (TCAM). Both CAM and TCAM are extremely fast access and allow for line-rate switching performance. CAM supports only two results: 0 or 1. Therefore, CAM is useful for Layer 2 forwarding tables.

TCAM provides three results: 0, 1, and don't care. TCAM is most useful for building tables for searching on longest matches, such as IP routing tables organized by IP prefixes. The TCAM table stores ACL, QoS, and other information generally associated with upper-layer processing. As a result of using TCAM, applying ACLs does not affect the performance of the switch.

This section only touches on the details and implementation of CAM and TCAM needed for the CCNP certification. For a more detailed description, review the following support document at Cisco.com:

<https://supportforums.cisco.com/document/60831/cam-vs-tcam>

<https://www.pagiantzis.com/cam/camintro>

The next subsection discusses Layer 3 (multilayer) switch operation in more detail.

Layer 3 (Multilayer) Switch Operation

Multilayer switches not only perform Layer 2 switching but also forward frames based on Layer 3 and 4 information. Multilayer switches not only combine the functions of a switch and a router but also add a flow cache component.

Multilayer switches apply the same behavior as Layer 2 switches but add an additional parallel lookup for how to route a packet, as illustrated in Figure 2-15.

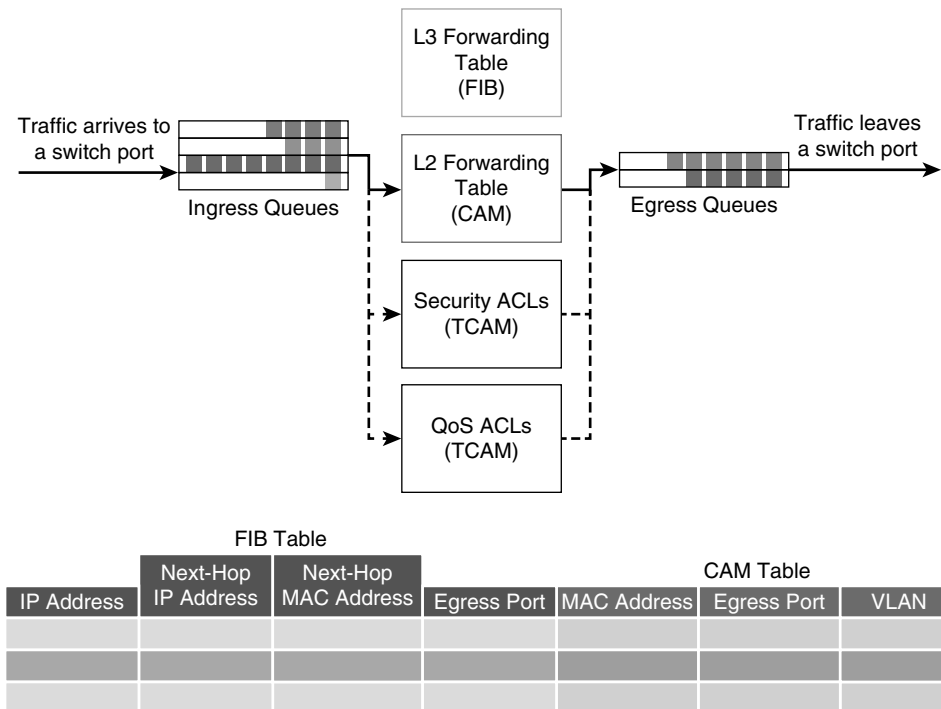


Figure 2-15 *Multilayer Switch Operation*

The associated table for Layer 3 lookups is called a *FIB table*. The FIB table contains not only egress ports and VLAN information but also MAC rewrite information. The ACL and QoS parallel lookups happen the same as Layer 2 switches, except there may be additional support for Layer 3 ACLs and QoS prioritization.

For example, a Layer 2 switch may only be able to apply to rate-limiting frames based on source or destination MAC addresses, whereas a multilayer switch generally supports rate-limiting frames on IP/MAC addresses.

Unfortunately, different models of Cisco switches support different capabilities, and some Layer 2-only switches actually support Layer 3 ACLs and QoS lookups. It is best to consult the product documentation at Cisco.com for clear information about what your switch supports. For the purpose of CCNP Switch and the context of this book, Layer 2 switches support ACLs and QoS based on MAC addresses, whereas Layer 3 switches support ACLs and QoS based on IP or MAC addresses.

Useful Commands for Viewing and Editing Catalyst Switch MAC Address Tables

There is one command for viewing the Layer 2 forwarding table on Catalyst and Nexus switches: **show mac address-table**. The table has many optional parameters to narrow the output to a more manageable result in large networks. The full command options are as follows: **show mac-address-table [aging-time | count | dynamic | static] [address *hw-addr*] [interface *interface-id*] [vlan *vlan-id*] [| {begin | exclude | include} *expression*]**.

Example 2-1 illustrates sample uses of the command and several useful optional uses.

Example 2-1 *Layer 2 Forwarding Table*

```
Switch1# show mac address-table
      Mac Address Table
-----
Vlan    Mac Address      Type      Ports
----    -
1       0000:0c00.9001   DYNAMIC   Et0/1
1       0000.0c00.9002   DYNAMIC   Et0/2
1       0000.0c00.9002   DYNAMIC   Et0/3
Total Mac Addresses for this criterion: 3

Switch1# show mac address-table interface ethernet 0/1
      Mac Address Table
-----
Vlan    Mac Address      Type      Ports
----    -
1       0000:0c00.9001   DYNAMIC   Et0/1
Total Mac Addresses for this criterion: 1

Switch1# show mac address-table | include 9001
1       0000:0c00.9001   DYNAMIC   Et0/1
```

Frame Rewrite

From your CCNA studies, you know that many fields of a packet must be rewritten when the packets are routed between subnets. These fields include both source and destination MAC addresses, the IP header checksum, the TTL (Time-to-Live), and the trailer checksum (Ethernet CRC). See Chapter 1, “Fundamentals Review,” for an example.

Distributed Hardware Forwarding

Network devices contain at least three planes of operation:

- Management plane
- Control plane
- Forwarding plane

The management plane is responsible for the network management, such as SSH access and SNMP, and may operate over an out-of-band (OOB) port. The control plane is responsible for protocols and routing decisions, and the forwarding plane is responsible for the actual routing (or switching) of most packets.

Multilayer switches must achieve high performance at line rate across a large number of ports. To do so, multilayer switches deploy independent control and forwarding planes. In this manner, the control plane will program the forwarding plane on how to route packets. Multilayer switches may also employ multiple forwarding planes. For example, a Catalyst 6800 uses forwarding planes on each line module, with a central control plane on the supervisor module.

To continue the example of the Catalyst 6800, each line module includes a microcoded processor that handles all packet forwarding. For the control plane on the supervisor to communicate with the line module, a control layer communication protocol exists, as shown in Figure 2-16.

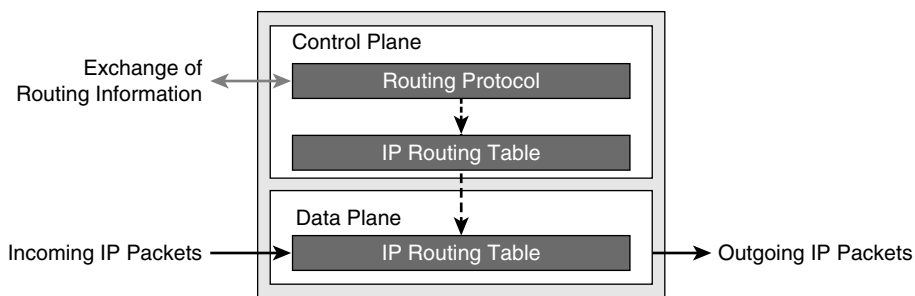


Figure 2-16 *Distributed Hardware Forwarding*

The main functions of this control layer protocol between the control plane and the forwarding plane are as follows:

- Managing the internal data and control circuits for the packet-forwarding and control functions
- Extracting the other routing and packet-forwarding-related control information from the Layer 2 and Layer 3 bridging and routing protocols and the configuration data, and then conveying the information to the interface module for control of the data path
- Collecting the data path information, such as traffic statistics, from the interface module to the route processor
- Handling certain data packets that are sent from the Ethernet interface modules to the route processor (for example, DHCP requests, broadcast packets, routing protocol packets)

Cisco Switching Methods

The term *Cisco switching methods* describes the route processor behavior found on Cisco IOS routers. Because multilayer switches are capable of routing and, in fact, contain a routing process, a review of these concepts is necessary.

A Cisco IOS-based router uses one of three methods to forward packets: process switching, fast switching, and Cisco Express Forwarding (CEF). Recall from your study of

routers that process switching is the slowest form of routing because the router processor must route and rewrite using software. Because speed and the number of cores limit the route processor, this method does not scale. The second method, fast switching, is a faster method by which the first packet in a flow is routed and rewritten by a route processor using software, and each subsequent packet is then handled by hardware. The CEF method uses hardware forwarding tables for most common traffic flows, with only a few exceptions. If you use CEF, the route processor spends its cycles mostly on other tasks.

The architecture of the Cisco Catalyst and Nexus switches both focus primarily on the Cisco router equivalents of CEF. The absolute last-resort switching method for Cisco Catalyst or Nexus switches is process switching. The route processors of these switches were never designed to switch or route packets, and by doing so, this will have an adverse effect on performance. Fortunately, the default behavior of these switches is to use fast switching or CEF, and process switching occurs only when necessary.

With Cisco Catalyst switching terminology, fast switching is referred to as *route caching*, and the application of CEF with distributed hardware forwarding is referred to as *topology-based switching*.

As a review, the following list summarizes route caching and topology-based forwarding on Cisco Catalyst switches:

- **Route caching:** Also known as *flow-based* or *demand-based switching*, route caching describes a Layer 3 route cache that is built within the hardware functions as the switch detects traffic flow into the switch. This method is functionally equivalent to fast switching in Cisco IOS Software.
- **Topology-based switching:** Information from the routing table is used to populate the route cache, regardless of traffic flow. The populated route cache is the FIB, and CEF is the facility that builds the FIB. This method is functionally equivalent to CEF in Cisco IOS Software.

The next subsections describe route caching and topology-based switching in more detail.

Route Caching

Route caching is the fast switching equivalent in Cisco Catalyst switches. For route caching to operate, the destination MAC address of an incoming frame must be that of a switch interface with Layer 3 capabilities. The first packet in a stream is switched in software by the route processor, because no cache entry exists yet for the new flow. The forwarding decision that is made by the route processor is then programmed into a cache table (the hardware forwarding table), and all subsequent packets in the flow are switched in the hardware, commonly referred to as using *application-specific interface circuits* (ASICs). Entries are created only in the hardware forwarding table as the switch detects new traffic flows, and entries will time out after they have been unused for a period of time.

Because entries are created only in the hardware cache as flows are detected by the switch, route caching will always forward at least one packet in a flow using software.

Route caching carries many other names, such as NetFow LAN switching, flow-based or demand-based switching, and route once, switch many.

Figure 2-17 briefly highlights this concept from a hardware perspective.

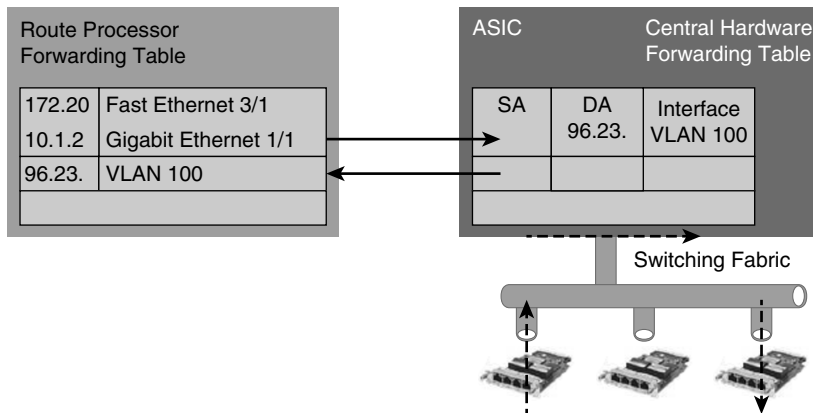


Figure 2-17 *Route Caching*

Topology-Based Switching

Topology-based switching is the CEF equivalent feature of Cisco Catalyst switches. Topology-based switching is ideal for Layer 3 switching over route caching because it offers the best performance and scalability. Fortunately, all Cisco Catalyst switches capable of Layer 3 routing leverage topology-based switching / CEF. For the purpose of CCNP Switch, focus primarily on the benefits and operation of topology-based switching.

CEF uses information in the routing table to populate a route cache (known as an FIB), without traffic flows being necessary to initiate the caching process. Because this hardware FIB exists regardless of traffic flow, assuming that a destination address has a route in the routing table, all packets that are part of a flow will be forwarded by the hardware. The FIB even handles the first packet of a flow. Figure 2-18 illustrates this behavior.

In addition, CEF adds enhanced support for parallel paths and thus optimizes load balancing at the IP layer. In most current-generation Catalyst switches, such as the Catalyst 4500 and 6800, CEF supports both load balancing based on source IP address and destination IP address combination and source and destination IP plus TCP/UDP port number.

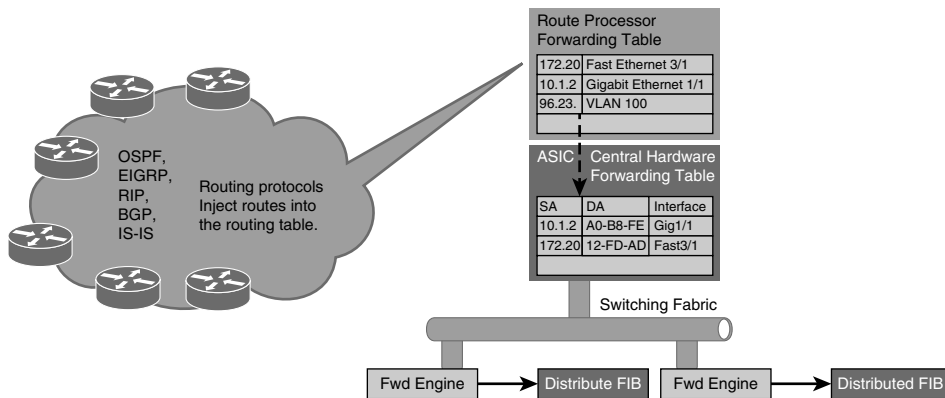


Figure 2-18 *Topology-Based Switching*

Note The load-balancing options and default behavior varies between different Catalyst switch models and software versions. Consult Cisco.com for the particular Catalyst switch you have in question for supported load-balancing methods and default configurations.

CEF load-balancing schemes allow for Layer 3 switches to use multiple paths to achieve load sharing. Packets for a given source-destination host pair are guaranteed to take the same path, even if multiple paths are available. This ensures that packets for a given host pair arrive in order, which in some cases may be the desired behavior with legacy applications.

Moreover, load balancing based only on source and destination IP address has a few shortcomings. Because this load-balancing method always selects the same path for a given host pair, a heavily used source-destination pair, such as a firewall to web server, might not leverage all available links. In other words, the behavior of this load-balancing scheme may “polarize” the traffic by using only one path for a given host pair, thus effectively negating the load-balancing benefit of the multiple paths for that particular host pair.

So, optimal use of any load-balancing scheme depends on the statistical distribution of traffic because source and destination IP load sharing becomes more effective as the number of source-destination IP pairs increases. In an environment where there is a broad distribution of traffic among host pairs, polarization is of minimal concern. However, in an environment where the data flow between a small number of host pairs creates a disproportionate percentage of the packets traversing the network, polarization can become a serious problem.

A popular alternative that is now the default behavior in new Catalyst switches is load balancing based on source and destination IP to include TCP/UDP port numbers. The more additional factors added to the load-balancing scheme, the less likely polarization will exist.

Cisco Catalyst supports additional load-balancing methods and features by which to tune load balancing based on hardware model and software version. Consult Cisco.com for such configuration optimizations if necessary.

Hardware Forward Details

The actual Layer 3 switching of packets occurs at two possible different locations on Catalyst switches. These possible locations are in a centralized manner, such as on a supervisor module, or in distributed fashion, where switching occurs on individual line modules. These methods are referred to as *centralized switching* and *distributed switching*, respectively.

The Catalyst 6500 was a perfect example where there was an option to centralize switch everything on the supervisor or place specific hardware versions of line modules in the chassis to gain distributed switching capability.

The benefits of centralized switching include lower hardware cost and lower complexity. For scaling and large enterprise core networks, distributed switching is optimal. Most small form-factor switches leverage centralized switching.

Note Some small form-factor switches may leverage a switch-on-chip (SOC) concept, where the entire intelligence and processing of the switch happens on a single low-cost ASIC. This practice has now become an industry standard for low-feature and low-cost switches and is found on specific fixed-port Cisco Catalyst and Nexus switches. In addition, newer generation modular switches such as the Nexus 9000 may leverage SOC in a hybrid capacity, whereas line modules may contain their own SOC and leverage distributed switching concepts.

In conclusion, the subsections of this chapter pertaining to switching methods and hardware forwarding included many specific details about routing and switching operations on Cisco switches. Among all the lengthy explanations and details, conclude this section with the following concepts:

- The control plane (CPU/route processor) of a Cisco Catalyst was never designed to route or switch frames. The control plane is intended only to populate hardware tables with routing information and maintain routing protocols. The control plane may route frames in a few exception conditions.
- Medium- to high-end Cisco Catalyst switches were designed based on the distributing forward model to scale to demands of campus and data center networks.
- Cisco Catalyst switches leverage CEF (topology-based switching) for routing of frames as a means to implement a distributing hardware forwarding model.
- Cisco Catalyst switches use either a centralized method or a distributed line module method of hardware forwarding, depending on specific platform model and configuration.

Study Tips

- The **show mac address-table** command displays the Layer 2 forwarding table of a Cisco switch.
- Layer 2 switches forward traffic based on the destination MAC address of a frame.
- Campus network designs are still built upon the hierarchical model, where end devices connect to the access layer, the distribution layer aggregates the access layer, and the core aggregates the entire enterprise network.
- Cisco switches leverage CEF (topology-based switching) for Layer 3 forwarding.

Summary

This chapter briefly introduced some concepts about campus networks, including the hierarchical model, benefits of Layer 3 routing the access, Cisco switches, and some hardware details related to Cisco Catalyst switches. The next chapters of this book go into more detail about specific feature and design elements of the campus network, such as VLANs, spanning tree, and security. The information in this chapter is summarized as follows:

- Flat Layer 2 networks are extremely limited in scale and in most cases will only scale to 10 to 20 end users before adverse conditions may occur.
- Despite its age, the hierarchical model continues to be a key design fundamental of any network design, including campus network designs.
- The hierarchical model consists of an access, distribution, and core layer, thus allowing for scalability and growth of a campus network in a seamless manner.
- The different models of Cisco Catalyst switches provide for a range of capabilities depending on need and placement within the hierarchical model.
- Cisco Catalyst switches leverage CAM for Layer 2 forwarding tables and TCAM for Layer 3 forwarding tables to achieve line-rate performance.
- Cisco Catalyst switches leverage CEF (topology-based switching) for routing, utilizing a distributed hardware forwarding model that is centralized or distributed per line card.

Review Questions

Use the questions in this section as a review of what you have learned in this chapter. The correct answers are found in Appendix A, “Answers to Chapter Review Questions.”

1. Which of the following statements is true about campus networks?
 - a. The campus network describes the interconnections of servers in a data center.
 - b. The campus network describes the WAN interconnectivity between two remote sites and head office.
 - c. The campus network describes the network devices that interconnect end users to applications such as e-mail, the intranet, or the Internet over wire or wireless connections.
2. Which of the following is a disadvantage to using flat Layer 2 networks?
 - a. Broadcast packets are flooded to every device in the network.
 - b. No IP boundary to administer IP-based access control.
 - c. A host flooding traffic onto the network effects every device.
 - d. Scalability is limited.
 - e. All the above
3. Why are networks designed with layers?
 - a. Allows focus within specific layers due to grouping, segmentation, and compartmentalization
 - b. Simplification of network design
 - c. Optimizes use of physical interconnects (links)
 - d. Optimizes application of policies and access control
 - e. Eases network management
 - f. All of the above
4. Identify the three layers of the hierarchical model for designing networks.
 - a. Core
 - b. Access
 - c. Distribution
 - d. Enterprise edge
 - e. WAN
 - f. Wireless

- 5.** What is another common name for the core layer?
 - a.** Backbone
 - b.** Campus
 - c.** Data center
 - d.** Routing layer
- 6.** In newer terminology, what layers are referred to as the spine layer and the leaf layer?
 - a.** The spine layer is the equivalent to the core layer, and the leaf layer is equivalent to the distribution layer.
 - b.** The spine layer is equivalent to the access layer, and the leaf layer is equivalent to the distribution layer.
 - c.** The spine layer is equivalent to the distribution layer, and the leaf layer is equivalent to the access layer.
 - d.** The spine layer is equivalent to the core layer, and the leaf layer is equivalent to the access layer.
- 7.** Match each layer to its definition.
 - a.** Core
 - b.** Distribution
 - c.** Access
 - 1.** Connects PCs, wireless access points, and IP phones
 - 2.** High-speed interconnectivity layer that generally supports routing capability
 - 3.** Aggregates access layer switches and provides for policy control
- 8.** Which of the following are generally true about recommended core layer designs?
 - a.** Requires high-availability and resiliency
 - b.** Connects critical application servers directly for optimal latency and bandwidth
 - c.** Leverages fixed form factor switches in large enterprises
- 9.** In which layer are you most likely to find fixed Catalyst switches?
 - a.** Access layer
 - b.** Core layer
 - c.** Distribution layer

- 10.** In which layer are you most likely to find modular Catalyst switches?
 - a.** Access layer
 - b.** Backbone layer
 - c.** Core layer
- 11.** Which of the following are benefits to using Layer 3 in the access layer? (Choose two.)
 - a.** Reduced cost
 - b.** Reduced Layer 2 domain
 - c.** Reduced spanning-tree domain
 - d.** Mobility
- 12.** Which of the following is the biggest disadvantage with using Layer 3 in the access layer using current technologies?
 - a.** More difficult troubleshooting
 - b.** Lack of broadcast forwarding
 - c.** Native mobility without additional features
 - d.** Lack of high availability
- 13.** A Layer 2-only switch makes forwarding decisions based on what?
 - a.** Source MAC address
 - b.** Destination MAC address
 - c.** Source IP address
 - d.** Destination IP address
- 14.** What does a switch do when it does not know how to forward a frame?
 - a.** Drops the frame
 - b.** Floods the frames on all ports in the same Layer 2 domain except the source port
 - c.** Stores the frame for later transmission
 - d.** Resends the frame out the port where it was received

- 15.** The Layer 2 forwarding table of Cisco switches is also referred to as which of the following?
 - a.** CAM table
 - b.** Routing table
 - c.** MAC address table
 - d.** FIB table
- 16.** Which of the following lookups does a Layer 2-only Cisco Catalyst switch perform on an ingress frame?
 - a.** Layer 2 forwarding for destination port
 - b.** ACL for access control
 - c.** NetFlow for statistics monitoring
 - d.** QoS for classification, marking, or policing
- 17.** Which of the following are true about CAM and/or TCAM? (Choose three.)
 - a.** TCAM stands for ternary content-addressable memory.
 - b.** CAM provides three results: 0, 1, and don't care.
 - c.** Leveraging CAM and TCAM ensures line-rate performance of the switch.
 - d.** CAM and TCAM are software-based tables.
 - e.** TCAM is leveraged by QoS and ACL tables.
- 18.** Why is TCAM necessary for IP routing tables over CAM?
 - a.** TCAM supports longest matching instead of match or not match.
 - b.** TCAM is faster than CAM.
 - c.** TCAM memory is cheaper than CAM.
- 19.** Cisco Catalyst switches leverage which of the following technologies for Layer 3 forwarding?
 - a.** Route caching
 - b.** Processor/CPU switching
 - c.** NetFlow
 - d.** CEF

- 20.** Cisco Catalyst switches relay routing information to hardware components for additional performance and scalability (line-rate forwarding). What are the two common hardware types that receive relayed routing information?
- a.** Centralized
 - b.** Distributed
 - c.** Aggregated
 - d.** Core-based
- 21.** With regard to load balancing, what term describes the situation where less than optimal use of all links occurs?
- a.** Reverse path forwarding (RPF)
 - b.** Polarization
 - c.** Inverse routing
 - d.** Unicast flooding
- 22.** What is the default load-balancing mechanism found on Cisco Catalyst switches?
- a.** Per-flow
 - b.** Per-destination IP address
 - c.** Per-packet
 - d.** Per-destination MAC address

This page intentionally left blank

Campus Network Architecture

This chapter covers the following topics:

- Implementing VLANs and trunks in campus switched architecture
- Understanding the concept of VTP and its limitation and configurations
- Implementing and configuring EtherChannel

This chapter covers the key concepts of VLANs, trunking, and EtherChannel to build the campus switched networks. Knowing the function of VLANs and trunks and how to configure them is the core knowledge needed for building a campus switched network. VLANs can span across the whole network, or they can be configured to remain local. Also, VLANs play a critical role in the deployment of voice and wireless networks. Even though you might not be a specialist at one of those two fields, it is important to understand basics because both voice and wireless often rely on a basic switched network.

Once VLANs are created, their names and descriptions are stored in a VLAN database, with the exception of specific VLANs such as VLANs in the extended range in Cisco IOS for the Catalyst 6500. A mechanism called *VLAN Trunking Protocol* (VTP) dynamically distributes this information between switches. However, even if network administrators do not plan to enable VTP, it is important to consider its consequences.

EtherChannel can be used to bundle physical links in one virtual link, thus increasing throughput. There are multiple ways traffic can be distributed over the physical link within the EtherChannel.

Implementing VLANs and Trunks in Campus Environment

Within the switched internetwork, VLANs provide segmentation and organizational flexibility. VLANs help administrators to have the end node or workstations group that

are segmented logically by functions, project teams, and applications, without regard to the physical location of the users. In addition, VLANs allow you to implement access and security policies to particular groups of users and limit the broadcast domain.

In addition, the voice VLAN feature enables access ports to carry IP voice traffic from an IP phone. Because the sound quality of an IP phone call can deteriorate if the data is unevenly sent, the switch supports quality of service (QoS).

This section discusses in detail how to plan, implement, and verify VLAN technologies and address schemes to meet the given business and technical requirements along with constraints. This ability includes being able to meet these objectives:

- Describe the different VLAN segmentation models
- Identify the basic differences between end-to-end and local VLANs
- Describe the benefits and drawbacks of local VLANs versus end-to-end VLANs
- Configure and verify VLANs
- Implement a trunk in a campus network
- Configure and verify trunks
- Explain switchport mode interactions
- Describe voice VLANs
- Configure voice VLANs

VLAN Overview

A VLAN is a logical broadcast domain that can span multiple physical LAN segments. Within the switched internetwork, VLANs provide segmentation and organizational flexibility. A VLAN can exist on a single switch or span multiple switches. VLANs can include (hosts or endnodes) stations in a single building or multiple-building infrastructures. As shown in Figure 3-1, sales, human resources, and engineering are three different VLANs spread across all three floors.

The Cisco Catalyst switch implements VLANs by only forwarding traffic to destination ports that are in the same VLAN as the originating ports. Each VLAN on the switches implements address learning, forwarding, and filtering decisions and loop-avoidance mechanisms, just as though the VLAN were a separate physical switch.

Ports in the same VLAN share broadcasts. Ports in different VLANs do not share broadcasts, as illustrated in Figure 3-2, where a PC 3 and PC 4 cannot ping because they are in different VLANs, whereas PC 1 and PC 2 can ping each other because they are part of the same VLAN. Containing broadcasts within a VLAN improves the overall performance of the network. Because a VLAN is a single broadcast domain, campus design best practices recommend mapping a VLAN generally to one IP subnet. To communicate between VLANs, packets need to pass through a router or Layer 3 device.


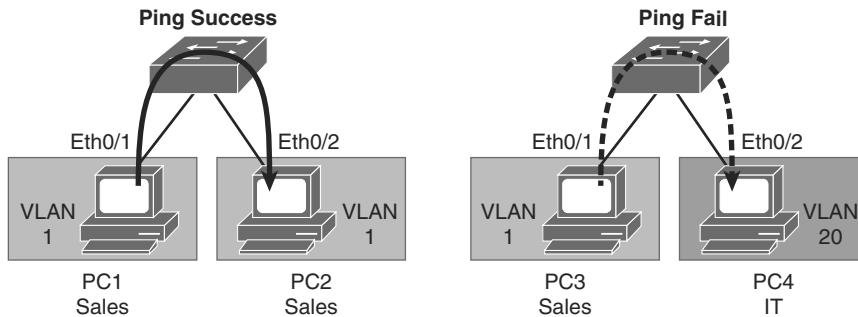


Figure 3-1 VLAN Overview



- VLAN is an independent LAN network.
- VLAN = broadcast domain.
- VLAN maps to logical network (subnet).
- VLANs provide segmentation, security, and network flexibility.

Figure 3-2 VLAN Broadcast Domain

Note Inter-VLAN routing is discussed in detail in Chapter 5, “Inter-VLAN Routing.” Generally, a port carries traffic only for the single VLAN. For a VLAN to span multiple switches, Catalyst switches use trunks. A trunk carries traffic for multiple VLANs by using Inter-Switch Link (ISL) encapsulation or IEEE 802.1Q. This chapter discusses trunking in more detail in later sections. Because VLANs are an important aspect of any campus design, almost all Cisco devices support VLANs and trunking.

Note Most of the Cisco products support only 802.1Q trunking because 802.1Q is the industry standard. This book focuses only on 802.1Q.

VLAN Segmentation

Larger flat networks generally consist of many end devices in which broadcasts and unknown unicast packets are flooded on all ports in the network. One advantage of using VLANs is the capability to segment the Layer 2 broadcast domain. All devices in a VLAN are members of the same broadcast domain. If an end device transmits a Layer 2 broadcast, all other members of the VLAN receive the broadcast. Switches filter the broadcast from all the ports or devices that are not part of the same VLAN.

In a campus design, a network administrator can design a campus network with one of two models: end-to-end VLANs or local VLANs. Business and technical requirements, past experience, and political motivations can influence the design chosen. Choosing the right model initially can help create a solid foundation upon which to grow the business. Each model has its own advantages and disadvantages. When configuring a switch for an existing network, try to determine which model is used so that you can understand the logic behind each switch configuration and position in the infrastructure.

End-to-End VLANs

The term *end-to-end VLAN* refers to a single VLAN that is associated with switch ports widely dispersed throughout an enterprise network on multiple switches. A Layer 2 switched campus network carries traffic for this VLAN throughout the network, as shown in Figure 3-3, where VLANs 1, 2, and 3 are spread across all three switches.




Figure 3-3 End to End VLAN

If more than one VLAN in a network is operating in the end-to-end mode, special links (Layer 2 trunks) are required between switches to carry the traffic of all the different VLANs.

An end-to-end VLAN model has the following characteristics:

- Each VLAN is dispersed geographically throughout the network.
- Users are grouped into each VLAN regardless of the physical location.
- As a user moves throughout a campus, the VLAN membership of that user remains the same, regardless of the physical switch to which this user attaches.
- Users are typically associated with a given VLAN for network management reasons. This is why they are kept in the same VLAN, therefore the same group, as they move through the campus.
- All devices on a given VLAN typically have addresses on the same IP subnet.
- Switches commonly operate in a server/client VTP mode.

Local VLANs

The campus enterprise architecture is based on the local VLAN model. In a local VLAN model, all users of a set of geographically common switches are grouped into a single VLAN, regardless of the organizational function of those users. Local VLANs are generally confined to a wiring closet, as shown in Figure 3-4. In other words, these VLANs are local to a single access switch and connect via a trunk to an upstream distribution switch. If users move from one location to another in the campus, their connection changes to the new VLAN at the new physical location.

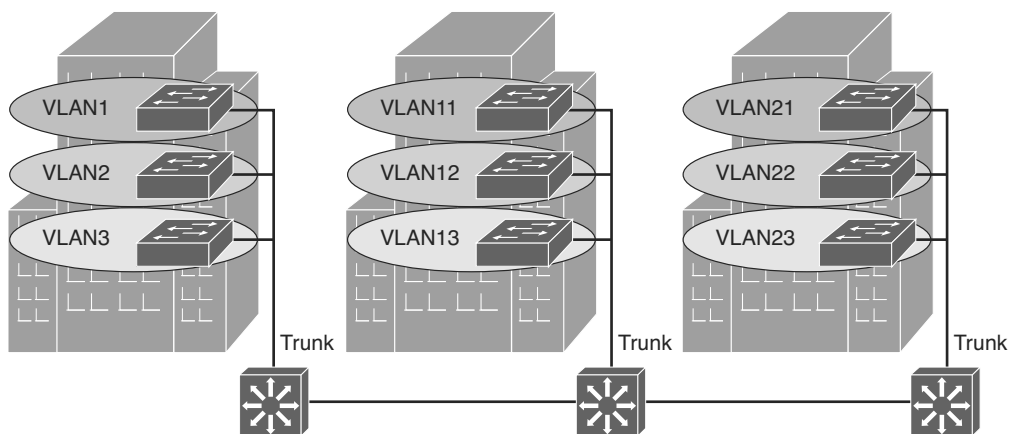


Figure 3-4 *Local VLANs*

In the local VLAN model, Layer 2 switching is implemented at the access level, and routing is implemented at the distribution and core level, as shown in Figure 2-4, to enable

users to maintain access to the resources they need. An alternative design is to extend routing to the access layer, and links between the access switches and distribution switches are routed links.

The following are some local VLAN characteristics and user guidelines:

- The network administrator should create local VLANs with physical boundaries in mind rather than the job functions of the users on the end devices.
- Generally, local VLANs exist between the access and distribution levels.
- Traffic from a local VLAN is routed at the distribution and core levels to reach destinations on other networks.
- Configure the VTP mode in transparent mode because VLANs on a given access switch should not be advertised to all other switches in the network, nor do they need to be manually created in any other switch VLAN databases.

Note VTP is discussed in more detail later in this chapter.

- A network that consists entirely of local VLANs can benefit from increased convergence times offered via routing protocols, instead of a spanning tree for Layer 2 networks. It is usually recommended to have one to three VLANs per access layer switch.

Comparison of End-to-End VLANs and Local VLANs

This subsection describes the benefits and drawbacks of local VLANs versus end-to-end VLANs.

Because a VLAN usually represents a Layer 3 segment, each end-to-end VLAN enables a single Layer 3 segment to be dispersed geographically throughout the network. The following could be some of the reasons for implementing the end-to-end design:

- **Grouping users:** Users can be grouped on a common IP segment, even though they are geographically dispersed. Recently, the trend has been moving toward virtualization. Solutions such as those from VMware need end-to-end VLANs to be spread across segments of the campus.
- **Security:** A VLAN can contain resources that should not be accessible to all users on the network, or there might be a reason to confine certain traffic to a particular VLAN.
- **Applying quality of service (QoS):** Traffic can be a higher- or lower-access priority to network resources from a given VLAN. Note that QoS may also be applied without the use of VLANs.

- **Routing avoidance:** If much of the VLAN user traffic is destined for devices on that same VLAN, and routing to those devices is not desirable, users can access resources on their VLAN without their traffic being routed off the VLAN, even though the traffic might traverse multiple switches.
- **Special-purpose VLAN:** Sometimes a VLAN is provisioned to carry a single type of traffic that must be dispersed throughout the campus (for example, multicast, voice, or visitor VLANs).
- **Poor design:** For no clear purpose, users are placed in VLANs that span the campus or even span WANs. Sometimes when a network is already configured and running, organizations are hesitant to improve the design because of downtime or other political reasons.

The following list details some considerations that the network administrators should consider when implementing end-to-end VLANs:

- Switch ports are provisioned for each user and associated with a given VLAN. Because users on an end-to-end VLAN can be anywhere in the network, all switches must be aware of that VLAN. This means that all switches carrying traffic for end-to-end VLANs are required to have those specific VLANs defined in each switch's VLAN database.
- Also, flooded traffic for the VLAN is, by default, passed to every switch even if it does not currently have any active ports in the particular end-to-end VLAN.
- Finally, troubleshooting devices on a campus with end-to-end VLANs can be challenging because the traffic for a single VLAN can traverse multiple switches in a large area of the campus, and that can easily cause potential spanning-tree problems.

Based on the data presented in this section, there are many reasons to implement end-to-end VLANs. The main reason to implement local VLANs is simplicity. Local VLAN configurations are quick and easy for small-scale networks.

Mapping VLANs to a Hierarchical Network

In the past, network designers have attempted to implement the 80/20 rule when designing networks. The rule was based on the observation that, in general, 80 percent of the traffic on a network segment was passed between local devices, and only 20 percent of the traffic was destined for remote network segments. Therefore, network architecture used to prefer end-to-end VLANs. To avoid the complications of end-to-end VLANs, designers now consolidate servers in central locations on the network and provide access to external resources, such as the Internet, through one or two paths on the network because the bulk of traffic now traverses a number of segments. Therefore, the paradigm now is closer to a 20/80 proportion, in which the greater flow of traffic leaves the local segment; so, local VLANs have become more efficient.

In addition, the concept of end-to-end VLANs was attractive when IP address configuration was a manually administered and burdensome process; therefore, anything that

reduced this burden as users moved between networks was an improvement. However, given the ubiquity of Dynamic Host Configuration Protocol (DHCP), the process of configuring an IP address at each desktop is no longer a significant issue. As a result, there are few benefits to extending a VLAN throughout an enterprise (for example, if there are some clustering and other requirements).

Local VLANs are part of the enterprise campus architecture design, as shown in Figure 3-4, in which VLANs used at the access layer should extend no further than their associated distribution switch. For example, VLANs 1, 10 and VLANs 2, 20 are confined to only a local access switch. Traffic is then routed out the local VLAN as to the distribution layer and then to the core depending on the destination. It is usually recommended to have two to three VLANs per access block rather than span all the VLANs across all access blocks. This design can mitigate Layer 2 troubleshooting issues that occur when a single VLAN traverses the switches throughout a campus network. In addition, because Spanning Tree Protocol (STP) is configured for redundancy, the switch limits the STP to only the access and distribution switches that help to reduce the network complexity in times of failure.

Implementing the enterprise campus architecture design using local VLANs provides the following benefits:

- **Deterministic traffic flow:** The simple layout provides a predictable Layer 2 and Layer 3 traffic path. If a failure occurs that was not mitigated by the redundancy features, the simplicity of the model facilitates expedient problem isolation and resolution within the switch block.
- **Active redundant paths:** When implementing Per-VLAN Spanning Tree (PVST) or Multiple Spanning Tree (MST) because there is no loop, all links can be used to make use of the redundant paths.
- **High availability:** Redundant paths exist at all infrastructure levels. Local VLAN traffic on access switches can be passed to the building distribution switches across an alternative Layer 2 path if a primary path failure occurs. Router redundancy protocols can provide failover if the default gateway for the access VLAN fails. When both the STP instance and VLAN are confined to a specific access and distribution block, Layer 2 and Layer 3 redundancy measures and protocols can be configured to failover in a coordinated manner.
- **Finite failure domain:** If VLANs are local to a switch block, and the number of devices on each VLAN is kept small, failures at Layer 2 are confined to a small subset of users.
- **Scalable design:** Following the enterprise campus architecture design, new access switches can be easily incorporated, and new submodules can be added when necessary.

Implementing a Trunk in a Campus Environment

A trunk is a point-to-point link that carries the traffic for multiple VLANs across a single physical link between the two switches or any two devices. Trunking is used to extend Layer 2 operations across an entire network, such as end-to-end VLANs, as shown in Figure 3-5. PC 1 in VLAN 1 can communicate with the host in VLAN 21 on another switch over the single trunk link, the same as a host in VLAN 20 can communicate with a host in another switch in VLAN 20.

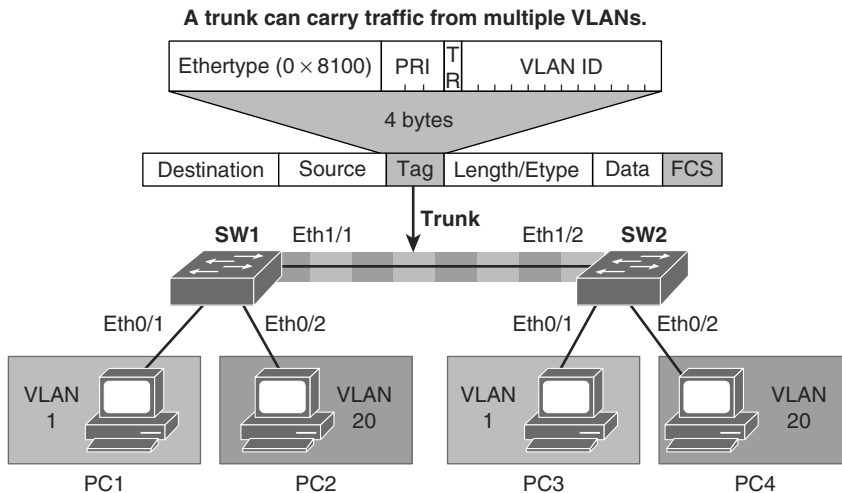


Figure 3-5 *Trunk Overview*

As discussed earlier in this chapter, to allow a switch port that connects two switches to carry more than one VLAN, it must be configured as a trunk. If frames from a single VLAN traverse a trunk link, a trunking protocol must mark the frame to identify its associated VLAN as the frame is placed onto the trunk link. The receiving switch then knows the frame's VLAN origin and can process the frame accordingly. On the receiving switch, the VLAN ID (VID) is removed when the frame is forwarded on to an access link associated with its VLAN.

A special protocol is used to carry multiple VLANs over a single link between two devices. There are two trunking technologies:

- **Inter-Switch Link (ISL):** A Cisco proprietary trunking encapsulation
- **IEEE 802.1Q:** An industry-standard trunking method

Note ISL is a Cisco proprietary implementation. It is not widely used anymore.

When configuring an 802.1Q trunk, a matching native VLAN must be defined on each end of the trunk link. A trunk link is inherently associated with tagging each frame with a VID. The purpose of the native VLAN is to enable frames that are not tagged with a VID to traverse the trunk link. Native VLAN is discussed in more detail in a later part of this section.

Because the ISL protocol is almost obsolete, this book focuses only on 802.1Q. Figure 3-6 depicts how ISL encapsulates the normal Ethernet frame. Currently, all Catalyst switches support 802.1Q tagging for multiplexing traffic from multiple VLANs onto a single physical link.

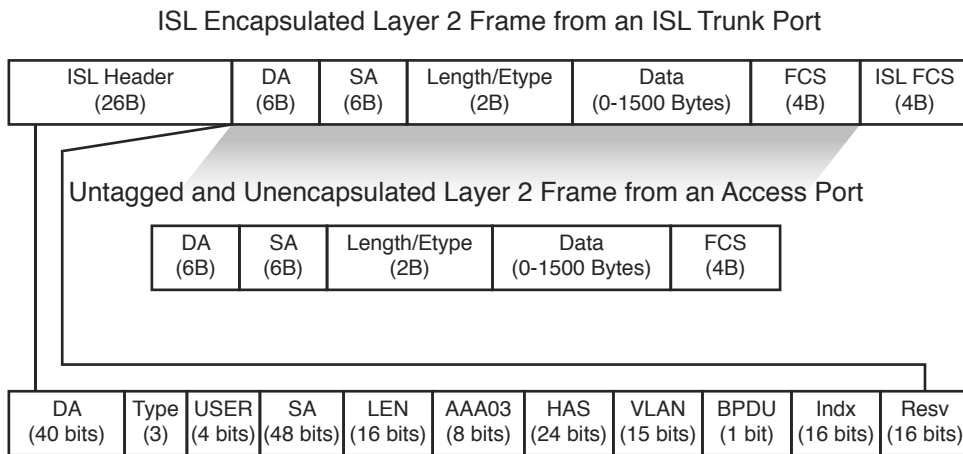


Figure 3-6 ISL Frame

IEEE 802.1Q trunk links employ the tagging mechanism to carry frames for multiple VLANs, in which each frame is tagged to identify the VLAN to which the frame belongs. Figure 3-7 shows the layout of the 802.1Q frame.

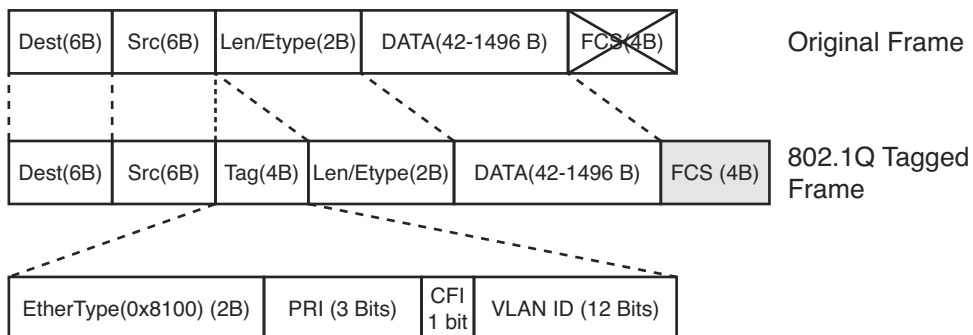


Figure 3-7 802.1Q Frame

The IEEE 802.1Q/802.1p standard provides the following inherent architectural advantages over ISL:

- 802.1Q has smaller frame overhead than ISL. As a result, 802.1Q is more efficient than ISL, especially in the case of small frames. 802.1Q overhead is 4 bytes, whereas ISL is 30 bytes.
- 802.1Q is a widely supported industry standard protocol.
- 802.1Q has the support for 802.1p fields for QoS.

The 802.1Q Ethernet frame header contains the following fields:

- **Dest:** Destination MAC address (6 bytes)
- **Src:** Source MAC address (6 bytes)
- **Tag:** Inserted 802.1Q tag (4 bytes, detailed here)
 - **EtherType(TPID):** Set to 0x8100 to specify that the 802.1Q tag follows.
 - **PRI:** 3-bit 802.1p priority field.
 - **CFI:** Canonical Format Identifier is always set to 0 for Ethernet switches and to 1 for Token Ring-type networks.
 - **VLAN ID:** 12-bit VLAN field. Of the 4096 possible VLAN IDs, the maximum number of possible VLAN configurations is 4094. A VLAN ID of 0 indicates priority frames, and value 4095 (FFF) is reserved. CFI, PRI, and VLAN ID are represented as Tag Control Information (TCI) fields.
- **Len/Etype:** 2-byte field specifying length (802.3) or type (Ethernet II)
- **Data:** Data itself
- **FCS:** Frame check sequence (4 bytes)

IEEE 802.1Q uses an internal tagging mechanism that modifies the original frame (as shown by the *X* over FCS in the original frame in Figure 3-7), recalculates the cyclic redundancy check (CRC) value for the entire frame with the tag, and inserts the new CRC value in a new FCS. ISL, in comparison, wraps the original frame and adds a second FCS that is built only on the header information but does not modify the original frame FCS.

IEEE 802.1p redefined the three most significant bits in the 802.1Q tag to allow for prioritization of the Layer 2 frame.

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

Baby giants are frames that are larger than the standard MTU of 1500 bytes but less than 2000 bytes. Because ISL and 802.1Q tagged frames increase the MTU beyond 1500 bytes, switches consider both frames as baby giants. ISL-encapsulated packets over Ethernet have an MTU of 1548 bytes, whereas 802.1Q has an MTU of 1522 bytes.

Understanding Native VLAN in 802.1Q Trunking

The IEEE 802.1Q protocol allows operation between equipment from different vendors. All frames, except native VLAN, are equipped with a tag when traversing the link, as shown in Figure 3-8.

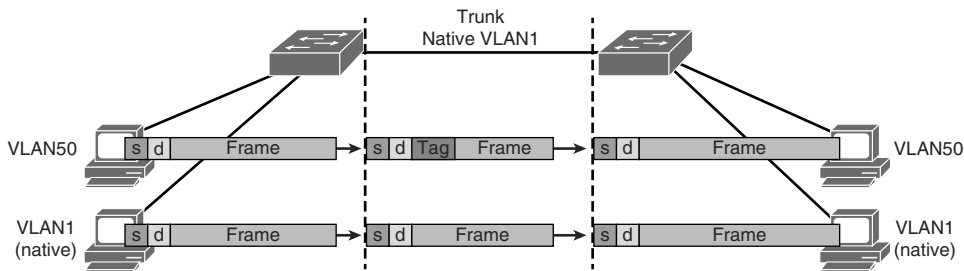


Figure 3-8 Native VLAN in 802.1Q

A frequent configuration error is to have different native VLANs. The native VLAN that is configured on each end of an 802.1Q trunk must be the same. If one end is configured for native VLAN 1 and the other for native VLAN 2, a frame that is sent in VLAN 1 on one side will be received on VLAN 2 on the other. VLAN 1 and VLAN 2 have been segmented and merged. There is no reason this should be required, and connectivity issues will occur in the network. If there is a native VLAN mismatch on either side of an 802.1Q link, Layer 2 loops may occur because VLAN 1 STP BPDUs are sent to the IEEE STP MAC address (0180.c200.0000) untagged.

Cisco switches use Cisco Discovery Protocol (CDP) to warn of a native VLAN mismatch. On select versions of Cisco IOS Software, CDP may not be transmitted or will be automatically turned off if VLAN 1 is disabled on the trunk.

By default, the native VLAN will be VLAN 1. For the purpose of security, the native VLAN on a trunk should be set to a specific VID that is not used for normal operations elsewhere on the network.

```
Switch(config-if)# switchport trunk native vlan vlan-id
```

Note Cisco ISL does not have a concept of native VLAN. Traffic for all VLANs is tagged by encapsulating each frame.

Understanding DTP

All recent Cisco Catalyst switches, except for the Catalyst 2900XL and 3500XL, use a Cisco proprietary point-to-point protocol called *Dynamic Trunking Protocol* (DTP) on trunk ports to negotiate the trunking state. DTP negotiates the operational mode of directly connected switch ports to a trunk port and selects an appropriate trunking protocol. Negotiating trunking is a recommended practice in multilayer switched networks because it avoids network issues resulting from trunking misconfigurations for initial configuration, but best practice is when the network is stable, change to permanent trunk.

Cisco Trunking Modes and Methods

Table 3-1 describes the different trunking modes supported by Cisco switches.

Table 3-1 *Trunking Modes*

Mode in Cisco IOS	Function
Access	Puts the interface into permanent nontrunking mode and negotiates to convert the link into a nontrunk link. The interface becomes a nontrunk interface even if the neighboring interface does not agree to the change.
Trunk	Puts the interface into permanent trunking mode and negotiates to convert the link into a trunk link. The interface becomes a trunk interface even if the neighboring interface does not agree to the change.
Nonegotiate	Prevents the interface from generating DTP frames. You must configure the local and neighboring interface manually as a trunk interface to establish a trunk link. Use this mode when connecting to a device that does not support DTP.
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 auto mode.
Dynamic auto	Makes the interface willing to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to trunk or desirable mode. This is the default mode for all Ethernet interfaces in Cisco IOS.

Note The Cisco Catalyst 4000 and 4500 switches run Cisco IOS or Cisco CatOS depending on the Supervisor Engine model. The Supervisor Engines for the Catalyst 4000 and 4500 do not support ISL encapsulation on a per-port basis. Refer to the product documentation on Cisco.com for more details.

Figure 3-9 shows the combination of DTP modes between the two links. A combination of DTP modes can either make the port as an access port or trunk port.

	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

Figure 3-9 Output from the SIMPLE Program

VLAN Ranges and Mappings

ISL supports VLAN numbers in the range of 1 to 1005, whereas 802.1Q VLAN numbers are in the range of 1 to 4094. The default behavior of VLAN trunks is to permit all normal and extended-range VLANs across the link if it is an 802.1Q interface and to permit normal VLANs in the case of an ISL interface.

VLAN Ranges

Cisco Catalyst switches support up to 4096 VLANs depending on the platform and software version. Table 3-2 illustrates the VLAN division for Cisco Catalyst switches. Table 3-3 shows VLAN ranges.

Note The Catalyst 2950 and 2955 support as many as 64 VLANs with the Standard Software image, and up to 250 VLANs with the Enhanced Software image. Cisco Catalyst switches do not support VLANs 1002 through 1005; these are reserved for Token Ring and FDDI VLANs. Furthermore, the Catalyst 4500 and 6500 families of switches do not support VLANs 1006 through 1024. In addition, several families of switches support more VLANs than the number of spanning-tree instances. For example, the Cisco Catalyst 2970 supports 1005 VLANs but only 128 spanning-tree instances. For information on the number of supported spanning-tree instances, refer to the Cisco product technical documentation.

Table 3-2 *VLAN Support Matrix for Catalyst Switches*

Type of Switch	Maximum Number of VLANs	VLAN ID Range
Catalyst 2940	4	1–1005
Catalyst 2950/2955	250	1–4094
Catalyst 2960	255	1–4094
Catalyst 2970/3550/3560/3750	1005	1–4094
Catalyst 2848G/2980G/4000/4500	4094	1–4094
Catalyst 6500	4094	1–4094

Table 3-3 *VLAN Ranges*

VLAN Range	Range Usage	Propagated via VTP
0, 4095	Reserved for system use only. You cannot see or use these VLANs.	—
1	Normal Cisco default. You can use this VLAN, but you cannot delete it.	Yes
2–1001	Normal For Ethernet VLANs. You can create, use, and delete these VLANs.	Yes
1002–1005	Normal Cisco defaults for FDDI and Token Ring. You cannot delete VLANs 1002–1005.	Yes
1006–1024	Reserved for system use only. You cannot see or use these VLANs.	—
1025–4094	Extended for Ethernet VLANs only.	Not supported in VTP Versions 1 and 2. The switch must be in VTP transparent mode to configure extended-range VLANs. This range is only supported in Version 3.

Configuring, Verifying, and Troubleshooting VLANs and Trunks

This section provides the configuration, verification, and troubleshooting steps for VLANs and trunking.

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

Step 1. Enter global configuration mode:

```
Switch# configure terminal
```

Step 2. Create a new VLAN with a particular ID number:

```
Switch(config)# vlan vlan-id
```

Step 3. (Optional.) Name the VLAN:

```
Switch(config-vlan)# name vlan-name
```

Example 3-1 shows how to configure a VLAN in global configuration mode.

Example 3-1 *Creating a VLAN in Global Configuration Mode in Cisco IOS*

```
Switch# configure terminal
Switch(config)# vlan 5
Switch(config-vlan)# name Engineering
Switch(config-vlan)# exit
```

To delete a VLAN in global configuration mode, delete the VLAN by referencing its ID number:

```
Switch(config)# no vlan vlan-id
```

Note After a VLAN is deleted, the access ports that belong to that VLAN move into the inactive state until the ports are moved to another VLAN. As a security measure, ports in the inactive state do not forward traffic.

Example 3-2 demonstrates deletion of a VLAN in global configuration mode.

Example 3-2 *Deleting a VLAN in Global Configuration Mode*

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

To assign a switch port to a previously created VLAN, follow these steps:

Step 1. From global configuration mode, enter the configuration mode for the particular port you want to add to the VLAN:

```
Switch(config)# interface interface-id
```

Step 2. Specify the port as an access port:

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

```
Switch(config-if)# switchport host
```

Note The `switchport host` command effectively configures a port for a host device, such as a workstation or server. This feature is a macro for enabling spanning-tree PortFast and disabling EtherChanneling on a per-port basis. These features are discussed in later chapters. The `switchport mode access` command is needed so that the interface doesn't attempt to negotiate trunking.

Step 3. Remove or place the port in a particular VLAN:

```
Switch(config-if)# [no] switchport access vlan vlan-id
```

Example 3-3 illustrates configuration of an interface as an access port in VLAN 200.

Example 3-3 *Assigning an Access Port to a VLAN*

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface FastEthernet 5/6
Switch(config-if)# description PC A
Switch(config-if)# switchport
Switch(config-if)# switchport host
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 200
Switch(config-if)# no shutdown
Switch(config-if)# end
```

Note Use the `switchport` command with no keywords to configure interfaces as Layer 2 interfaces on Layer 3 switches. After configuring the interface as a Layer 2 interface, use additional `switchport` commands with keywords to configure Layer 2 properties, such as access VLANs or trunking.

Verifying the VLAN Configuration

As previously discussed, after you configure the VLANs, one of the important steps is to be able to verify the configuration. To verify the VLAN configuration of a Catalyst switch, use `show` commands. The `show vlan` command from privileged EXEC mode displays information about a particular VLAN. Table 3-4 documents the fields displayed by the `show vlan` command.

Table 3-4 *show vlan Field Descriptions*

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

Example 3-4 displays information about a VLAN identified by number in Cisco IOS.

Example 3-4 *Displaying Information About a VLAN by Number in Cisco IOS*

```

SW1#show vlan id 3

VLAN Name                Status    Ports
-----
3      VLAN0003                active    Et1/1

VLAN Type  SAID      MTU   Parent  RingNo BridgeNo  Stp   BrdgMode Trans1 Trans2
-----
3      enet    100003   1500   -       -        -    -        0      0

Primary Secondary Type            Ports
-----
SW1#

```

Example 3-5 displays information about a VLAN identified by name in Cisco IOS.

Example 3-5 *Displaying Information About a VLAN by Name in Cisco IOS*

```
SW1# show vlan name VLAN0003
```

VLAN Name	Status	Ports
3 VLAN0003	active	Et1/1

VLAN Type	SAID	MTU	Parent	RingNo	BridgeNo	Stp	BrdgMode	Trans1	Trans2
3 enet	100003	1500	-	-	-	-	-	0	0

Primary	Secondary	Type	Ports

```
SW1#
```

To display the current configuration of a particular interface, use the **show running-config interface *interface-type slot/port*** command. To display detailed information about a specific switch port, use the **show interfaces** command. The command **show interfaces *interface-type slot/port*** with the **switchport** keyword displays not only a switch port's characteristics but also private VLAN and trunking information. The **show mac address-table interface *interface-type slot/port*** command displays the MAC address table information for the specified interface in specific VLANs. During troubleshooting, this command is helpful in determining whether the attached devices are sending packets to the correct VLAN.

Example 3-6 displays the configuration of a particular interface. Example 3-6 shows that the interface Ethernet 5/6 is configured with the VLAN 200 and in an access mode so that the port does not negotiate for trunking.

Example 3-6 *Displaying Information About the Interface Config*

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

Example 3-7 displays detailed switch port information as the port VLAN and operation modes. As shown in Example 3-7, the Ethernet port 4/1 is configured as the switch port means Layer 2 port, working as an access port in VLAN 2.

Example 3-7 *Displaying Detailed Switch Port Information*

```

BXB-6500-10:8A# SW1# show int ethernet 4/1 switchport
Name: Et4/1
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: native
Operational Dot1q Ethertype: 0x8100
Negotiation of Trunking: Off
Access Mode VLAN: 200 (Inactive)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Operational Native VLAN tagging: disabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL

Voice VLAN: none (Inactive)
Appliance trust: none

```

Example 3-8 displays the MAC address table information for a specific interface in VLAN 1.

Example 3-8 *Displaying MAC Address Table Information*

```

Switch# show mac-address-table interface GigabitEthernet 0/1 vlan 1
SW1# show mac address-table interface GigabitEthernet 0/1
      Mac Address Table
-----
Vlan    Mac Address      Type        Ports
----    -
1       aabb.cc01.0600   DYNAMIC     Gi0/1
Total Mac Addresses for this criterion: 1

```

Note In this book, the configuration and verification are shown as the part of the scenarios that will be shown in a particular topology.

To configure the VLANs on switches SW1 and SW2 and enable trunking between the switches, use the topology shown in Figure 3-10.

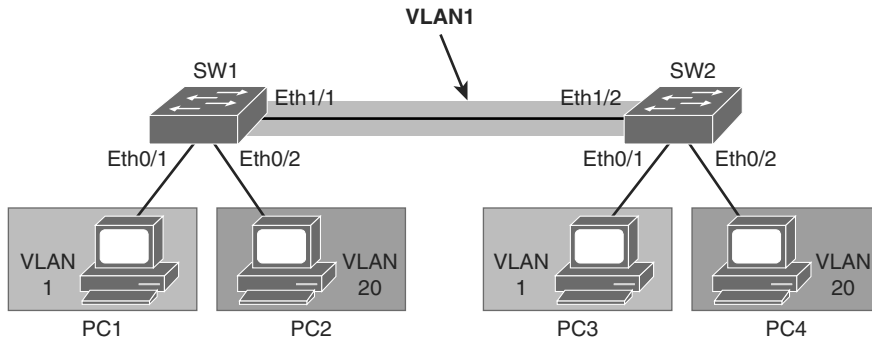


Figure 3-10 *Topology to Configure VLAN and Trunking*

Table 3-5 outlines the IP addressing scheme that will be used for this topology.

Table 3-5 *IP Addressing*

Device	Device IP	Device Interface	Device Neighbor	Interface on the Neighbor
PC1	192.168.1.100	Eth0/0	SW1	Eth0/1
PC2	192.168.20.101	Eth0/0	SW1	Eth0/2
PC3	192.168.1.110	Eth0/0	SW2	Eth0/1
PC4	192.168.20.110	Eth0/0	SW2	Eth0/2

Configuring VLANs and Trunks

To configure a port to belong to a certain VLAN, you have the following two options:

- Static VLAN configuration
- Dynamic VLAN configuration

With static VLAN configuration, switch ports are assigned to a specific VLAN. End devices become members in a VLAN based on the physical port to which they are connected. The end device is not even aware that a VLAN exists. Each port that is assigned to a VLAN receives a port VLAN ID (PVID).

With dynamic VLAN configuration, membership is based on the MAC address of the end device. When a device is connected to a switch port, the switch must query a database to figure out what VLAN needs to be configured. With dynamic VLANs, you need to assign a user's MAC address to VLAN in the database of a VLAN Management Policy Server (VMPS). With dynamic VLANs, users can connect to any port on the switch, and they will be automatically assigned into the VLAN they belong to.

Note This book focuses only on configuring VLANs statically. All Cisco Catalyst switches support VLANs. That said, each Cisco Catalyst switch supports a different number of VLANs, with high-end Cisco Catalyst switches supporting as many as 4096 VLANs. Table 3-2 notes the maximum number of VLANs supported by each model of Catalyst switch. With static VLAN configuration, you first need to create a VLAN on the switch, if it does not yet exist. VLANs are identified by the VLAN number that runs 1 through 4094.

Step 1. Create VLAN 20 on both switches.

```
SW1(config)# vlan 20
SW1(config-vlan)# exit
% Applying VLAN changes may take few minutes. Please
wait...SW1(config)#
```

Step 2. As shown in Figure 3-10, on SW1 configure port Ethernet 0/2 to be an access port and assign it to VLAN 20. By default, it is part of VLAN 1:

```
SW1(config)# interface ethernet 0/2
SW1(config-if)# switchport mode access
SW1(config-if)# switchport access vlan 20
```

The **switchport mode access** command explicitly tells the port to be assigned only a single VLAN, providing connectivity to an end user. When you assign a switch port to a VLAN using this method, it is known as a *static access port*.

Step 3. On SW1, verify membership of port Ethernet 0/2.

Use the **show vlan** command to display information on all configured VLANs. The command displays configured VLANs, their names, and the ports on the switch that are assigned to each VLAN:

```
SW1# show vlan
```

VLAN	Name	Status	Ports
1	default	active	Et0/0, Et0/1, Et0/3, Et1/0 Et1/2, Et1/3, Et2/0, Et2/1 Et2/2, Et2/3, Et3/0, Et3/1 Et3/2, Et3/3, Et4/0, Et4/1


```

                                Et4/2, Et4/3, Et5/0, Et5/1
                                Et5/2, Et5/3
20  IT                            active  Et0/2
1002 fddi-default                 act/unsup
1003 token-ring-default          act/unsup
1004 fddinet-default             act/unsup
1005 trnet-default               act/unsup

VLAN Type  SAID      MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
-----
1   enet   100001    1500  -      -      -      -      -      0      0
20  enet   100020    1500  -      -      -      -      -      0      0
1002 fddi   101002    1500  -      -      -      -      -      0      0
1003 tr    101003    1500  -      -      -      -      -      0      0
1004 fdnet 101004    1500  -      -      -      ieee  -      0      0
1005 trnet 101005    1500  -      -      -      ibm   -      0      0

Primary Secondary Type          Ports
-----

```

In the **show vlan** output, you can see that VLAN 20, named IT, is created. Also notice that Ethernet 0/2 is assigned to VLAN 20.

Use the **show vlan id *vlan-number*** or the **show vlan name *vlan-name*** command to display information about a particular VLAN.

Note If you do not see a port listed in the output, this is probably because it is not configured as an access port.

Step 4. Ping from PC 1 to PC 3. The ping should be successful:

```

PC1# ping 192.168.1.110
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.110, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
First few pings might fail because of the Address Resolution Protocol (ARP)
process.

```

PC 1 and PC 3 belong to the same VLAN. The configuration on the two ports that connect switches SW1 and SW2 is default; both ports belong to VLAN 1. So PCs 1 and 3 belong to the same LAN-Layer 2 network.

Step 5. Ping from PC 2 to PC 4. The ping should not be successful.

The ping should not be successful because the link between SW1 and SW2 is an access link and carries only data for VLAN 1:

```
PC2# ping 192.168.20.110
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.20.110, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Step 6. Configure ports that connect SW1 and SW2 as trunks. Use the dot1Q encapsulation. Allow only VLANs 1 and 20 to traverse the trunk link.

Trunk configuration on SW1:

```
SW1(config)# interface Ethernet 1/1
SW1(config-if)# switchport trunk encapsulation dot1q
SW1(config-if)# switchport trunk allowed vlan 1,20
SW1(config-if)# switchport mode trunk
```

Trunk configuration on SW2:

```
SW2(config)# interface Ethernet 1/2
SW2(config-if)# switchport trunk encapsulation dot1q
SW2(config-if)# switchport trunk allowed vlan 1,20
SW2(config-if)# switchport mode trunk
```

If you do not explicitly allow VLANs to traverse the trunk, all traffic will be allowed to cross the link. This includes broadcasts for all VLANs, using unnecessary bandwidth.

Step 7. Verify that Ethernet 1/1 on SW1 is now trunking:

```
SW1# show interfaces trunk
```

Port	Mode	Encapsulation	Status	Native vlan
Et1/1	on	802.1q	trunking	1

Port	Vlans allowed on trunk
Et1/1	1,20

Port	Vlans allowed and active in management domain
Et1/1	1,20

Port	Vlans in spanning tree forwarding state and not pruned
Et1/1	1,20

Also notice that only VLANs 1 and 20 are allowed on the trunk.

Step 8. Issue a ping from PC2 to PC4. The ping should be successful.

You have configured the link between SW1 and SW2 to carry data for both VLAN 1 and VLAN 20:

```
PC2# ping 192.168.20.110
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.20.110, timeout is 2 seconds:
..!!!
Success rate is 60 percent (3/5), round-trip min/avg/max = 1/1/1 ms
```

Best Practices for VLANs and Trunking

Usually, network designers design and implement the VLANs and their components depending on the business needs and requirements, but this section provides general best practices for implementing VLAN in a campus network.

Following are some of the practices for VLAN design:

- For the Local VLANs model, it is usually recommended to have only one to three VLANs per access module and, as discussed, limit those VLANs to a couple of access switches and the distribution switches.
- Avoid using VLAN 1 as the black hole for all unused ports. Use any other VLAN except 1 to assign all the unused ports to it.
- Try to always have separate voice VLANs, data VLANs, management VLANs, native VLANs, black hole VLANs, and default VLANs (VLAN 1).
- In the local VLANs model, avoid VTP; it is feasible to use manually allowed VLANs in a network on trunks.
- For trunk ports, turn off DTP and configure it manually. Use IEEE 802.1Q rather than ISL because it has better support for QoS and is a standard protocol.
- Manually configure access ports that are not specifically intended for a trunk link.
- Prevent all data traffic from VLAN 1; only permit control protocols to run on VLAN 1 (DTP, VTP, STP bridge protocol data units [BPDUs], Port Aggregation Protocol [PAgP], Link Aggregation Control Protocol [LACP], Cisco Discovery Protocol [CDP], and such.).
- Avoid using Telnet because of security risks; enable Secure Shell (SSH) support on management VLANs.
- In a hierarchical design, access layer switches connect to distribution layer switches. This is where the trunks are implemented, as illustrated in Figure 3-11, where the links from each access switch to the distribution switches are the trunk links because they must carry two VLANs from each switch. Links between distribution and core layers are usually Layer 3. Also, to avoid spanning-tree problems, it is usually recommended not to link the two distribution switches as Layer 2 trunk links or have no link between them. In this way, the access layer switches are configured as a

spanning-tree, loop-free V topology if one distribution link fails, using the Hot Standby Router Protocol (HSRP) or Virtual Router Redundancy Protocol (VRRP) for creating a virtual default gateway. Spanning tree, HSRP, and VRRP are discussed more in later chapters.

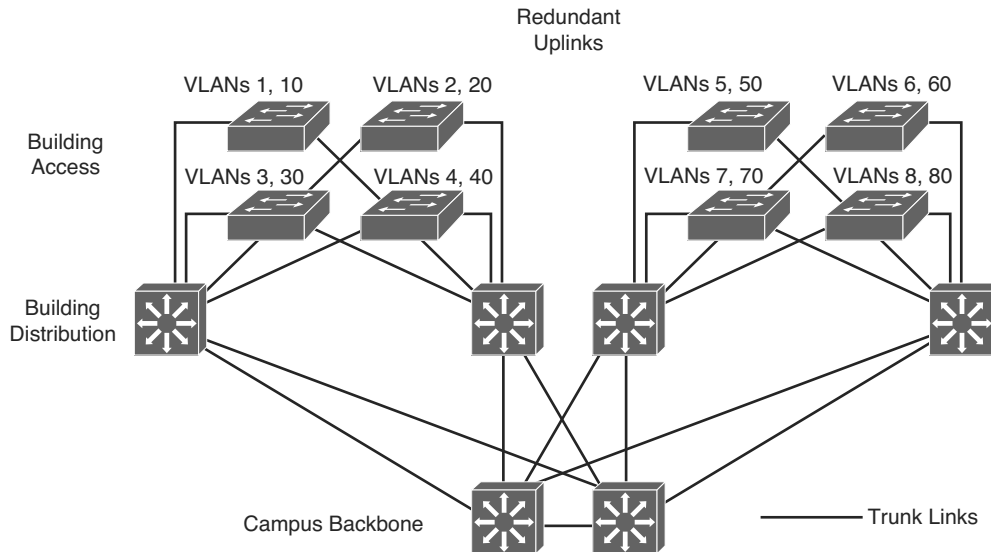


Figure 3-11 *Trunk Implementations*

- DTP is useful when the status of the switch on the other end of the link is uncertain or might be changing over time. When the link is to be set to trunk in a stable manner, changing both ends to trunk `nonnegotiate` accelerates the convergence time, saving up to 2 seconds upon boot time. We recommend this mode on stable links between switches that are part of the same core infrastructure.
- On trunk links, it is recommended to manually prune the VLANs that are not used. You can use VTP pruning if VTP is in use, but manual pruning (using a switchport trunk allowed VLAN) is a secure way of allowing only those VLANs that are expected and allowed on the link. In addition to this, it is also a good practice to have an unused VLAN as a native VLAN on the trunk links to prevent DTP spoofing.
- If trunking is not used on a port, you can disable it with the interface level command `switchport host`. This command is a macro that sets the port to access mode (switchport mode access) and enables portfast.

Voice VLAN Overview

Some Cisco Catalyst switches offer a unique feature called *voice VLAN*, which lets you overlay a voice topology onto a data network. You can segment phones into separate logical networks even though the data and voice infrastructure are physically the same.

The voice VLAN feature places the phones into their own VLANs without any end-user intervention. These VLAN assignments can be seamlessly maintained even if the phone is moved to a new location.

The user simply plugs the phone into the switch, and the switch provides the phone with the necessary VLAN information. By placing phones into their own VLANs, network administrators gain the advantages of network segmentation and control. Furthermore, network administrators can preserve their existing IP topology for the data end stations. IP phones can be easily assigned to different IP subnets using standards-based DHCP operation.

With the phones in their own IP subnets and VLANs, network administrators can more easily identify and troubleshoot network problems. In addition, network administrators can create and enforce QoS or security policies.

With the voice VLAN feature, Cisco enables network administrators to gain all the advantages of physical infrastructure convergence while maintaining separate logical topologies for voice and data terminals. This ability offers the most effective way to manage a multiservice network.

Multiservice switches support a new parameter for IP telephony support that makes the access port a multi-VLAN access port. The new parameter is called a *voice* or *auxiliary VLAN*. Every Ethernet 10/100/1000 port in the switch is associated with two VLANs:

- A native VLAN for data service that is identified by the PVID
- A voice VLAN that is identified by the voice VLAN ID (VVID)

During the initial CDP exchange with the access switch, the IP phone is configured with a VVID.

The IP phone is also supplied with a QoS configuration using CDP.

Data packets between the multiservice access switch and the PC or workstation are on the native VLAN. All packets going out on the native VLAN of an IEEE 802.1Q port are sent untagged by the access switch. The PC or workstation connected to the IP phone usually sends untagged packets, as shown in Figure 3-12, whereas a PC VLAN that connected directly to the phone sends untagged packets because this considers the native VLAN and voice VLAN as VVID 110. The IP phone tags voice packets based on the CDP information from the access switch.

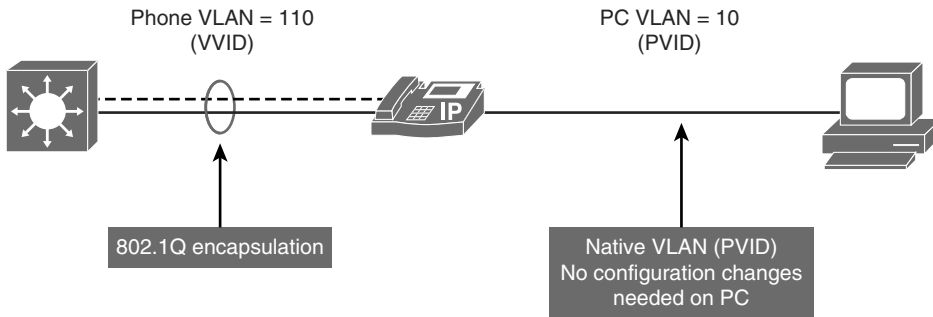


Figure 3-12 Voice VLAN Overview

The multi-VLAN access ports are not trunk ports, even though the hardware is set to the dot1Q trunk. The hardware setting is used to carry more than two VLANs, but the port is still considered an access port that is able to carry one native VLAN and the voice VLAN.

The **switchport host** command can be applied to a multi-VLAN access port on the access switch.

As shown in Figure 3-13, interface Fa0/1 is configured to set data devices in data VLAN 10 and VoIP devices in voice VLAN 110.




Figure 3-13 Voice VLAN Configuration

When you run the **show vlan** command, both the voice and the data VLAN are seen applied to the interface Fa0/1 as demonstrated in Example 3-9.

Example 3-9 *show vlan Command Output Provides Information About the Voice and Data VLAN*

```
Switch# show vlan

VLAN  Name                Status              Ports
-----
1     default              active              Fa0/6,Fa0/7,Fa0/8,Fa0/9,Fa0/10

10    VLAN0010             active              Fa0/1
110   VLAN0110             active              Fa0/1
<... output omitted ...>
```

Verify the switchport mode and the voice VLAN by using the `show interface interface-slot/number switchport` command.

Switch Configuration for Wireless Network Support

Cisco offers the following two WLAN implementations:

- The standalone WLAN solution is based on autonomous (standalone) access points (APs).
- The controller-based WLAN solution is based on controller-based APs and WLCs (Wireless LAN Controllers).

In the autonomous (or standalone) solution, each AP operates independently and acts as a transition point between the wireless media and the 802.3 media. The data traffic between two clients flows via the Layer 2 switch when on the same subnet from a different AP infrastructure. As the AP converts the IEEE 802.11 frame into an 802.3 frame, the wireless client MAC address is transferred to the 802.3 headers and appears as the source for the switch. The destination, also a wireless client, appears as the destination MAC address. For the switch, the APs are relatively transparent, as illustrated in Figure 3-14.

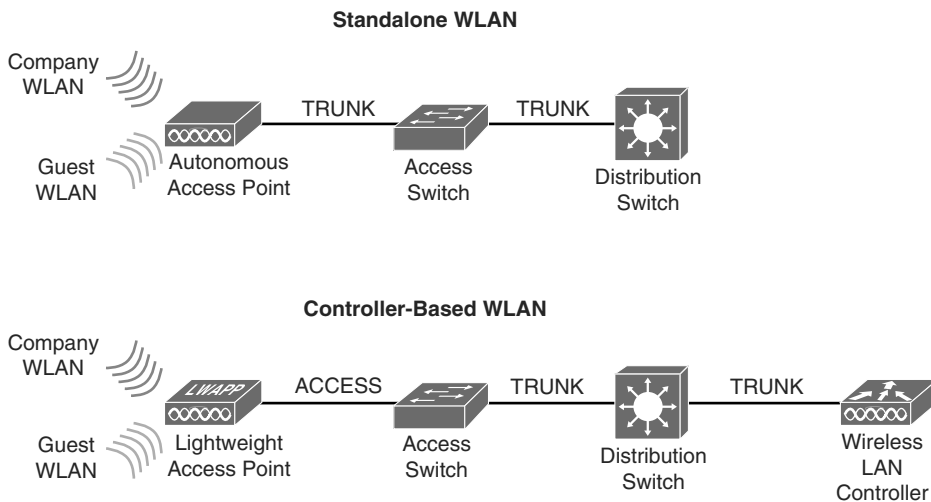


Figure 3-14 *Wireless Configurations Options*

In a controller-based solution, management, control, deployment, and security functions are moved to a central point: the wireless controller, as shown in Figure 3-14. Controllers are combined with lightweight APs that perform only the real-time wireless operation. Controllers can be standalone devices, integrated into a switch, or a WLC can be virtualized.

Both standalone and lightweight APs connect to a switch. It is common that the switch is Power over Ethernet (PoE)-able and so APs get power and data through the Ethernet cable. This makes the wireless network more scalable and easier to manage.

To implement a wireless network, APs and switches need to be configured. APs can be configured directly (autonomous APs) or through a controller (lightweight APs). Either way, configuring APs is a domain of the WLAN specialist. On the switch side, just configure VLANs and trunks on switches to support WLAN.

VLAN Trunking Protocol

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

This section discusses in detail how to plan, implement, and verify VTP in campus networks. The following subsections cover these topics:

- VTP overview
- VTP modes
- VTP versions
- VTP pruning
- VTP authentication
- VTP advertisements
- VTP configuration and verifications
- VTP configuration overwriting
- VTP best practices

VTP Overview

VTP is a Layer 2 protocol that maintains VLAN configuration consistency by managing the additions, deletions, and name changes of VLANs across networks, as shown in Figure 3-15. Switches transmit VTP messages only on 802.1Q or ISL trunks. Cisco switches transmit VTP summary advertisements over the management VLAN (VLAN 1 by default) using a Layer 2 multicast frame every 5 minutes. VTP packets are sent to the destination MAC address 01-00-0C-CC-CC-CC with a logical link control (LLC) code of Subnetwork Access Protocol (SNAP) (AAAA) and a type of 2003 (in the SNAP header).




Figure 3-15 VTP Overview

In Figure 3-15, configurations made to a single VTP server propagate across trunk links to all connected switches in the network in the following manner:

- Step 1.** An administrator adds a new VLAN definition.
- Step 2.** VTP propagates the VLAN information to all switches in the VTP domain.
- Step 3.** Each switch synchronizes its configuration to incorporate the new VLAN data.

VTP domain is one switch or several interconnected switches sharing the same VTP environment but switch can be only in one VTP domain at any time. By default, a Cisco Catalyst switch is in the no-management-domain state or <null> until it receives an advertisement for a domain over a trunk link or until you configure a management domain. Configurations that are made on a single VTP server are propagated across trunk links to all of the connected switches in the network. Configurations will be exchanged if VTP domain and VTP passwords match.

VTP is a Cisco proprietary protocol.

VTP Modes

VTP operates in one of three modes: server, transparent, or client. On some switches, VTP can also be completely disabled. Figure 3-16 shows the brief description of each of the VTP modes.

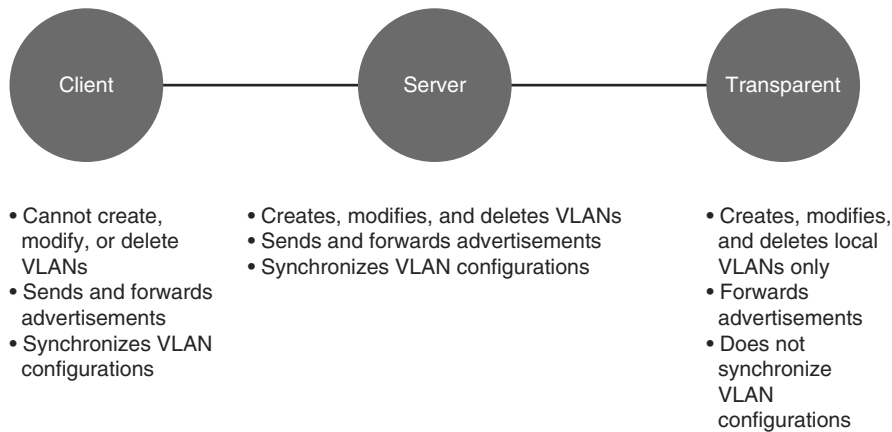


Figure 3-16 VTP Modes and Its Characteristics

The characteristics of the three VTP modes are as follows:

- **Server:** The default VTP mode is server mode, but VLANs are not propagated over the network until a management domain name is specified or learned. When you make a change to the VLAN configuration on a VTP server, the change is propagated to all switches in the VTP domain. VTP messages are transmitted out of all the trunk connections.
- **Transparent:** When you make a change to the VLAN configuration in VTP transparent mode, the change affects only the local switch. The change does not propagate to other switches in the VTP domain. VTP transparent mode does forward VTP advertisements that it receives within the domain.
- **Client:** A VTP client behaves like a VTP server and transmits and receives VTP updates on its trunks, but you cannot create, change, or delete VLANs on a VTP client. VLANs are configured on another switch in the domain that is in server mode.

Note In VTP Version 3, there is a concept of a primary server and a secondary server. VTP Version 3 is not within the scope of this book; for more information, refer to documents on Cisco.com.

In the server, transparent, and client modes, VTP advertisements are received and transmitted as soon as the switch enters the management domain state. In the VTP off mode, switches behave the same as in VTP transparent mode with the exception that VTP advertisements are not forwarded. Off mode is not available in all releases.

By default, Cisco IOS VTP servers and clients save VLANs to the `vlan.dat` file in flash memory, causing them to retain the VLAN table and revision number.

Switches that are in VTP transparent mode display the VLAN and VTP configurations in the **show running-config** command output because this information is stored in the configuration text file. If you perform **erase startup-config** on a VTP transparent switch you will delete its VLANs.

Note The **erase startup-config** command does not affect the `vlan.dat` file on switches in VTP client and server modes. Delete the `vlan.dat` file and reload the switch to clear the VTP and VLAN information. See documentation for your specific switch model to determine how to delete the `vlan.dat` file.

VTP Versions

Cisco Catalyst switches support three different versions of VTP: 1, 2, and 3. It is important to decide which version to use because they are not interoperable. In addition, Cisco recommends running only one VTP version for network stability. This chapter emphasizes VTP Versions 1 and 2 because VTP Version 3 is not the most frequently used version of the VTP.

The default VTP version that is enabled on a Cisco switch is Version 1. If you do need to change the version of VTP in the domain, the only thing that you need to do is to enable it on the VTP server; the change will propagate throughout the network.

VTP Version 2 offers the following features that Version 1 does not:

- **Version-dependent transparent mode:** In VTP Version 1, a VTP transparent network device inspects VTP messages for the domain name and version, and forwards a message only if the version and domain name match. Because only one domain is supported in the Supervisor Engine software, VTP Version 2 forwards VTP messages in transparent mode, without checking the version.
- **Consistency check:** In VTP Version 2, VLAN consistency checks, such as VLAN names and values, are performed. However, this is only done when you enter information through the command-line interface (CLI) or Simple Network Management Protocol (SNMP). Consistency checks are not performed when new information is obtained from a VTP message or when information is read from NVRAM. If the digest on a received VTP message is correct, its information is accepted without consistency checks.
- **Token ring support:** VTP Version 2 supports Token Ring LAN switching and VLANs.
- **Unrecognized type-length-value support:** VTP Version 2 switches propagate received configuration change messages out other trunk links, even if they are not able to understand the message. Instead of dropping the unrecognized VTP message, Version 2 still propagates the information and keeps a copy in NVRAM.

VTP Version 3 brings the following properties:

- **Extended VLAN support:** VTP also can be used to propagate VLANs with numbers 1017–4094 (1006–1017 and 4095–2096 are reserved).
- **Domain name is not automatically learned:** With VTPv2, a factory default switch that receives a VTP message will adapt the new VTP domain name. Because this is a very dangerous behavior, VTPv3 forces manual configuration.
- **Better security:** VTP domain password is secure during transmission and in the switch's database.
- **Better database propagation.** Only the primary server is allowed to update other devices and only one server per VTP domain is allowed to have this role.
- **Multiple Spanning Tree (MST) support:** VTPv3 adds support for propagation of MST instances.

Note VTPv3 is not compatible with VTPv1. VTPv3 is compatible with VTPv2 as long as you are not using it to propagate private or extended VLANs.

Note This book focuses only on VTP Versions 1 and 2 because VTP Version 3 is still not common in the field and is not the focus of the exam.

VTP Pruning

VTP pruning uses VLAN advertisements to determine when a trunk connection is flooding traffic needlessly. By default, a trunk connection carries traffic for all VLANs in the VTP management domain. Commonly, some switches in an enterprise network do not have local ports configured in each VLAN. In Figure 3-17, Switches 1 and 4 support ports statically configured in the red VLAN.

VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices. Figure 3-17 shows a switched network with VTP pruning enabled. The broadcast traffic from Hosts or workstation in red VLAN is not forwarded to Switches 3, 5, and 6, because traffic for the red VLAN has been pruned on the links indicated on Switches 2 and 4.

Regardless of whether you use VTP pruning support, Catalyst switches run an instance of STP for each VLAN. An instance of STP exists for each VLAN even if no ports are active in the VLAN or if VTP pruning removes the VLANs from an interface. As a result, VTP pruning prevents flooded traffic from propagating to switches that do not have members in specific VLANs. However, VTP pruning does not eliminate the switches' knowledge of pruned VLANs.




Figure 3-17 VTP Pruning

VTP Authentication

VTP domains can be secured by using the VTP password feature. It is important to make sure that all the switches in the VTP domain have the same password and domain name; otherwise, a switch will not become a member of the VTP domain. Cisco switches use the message digest 5 (MD5) algorithm to encode passwords in 16-byte words. These passwords propagate inside VTP summary advertisements. In VTP, passwords are case sensitive and can be 8 to 64 characters in length. The use of VTP authentication is a recommended practice.

VTP Advertisements

VTP advertisements are flooded throughout the management domain. VTP advertisements are sent every 5 minutes or whenever there is a change in VLAN configurations. Advertisements are transmitted (untagged) over the native VLAN (VLAN 1 by default) using a multicast frame. A configuration revision number is included in each VTP advertisement. A higher configuration revision number indicates that the VLAN information being advertised is more current than the stored information.

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

being advertised. As shown in Figure 3-18, when the VLAN was added into the database on the VTP server switch, it increased the revision to 4 and advertised the rest of the domain switches that are in client or server VTP mode. However, the switch in transparent mode does not change its revision number or its database.

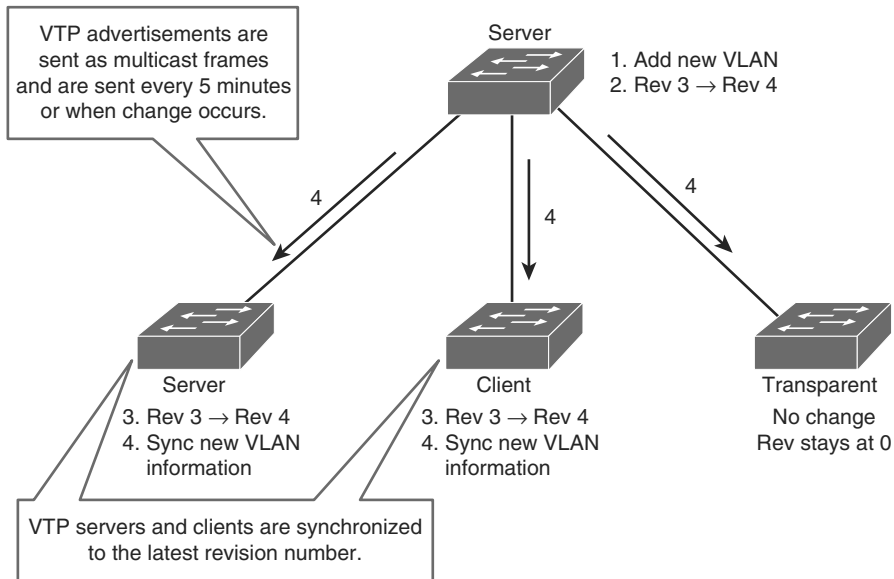


Figure 3-18 VTP Advertisement

The configuration revision number in VTP transparent mode is always zero. Because a VTP-transparent switch does not participate in VTP, that switch does not advertise its VLAN configuration or synchronize its VLAN database upon receipt of a VTP advertisement.

Note In the overwrite process, if the VTP server were to delete all the VLANs and have the higher revision number, the other devices in the VTP domain would also delete their VLANs.

A device that receives VTP advertisements must check various parameters before incorporating the received VLAN information. First, the management domain name and password in the advertisement must match those values that are configured on the local switch. Next, if the configuration revision number indicates that the message was created after the configuration currently in use, the switch incorporates the advertised VLAN information.

Note On many Cisco Catalyst switches, you can change the VTP domain to another name and then change it back to reset the configuration revision number; alternatively, you can change the mode to transparent and then back to the previous setting.

VTP Messages Types

VTP uses various message types for its communication. The subsections that follow describe the message types for VTP.

Summary Advertisements

By default, Catalyst switches issue summary advertisements in 5-minute increments. Summary advertisements inform adjacent Catalysts of the current VTP domain name and the configuration revision number.

When the switch receives a summary advertisement packet, the switch compares the VTP domain name to its own VTP domain name. If the name differs, 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 is higher or equal, the packet is ignored. If it is lower, an advertisement request is sent.

Subset Advertisements

When you add, delete, or change a VLAN in a Catalyst server, the Catalyst server where the changes are made increments the configuration revision and issues a summary advertisement. One or several subset advertisements follow the summary advertisement. A subset advertisement contains a list of VLAN information. If there are several VLANs, more than one subset advertisement can be required to advertise all the VLANs.

Advertisement Requests

A switch needs a VTP advertisement request in these situations:

- The switch has been reset.
- The VTP domain name has been changed.
- The switch has received a VTP summary advertisement with a higher configuration revision than its own.
- Upon receipt of an advertisement request, a VTP device sends a summary advertisement. One or more subset advertisements follow the summary advertisement.

Configuring and Verifying VTP

When creating VLANs, one must decide whether to use VTP in your network. With VTP, changes made on one or more switches propagate automatically to all other switches in the same VTP domain.

Note The domain name and password are case sensitive.

The VTP domain name can be specified or learned. By default, the domain name is <null>. You can specify the password for the VTP management domain. However, if the same password for each switch is not used in the domain, VTP will not function properly. MD5 hashing is used for VTP passwords.

Note The domain name cannot be reset to <null> except if the database is deleted. The domain name can only be reassigned.

To configure VTP, use the topology layout shown in Figure 3-19. In this scenario, Switch 1 will be configured as client, Switch 2 as server, and Switch 3 for transparent mode.

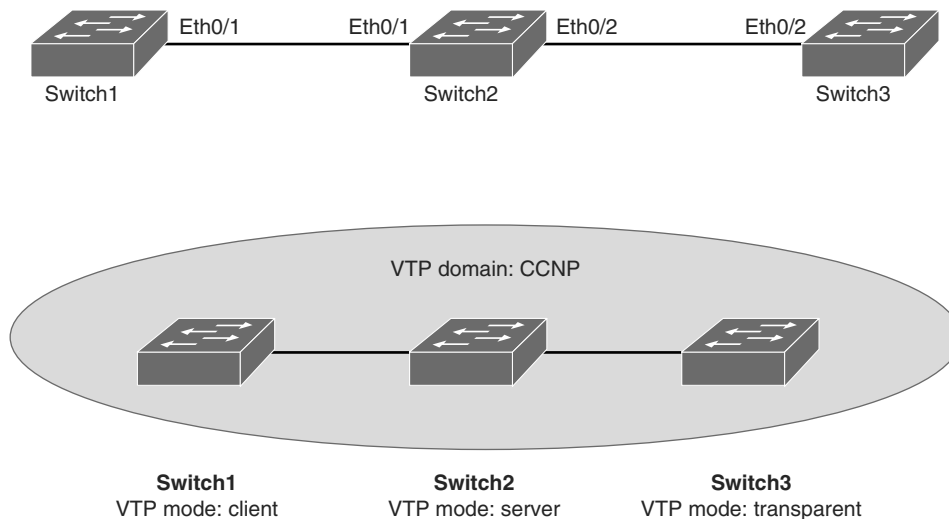


Figure 3-19 VTP Configuration Topology

Complete the following steps to configure the VTP on the switches shown in the topology in Figure 3-19:

- Step 1.** Configure VTP on all the switches, Switch 1 and Switch 3 as client mode where as Switch2 as server mode

```
Switch1 (config)# vtp password Cisco
Switch1 (config)#vtp mode client
Switch1 (config)#vtp domain CCNP
Switch1 (config)#vtp version 1
-----
Switch3 (config)# vtp password Cisco
Switch3 (config)#vtp mode client
Switch3 (config)#vtp domain CCNP
Switch3 (config)#vtp version 1
-----
Switch2 (config)# vtp password Cisco
Switch2 (config)#vtp mode server
Switch2 (config)#vtp domain CCNP
Switch2 (config)#vtp version 1
```

Note By default the switches run VTP Version 1. Vtp Version 2 is not covered in this book. You can look on Cisco.com for more information and capabilities of Version 2.

- Step 2.** Issue the `show vtp status` command on Switch 1 to view the default configuration.

Switch 1 is configured as a VTP client.

Switch 1 is in VTP domain CCNP:

```
Switch1# show vtp status
VTP Version capable          : 1 to 3
VTP version running         : 1
VTP Domain Name              :CCNP
VTP Pruning Mode             : Disabled
VTP Traps Generation         : Disabled
Device ID                    : aabb.cc00.5600
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00

Feature VLAN:
-----
VTP Operating Mode           : Client
Maximum VLANs supported locally : 1005
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

```

As you notice, there are only five default VLANs present on the switch. VLAN 1 and 1002–1005. The VTP revision is 0. Revision 0 means that no changes were made to the VLAN database on this switch so far. Every time that you make a change to the VLAN database (add, remove, modification), the revision will increase by one.

Step 3. Issue the `show vtp status` command on Switch 2.

Switch 2 is configured as VTP server.

Like on Switch 1, only default VLANs are present, VTP revision is 0, and the VTP domain is set to CCNP:

```

Switch2# show vtp status
VTP Version capable        : 1 to 3
VTP version running        : 1
VTP Domain Name            : CCNP
VTP Pruning Mode           : Disabled
VTP Traps Generation       : Disabled
Device ID                  : aabb.cc00.6300
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00
Local updater ID is 0.0.0.0 (no valid interface found)

Feature VLAN:
-----
VTP Operating Mode         : Server
Maximum VLANs supported locally : 1005
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

```

Step 4. Issue the `show vtp status` command on Switch 3.

Switch 3 is configured for VTP transparent mode.

Like on Switch 1 and Switch 2, only default VLANs are present, VTP revision is 0, and the VTP domain is set to CCNP:

```

Switch3# show vtp status
VTP Version capable        : 1 to 3
VTP version running        : 1

```

```

VTP Domain Name           :CCNP
VTP Pruning Mode          : Disabled
VTP Traps Generation      : Disabled
Device ID                  : aabb.cc00.6400
Configuration last modified by 0.0.0.0 at 0-0-00 00:00:00

Feature VLAN:
-----
VTP Operating Mode        : Transparent
Maximum VLANs supported locally : 1005
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

```

Step 5. Create VLAN 10 on Switch 2.

Switch 2 is in VTP server mode. You should be allowed to add VLAN 10 to the Switch 2 database:

```

Switch2# configure terminal
Switch2(config)# vlan 10

```

Note If you try to add a VLAN on a VTP client, you will not be allowed. For example, if you try to add VLAN 5 to Switch 1, you will get the following message:

```
Switch1(config)# vlan 5
```

VTP VLAN configuration not allowed when device is in client mode.

Step 6. Verify VLAN database and VTP status on Switch 2.

Use the commands **show vlan** and **show vtp status**.

Switch 2 now has VLAN 10 in the database:

```
Switch2# show vlan
```

VLAN Name	Status	Ports
1 default	active	Et0/0, Et0/3, Et1/0, Et1/1 Et1/2, Et1/3, Et2/0, Et2/1

```

Et2/2, Et2/3, Et3/0, Et3/1
Et3/2, Et3/3, Et4/0, Et4/1
Et4/2, Et4/3, Et5/0, Et5/1
Et5/2, Et5/3

10 VLAN0010 active
1002 fddi-default act/unsup
1003 token-ring-default act/unsup
1004 fddinet-default act/unsup
1005 trnet-default act/unsup

VLAN Type SAID MTU Parent RingNo BridgeNo Stp BrdgMode Trans1 Trans2
-----
1 enet 100001 1500 - - - - - 0 0
10 enet 100010 1500 - - - - - 0 0
1002 fddi 101002 1500 - - - - - 0 0
1003 tr 101003 1500 - - - - - 0 0
1004 fdnet 101004 1500 - - - ieee - 0 0
1005 trnet 101005 1500 - - - - ibm - 0 0

Primary Secondary Type Ports
-----

```

The revision number increased by one on Switch 2:

```

Switch2# 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 : aabb.cc00.6300
Configuration last modified by 0.0.0.0 at 9-23-13 08:33:48
Local updater ID is 0.0.0.0 (no valid interface found)

Feature VLAN:
-----
VTP Operating Mode : Server
Maximum VLANs supported locally : 1005
Number of existing VLANs : 6
Configuration Revision : 1
MD5 digest : 0xB1 0xBE 0x72 0x49 0x96 0x6D 0x99
0xA4
0xBB 0xB4 0xDC 0x94 0x56 0xD4 0xC2 0x6A
0xBB

```

But the real question now is did changes in Switch 2's database propagate to Switch 1 and Switch 3?

Step 7. Verify changes in VLAN database and VTP status on Switch 1.

Use the commands **show vlan** and **show vtp status**.

Because Switch 1 is a VTP client, VLAN 10 got replicated from Switch 2:

```
Switch1# show vlan
```

```

VLAN Name                Status    Ports
-----
1    default                active   Et0/0, Et0/2, Et0/3, Et1/0
                                   Et1/1, Et1/2, Et1/3, Et2/0
                                   Et2/1, Et2/2, Et2/3, Et3/0
                                   Et3/1, Et3/2, Et3/3, Et4/0
                                   Et4/1, Et4/2, Et4/3, Et5/0
                                   Et5/1, Et5/2, Et5/3
10   VLAN0010               active
1002 fddi-default         act/unsup
1003 token-ring-default   act/unsup
1004 fddinet-default      act/unsup
1005 trnet-default        act/unsup

VLAN Type  SAID      MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
-----
1    enet  100001   1500  -     -     -     -     -     0     0
10   enet  100010   1500  -     -     -     -     -     0     0
20   enet  100020   1500  -     -     -     -     -     0     0
1002 fddi  101002   1500  -     -     -     -     -     0     0
1003 tr   101003   1500  -     -     -     -     srb   0     0

VLAN Type  SAID      MTU   Parent RingNo BridgeNo Stp  BrdgMode Trans1 Trans2
-----
1004 fdnet 101004   1500  -     -     -     ieee -     0     0
1005 trnet 101005   1500  -     -     -     ibm  -     0     0

Primary Secondary Type          Ports
-----

```

The revision number on Switch 1 is now the same as on Switch 2. This indicates that they have an identical VLAN database:

```

Switch1# show vtp status
VTP Version capable      : 1 to 3
VTP version running      : 1

```

```

VTP Domain Name           : CCNP
VTP Pruning Mode          : Disabled
VTP Traps Generation      : Disabled
Device ID                  : aabb.cc00.5600
Configuration last modified by 0.0.0.0 at 9-23-13 08:59:42

```

Feature VLAN:

```

VTP Operating Mode        : Client
Maximum VLANs supported locally : 1005
Number of existing VLANs   : 6
Configuration Revision     : 1
MD5 digest                 : 0xDF 0x2B 0x3B 0x5D 0x0E 0x8E 0x10
0x17
                               0x6D 0xDD 0xE2 0x45 0x7F 0x91 0x95
0x9E

```

Step 8. Verify changes in VLAN database and VTP status on Switch 3.

Use the commands `show vlan` and `show vtp status`.

Switch 3 is in VTP transparent mode. A switch in transparent mode never synchronizes its database to that of the VTP server. In essence, enabling VTP transparent mode disables VTP.

Notice that there is no VLAN 10 on Switch 3:

```
Switch3# show vlan
```

```

VLAN Name                Status    Ports
-----
1    default                active    Et0/0, Et0/2, Et0/3, Et1/0
                                         Et1/1, Et1/2, Et1/3, Et2/0
                                         Et2/1, Et2/2, Et2/3, Et3/0
                                         Et3/1, Et3/2, Et3/3, Et4/0
                                         Et4/1, Et4/2, Et4/3, Et5/0
                                         Et5/1, Et5/2, Et5/3

1002 fddi-default         act/unsup
1003 token-ring-default   act/unsup
1004 fddinet-default      act/unsup
1005 trnet-default        act/unsup

VLAN Type  SAID          MTU    Parent RingNo BridgeNo Stp    BrdgMode Trans1 Trans2
-----
1    enet  100001       1500   -     -     -     -     -     0     0
1002 fddi  101002       1500   -     -     -     -     -     0     0

```



```

10    VLAN0010                active
1002 fddi-default            act/unsup
<... output omitted ...>

```

Switch2# **show vlan**

VLAN Name	Status	Ports
1 default	active	Et0/0, Et0/3, Et1/0, Et1/1 Et1/2, Et1/3, Et2/0, Et2/1 Et2/2, Et2/3, Et3/0, Et3/1 Et3/2, Et3/3, Et4/0, Et4/1 Et4/2, Et4/3, Et5/0, Et5/1 Et5/2, Et5/3
10 VLAN0010	active	
1002 fddi-default	act/unsup	
1003 token-ring-default	act/unsup	
1004 fddinet-default	act/unsup	
1005 trnet-default	act/unsup	

<... output omitted ...>

Switch3# **show vlan**

VLAN Name	Status	Ports
1 default	active	Et0/0, Et0/1, Et0/3, Et1/0 Et1/1, Et1/2, Et1/3, Et2/0 Et2/1, Et2/2, Et2/3, Et3/0 Et3/1, Et3/2, Et3/3, Et4/0 Et4/1, Et4/2, Et4/3, Et5/0 Et5/1, Et5/2, Et5/3
20 VLAN0020	active	
1002 fddi-default	act/unsup	
1003 token-ring-default	act/unsup	
1004 fddinet-default	act/unsup	
1005 trnet-default	act/unsup	

<... output omitted ...>

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

In this example, VLAN 20 is only present in the VLAN database of Switch 3 (the VTP transparent switch, on which you created the VLAN).

Overwriting VTP Configuration (Very Common Issue with VTP)

One of the common issues with VTP is that if you are not careful you can easily wipe out the configuration of the VLAN database across the entire network. Therefore, when a switch is added to a network, it is important that it does not inject spurious information into the domain. Let's review the scenarios illustrated in Figure 3-20, where the SW1 is a VTP server, and SW2 and SW3 are in the VTP client mode. They are all synced to the same configuration revision number '12' and have VLANs 10, 20, 30, and 40. In addition, each switch has hosts connected to multiple VLANs, like SW1 has hosts in VLAN 10 and 20, as depicted in Figure 3-20.

Example 3-10 shows the VTP and VLAN configuration of the switch SW1. Note that SW2 and SW3 would have the similar revision number and VLANs because they are completely synced.

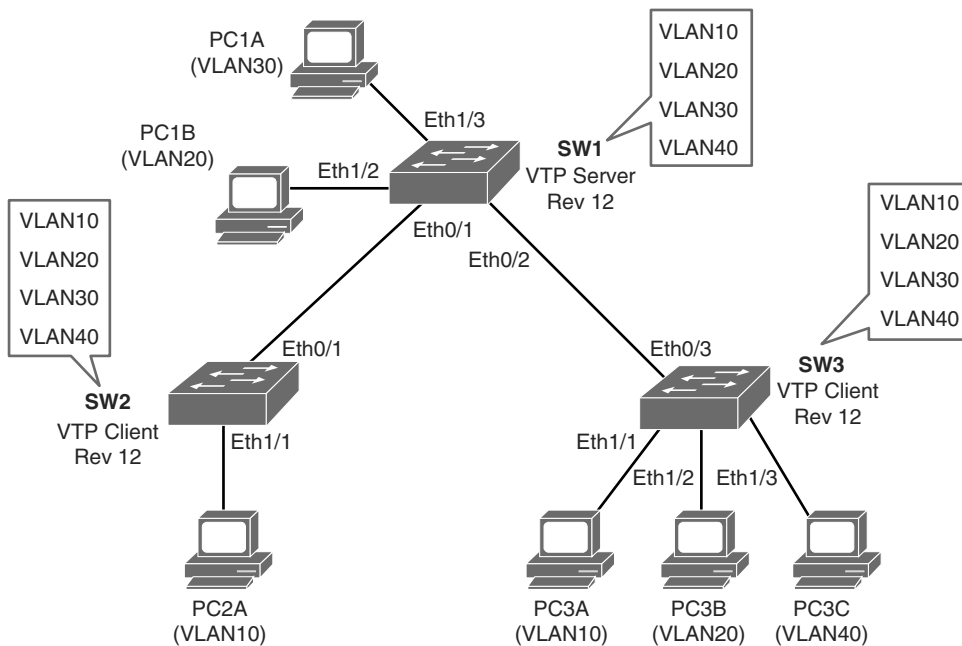


Figure 3-20 *Overwriting VTP Configuration*

Example 3-10 *VLAN and VTP Outputs from Switch SW1*

```

SW1# show vtp status
VTP Version capable      : 1 to 3
VTP version running     : 1
VTP Domain Name         : CCNP
VTP Pruning Mode        : Disabled
VTP Traps Generation    : Disabled
Device ID                : aabb.cc00.5a00
Configuration last modified by 0.0.0.0 at 9-24-13 07:33:33
Local updater ID is 0.0.0.0 (no valid interface found)

Feature VLAN:
-----
VTP Operating Mode      : Server
Maximum VLANs supported locally : 1005
Number of existing VLANs : 9
Configuration Revision  : 12
MD5 digest              : 0x11 0x31 0x4F 0x6A 0x96 0x0D 0xB6 0xB9
                        : 0xAE 0xF4 0xD4 0x85 0x4D 0x58 0xC8 0x4D

SW1# show vlan

```

VLAN Name	Status	Ports
1 default	active	Et0/0, Et1/0, Et2/0, Et2/1 Et2/2, Et2/3, Et3/0, Et3/1 Et3/2, Et3/3, Et4/0, Et4/1 Et4/2, Et4/3, Et5/0, Et5/1 Et5/2, Et5/3
10 VLAN0010	active	
20 VLAN0020	active	Et1/2
30 VLAN0030	active	Et1/3
40 VLAN0040	active	

Now assume that SW2 failed and was replaced by another new switch in the closet, as shown in Figure 3-21.

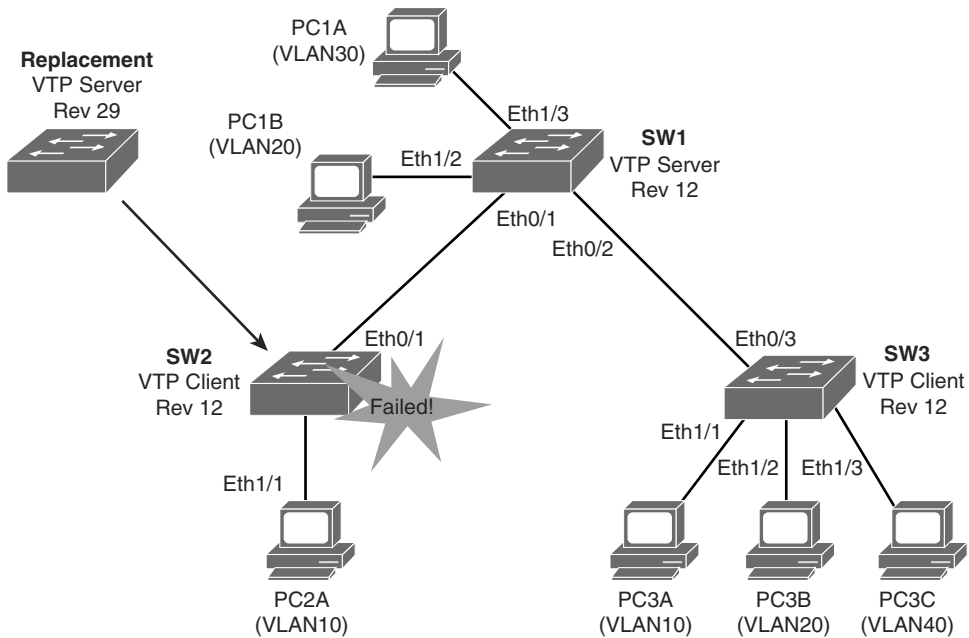


Figure 3-21 *Overwriting VTP Configuration: Switch Failure*

However, the network administrator forgot to erase the configuration and VLAN database.

The replacement switch has the same VTP domain name configured as the other two switches. The VTP revision number on the replacement switch is 29, higher than the revision on the other two switches.

Note The VTP revision number is stored in NVRAM and is not reset if you erase switch configuration and reload it.

Example 3-11 shows the output of VLANs and VTP on the new replacement switch to show the revision number and its VLAN database.

Example 3-11 *VTP and VLAN Output from the New Replacement Switch*

```

Replacement# show vtp status
VTP Version capable      : 1 to 3
VTP version running     : 1
VTP Domain Name         : CCNP
VTP Pruning Mode        : Disabled
VTP Traps Generation    : Disabled
Device ID                : aabb.cc00.5a00
  
```

```

Configuration last modified by 0.0.0.0 at 9-24-13 08:15:44
Local updater ID is 0.0.0.0 (no valid interface found)

Feature VLAN:
-----
VTP Operating Mode           : Server
Maximum VLANs supported locally : 1005
Number of existing VLANs     : 10
Configuration Revision       : 29
MD5 digest                   : 0x29 0xF2 0x1F 0xA5 0x41 0x44 0x04 0xAC
                               0x08 0x3B 0x9A 0x2C 0x73 0x8A 0xA2 0xBD
! The replacement switch does not have VLANs 20, 30, and 40 in its database.
Replacement# show vlan

VLAN Name                Status    Ports
-----
1    default                active    Et0/0, Et0/1, Et0/2, Et1/0
                               Et2/0, Et2/1, Et2/2, Et2/3
                               Et3/0, Et3/1, Et3/2, Et3/3
                               Et4/0, Et4/1, Et4/2, Et4/3
                               Et5/0, Et5/1, Et5/2, Et5/3
10   VLAN0010                 active    Et1/1
11   VLAN0011                 active
22   VLAN0022                 active
33   VLAN0033                 active
44   VLAN0044                 active
<... output omitted ...>

```

Because SW2 has a higher revision number, SW1 and SW3 will sync to the latest revision.

The consequence is that VLANs 20, 30, and 40 no longer exist on SW1 and SW2. This leaves the clients that are connected to ports belonging to nonexistent VLANs without connectivity, as shown in Figure 3-22.

Example 3-12 shows the output of **show vtp status** and **show vlan** of the SW1 and SW3 to show how the VLAN database is updated with new switch database.

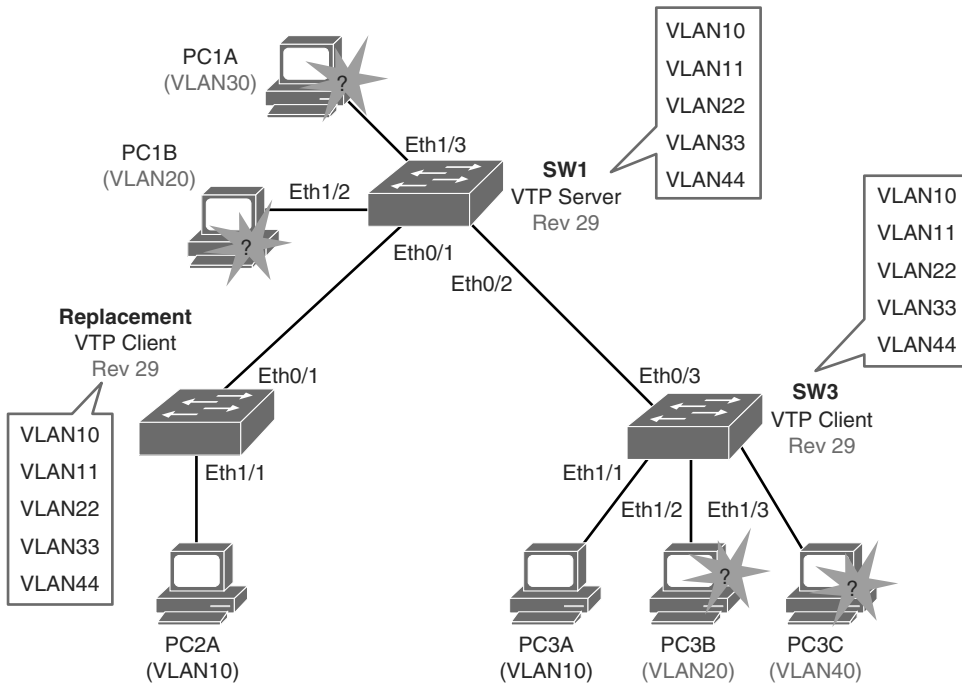


Figure 3-22 VTP Overwriting Advertisement

Example 3-12 Show VTP Status and Show VLAN Outputs from SW1 and SW3

```

SW1# show vtp status
VTP Version capable      : 1 to 3
VTP version running     : 1
VTP Domain Name         : CCNP
VTP Pruning Mode        : Disabled
VTP Traps Generation    : Disabled
Device ID                : aabb.cc00.5900
Configuration last modified by 0.0.0.0 at 9-24-13 08:15:44
Local updater ID is 0.0.0.0 (no valid interface found)

Feature VLAN:
-----
VTP Operating Mode      : Server
Maximum VLANs supported locally : 1005
Number of existing VLANs : 10
Configuration Revision  : 29
    
```

```
MD5 digest                : 0x29 0xF2 0x1F 0xA5 0x41 0x44 0x04 0xAC
                          : 0x08 0x3B 0x9A 0x2C 0x73 0x8A 0xA2 0xBD
```

```
SW1# show vlan
```

VLAN Name	Status	Ports
1 default	active	Et0/0, Et1/0, Et2/0, Et2/1 Et2/2, Et2/3, Et3/0, Et3/1 Et3/2, Et3/3, Et4/0, Et4/1 Et4/2, Et4/3, Et5/0, Et5/1 Et5/2, Et5/3
10 VLAN0010	active	
11 VLAN0011	active	
22 VLAN0022	active	
33 VLAN0033	active	
44 VLAN0044	active	

```
<... output omitted ...>
```

```
SW3# show vtp status
```

```
VTP Version capable      : 1 to 3
VTP version running      : 1
VTP Domain Name          : CCNP
VTP Pruning Mode         : Disabled
VTP Traps Generation     : Disabled
Device ID                 : aabb.cc00.5600
Configuration last modified by 0.0.0.0 at 9-24-13 08:15:44
```

```
Feature VLAN:
```

```
-----
VTP Operating Mode       : Client
Maximum VLANs supported locally : 1005
Number of existing VLANs : 10
Configuration Revision   : 29
MD5 digest                : 0x29 0xF2 0x1F 0xA5 0x41 0x44 0x04 0xAC
                          : 0x08 0x3B 0x9A 0x2C 0x73 0x8A 0xA2 0xBD
```

```
SW3# show vlan
```

VLAN Name	Status	Ports
1 default	active	Et0/0, Et0/2, Et1/0, Et2/0 Et2/1, Et2/2, Et2/3, Et3/0 Et3/1, Et3/2, Et3/3, Et4/0 Et4/1, Et4/2, Et4/3, Et5/0

```

Et5/1, Et5/2, Et5/3
10 VLAN0010 active Et1/1
11 VLAN0011 active
22 VLAN0022 active
33 VLAN0033 active
44 VLAN0044 active
<... output omitted ...>

```

Also, when the new switch is added with a VTP client with a higher revision number, it can cause the same havoc as a switch with the VTP server, as discussed earlier. The VTP client, as a general rule, just listens to VTP advertisements from VTP servers, and it does not do its own advertisements. However, when the switch with the VTP client is added to a network, it will send a summary advertisement from its own stored database. If the VTP client gets an inferior advertisement from the VTP server, it will assume it has better, more current information. The VTP client will now send out advertisements with a higher revision number. The VTP server and all directly connected VTP clients will accept these as more current. This will not only delete the old VLANs but also can add new VLANs into the network and create network instability.

Remember the revision configuration and how to reset it each time a new switch is inserted so that it does not bring down the entire network. Following are some of the key points:

- Avoid, as much as possible, VLANs that span the entire network.
- The VTP revision number is stored in NVRAM and is not reset if you erase the switch configuration and reload it. To reset the VTP revision number to zero, use the following 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.

Best Practices for VTP Implementation

VTP is often used in a new network to facilitate the implementation of VLANs. However, as the network grows larger, this benefit can turn into a liability. If a VLAN is deleted by accident on one server, it is deleted throughout the network. If a switch that already has a VLAN database defined is inserted into the network, it can hijack the VLAN database by deleting added VLANs. Because of this, it is the recommended practice to configure all switches to transparent VTP mode and manually add VLANs as needed, especially in a larger campus network. VTP configuration is usually good for small environments.

Implementing EtherChannel in a Switched Network

In networks where resources may be located far from where users might need them, some links between switches or between switches and servers become heavily solicited. The speed of these links can be increased, but only to a certain point. EtherChannel is a technology that allows you to circumvent the bandwidth issue by creating logical links that are made up of several physical links.

This section examines the benefits of EtherChannel and the various technologies available to implement it and also the types of EtherChannel protocol. In addition, it explains how to configure Layer 2 EtherChannels and how to load balance traffic between physical links inside a given EtherChannel bundle. EtherChannels can also operate in a Layer 3 mode, but this is discussed later in Chapter 5. The following topics are discussed in detail in the following subsections:

- The need for EtherChannel technology
- Port aggregation negotiation protocols
- Configuration steps for bundling interfaces into a Layer 2 EtherChannel
- Configuring EtherChannel
- Changing EtherChannel load-balancing behavior
- How EtherChannel load-balancing works
- The role of EtherChannel Guard

The Need for EtherChannel

Any-to-any communications of intranet applications, such as video to the desktop, interactive messaging, Voice over IP (VoIP), and collaborative whiteboard use, are increasing the need for scalable bandwidth within the core and at the edge of campus networks. At the same time, mission-critical applications call for resilient network designs. With the wide deployment of faster switched Ethernet links in the campus, users need to either aggregate their existing resources or upgrade the speed in their uplinks and core to scale performance across the network backbone.

In Figure 3-23, traffic coming from several VLANs at 100 Mbps aggregate on the access switches at the bottom and need to be sent to distribution switches in the middle. Obviously, bandwidth larger than 100 Mbps must be available on the link between two switches to accommodate the traffic load coming from all the VLANs. A first solution is to use a faster port speed, such as 1 or 10 Gbps. As the speed increases on the VLANs links, this solution finds its limitation where the fastest possible port is no longer fast enough to aggregate the traffic coming from all VLANs. A second solution is to multiply the numbers of physical links between both switches to increase the overall speed of the switch-to-switch communication. A downside of this method is that there must be a strict consistency in each physical link configuration. A second issue is that spanning tree may block one of the links, as shown in Figure 3-23.




Figure 3-23 *Network Without EtherChannel*

EtherChannel is a technology that was originally developed by Cisco as a LAN switch-to-switch technique of grouping several Fast or Gigabit Ethernet ports into one logical channel. This technology has many benefits:

- It relies on the existing switch ports. There is no need to upgrade the switch-to-switch link to a faster and more expensive connection.
- Most of the configuration tasks can be done on the EtherChannel interface instead of on each individual port, thus ensuring configuration consistency throughout the switch-to-switch links.
- Load balancing is possible between the links that are part of the same EtherChannel. Depending on the hardware platform, you can implement one or several methods, such as source-MAC to destination-MAC or source-IP to destination-IP load balancing across the physical links.

Keep in mind that the logic of EtherChannel is to increase the speed between switches, as illustrated in Figure 3-24. This concept was extended as the EtherChannel technology became more popular, and some hardware nonswitch devices support link aggregation into an EtherChannel link. In any case, EtherChannel creates a one-to-one relationship. You can create an EtherChannel link between two switches or between an EtherChannel-enabled server and a switch, but you cannot send traffic to two different switches through the same EtherChannel link. One EtherChannel link always connects the same two devices only. The individual EtherChannel group member port configuration must be consistent on both devices. EtherChannel technology only bundles ports of the same

type. On a Layer 2 switch, EtherChannel is used to aggregate access ports or trunks. For example, if the physical ports of one side are configured as trunks, the physical ports of the other side must also be configured as trunks. Each EtherChannel has a logical port channel interface. A configuration that is applied to the port channel interface affects all physical interfaces that are assigned to that interface. (Such commands can be STP commands or commands to configure a Layer 2 EtherChannel as a trunk or an access port.)

Note Using new technologies like VSS (Virtual Switching System) and vPC (Virtual Port Channel), a port channel can be created across two aggregation switches from the same access layer to provide better redundancy.




Figure 3-24 Network with EtherChannel

Keep in mind that EtherChannel creates an aggregation that is seen as one logical link. When several EtherChannel bundles exist between two switches, spanning tree may block one of the bundles to prevent redundant links. When spanning tree blocks one of the redundant links, it blocks one EtherChannel, thus blocking all the ports belonging to this EtherChannel link. Where there is only one EtherChannel link, all physical links in the EtherChannel are active because spanning tree sees only one (logical) link. If one link in EtherChannel goes down, the bandwidth of the EtherChannel will be automatically updated, and thus the STP cost will change as well.

Note On Layer 3 switches, you can convert switched ports to routed ports. You can also create EtherChannel links on Layer 3 links. Layer 3 port channels are discussed in more detail in Chapter 5.

Note Also, with technologies like VSS and VPC (which are discussed in more detail in Chapter 9, “High Availability,”) you can create the EtherChannel between the access layer and two different aggregation switches.

EtherChannel Mode Interactions

EtherChannel can be established using one of the following three mechanisms, as shown in Figure 3-25:

- **LACP:** IEEE’s negotiation protocol
- **PAgP:** Cisco’s negotiation protocol
- **Static persistence:** No negotiation protocol

LACP			PAgP			Static Persistence	
	Active	Passive		Desirable	Auto		On
Active	Yes	Yes	Desirable	Yes	Yes	On	
Passive	Yes	No	Auto	Yes	No	On	Yes

Figure 3-25 *EtherChannel Modes Interactions*

LACP

Link Aggregation Control Protocol (LACP) is part of an IEEE specification (802.3ad) that allows several physical ports to be bundled together to form a single logical channel. LACP allows a switch to negotiate an automatic bundle by sending LACP packets to the peer. Because LACP is an IEEE standard, you can use it to facilitate EtherChannels in mixed-switch environments. LACP checks for configuration consistency and manages link additions and failures between two switches. It ensures that when EtherChannel is created, all ports have the same type of configuration speed, duplex setting, and VLAN information. Any port modification after the creation of the channel will also change all the other channel ports.

LACP packets are exchanged between switches over EtherChannel-capable ports. Port capabilities are learned and compared with local switch capabilities. LACP assigns roles to EtherChannel’s ports. The switch with the lowest system priority is allowed to make decisions about what ports actively participate in EtherChannel. Ports become active

according to their port priority. A lower number means higher priority. Commonly up to 16 links can be assigned to an EtherChannel, but only 8 can be active at a time. Nonactive links are placed into a standby state and are enabled if one of the active links goes down.

The maximum number of active links in an EtherChannel varies between switches.

These are the LACP modes of operation:

- **Active:** Enable LACP
- **Passive:** Enable LACP only if an LACP device is detected

The following are some additional parameters that you can use when configuring LACP:

- **System priority:** Each switch running LACP must have a system priority. The system priority can be specified automatically or through the CLI. The switch uses the MAC address and the system priority to form the system ID.
- **Port priority:** Each port in the switch must have a port priority. The port priority can be specified automatically or through the CLI. The port priority and the port number form the port identifier. The switch uses the port priority to decide which ports to put in standby mode when a hardware limitation prevents all compatible ports from aggregating.
- **Administrative key:** Each port in the switch must have an administrative key value, which can be specified automatically or through the CLI. The administrative key defines the capability of a port to aggregate with other ports, determined by these factors: the port's physical characteristics, such as data rate, duplex capability, and point-to-point or shared medium.

All the preceding options of LACP are optional to configure. Usually, defaults are the best to use. To configure any of these options, refer to your configuration guide.

PAGP

Port Aggregation Protocol (PAgP) provides the same negotiation benefits as LACP. PAgP is a Cisco proprietary protocol, and it will work only on Cisco devices. PAgP packets are exchanged between switches over EtherChannel-capable ports. Neighbors are identified and capabilities are learned and compared with local switch capabilities. Ports that have the same capabilities are bundled together into an EtherChannel. PAgP forms an EtherChannel only on ports that are configured for identical VLANs or trunking. PAgP will automatically modify parameters of the EtherChannel if one of the ports in the bundle is modified. For example, if configured speed, duplex, or VLAN of a port in a bundle is changed, PAgP reconfigures that parameter for all ports in the bundle. PAgP and LACP are not compatible.

These are the following two PAgP modes of operation:

- **Desirable:** Enable PAgP
- **Auto:** Enable PAgP only if a PAgP device is detected

Note Negotiation with either LACP or PAgP introduces overhead and delay in initialization. As an alternative, you can statically bundle links into an EtherChannel. This method introduces no delays but can cause problems if not properly configured on both ends.

Layer 2 EtherChannel Configuration Guidelines

Before implementing EtherChannel in a network, plan the following steps necessary to make it successful:

- The first step is to identify the ports that you will use for the EtherChannel on both switches. This task helps identify any issues with previous configurations on the ports and ensures that the proper connections are available.
- Each interface should have the appropriate protocol identified (PAgP or LACP), have a channel group number to associate all the given interfaces with a port group, and be configured whether negotiation should occur.
- After the connections are established, make sure that both sides of the EtherChannel have formed and are providing aggregated bandwidth.

Follow these guidelines and restrictions when configuring EtherChannel interfaces:

- **EtherChannel support:** All Ethernet interfaces on all modules support EtherChannel, with no requirement that interfaces be physically contiguous or on the same module.
- **Speed and duplex:** Configure all interfaces in an EtherChannel to operate at the same speed and in the same duplex mode. Also, if one interface in the bundle is shut down, it is treated as a link failure, and traffic will traverse other links in the bundle.
- **VLAN match:** All interfaces in the EtherChannel bundle must be assigned to the same VLAN or be configured as a trunk.
- **Range of VLANs:** An EtherChannel supports the same allowed range of VLANs on all the interfaces in a trunking Layer 2 EtherChannel.

If the allowed range of VLANs is not the same, the interfaces do not form an EtherChannel, even when set to auto or desirable mode. For Layer 2 EtherChannels, either assign all interfaces in the EtherChannel to the same VLAN or configure them as trunks.

- **STP path cost:** Interfaces with different STP port path costs can form an EtherChannel as long as they are compatibly configured. Setting different STP port path costs does not, by itself, make interfaces incompatible for the formation of an EtherChannel.
- **Port channel versus interface configuration:** After you configure an EtherChannel, any configuration that you apply to the port channel interface affects the EtherChannel. Any configuration that you apply to the physical interfaces affects only the specific interface that you configured.

Note If you do not specify any protocol, it will be static binding. That topic is not within the scope of this book.

EtherChannel Load-Balancing Options

EtherChannel load balances traffic across links in the bundle. However, traffic is not necessarily distributed equally among all the links.

Frames are forwarded over an EtherChannel link that is based on results of a hashing algorithm. Options that switch can use to calculate this hash depends on the platform.

Table 3-6 shows the comment set of options for EtherChannel load balancing.

Table 3-6 *EtherChannel Load-Balancing Options*

Hash Input Code	Hash Input Decision	Switch Model
dst-ip	Destination IP address	All models
dst-mac	Destination MAC address	All models
src-dst-ip	Source and destination IP address	All models
src-dst-mac	Source and destination MAC address	All models
src-ip	Source IP address	All models
src-mac	Source MAC address	All models
src-port	Source port number	4500, 6500
dst-port	Destination port number	4500, 6500
src-dst-port	Source and destination port number	4500, 6500

To verify load-balancing options available on the device, use the **port-channel load-balance ?** global configuration command.

The hash algorithm calculates a binary pattern that selects a link within the EtherChannel bundle to forward the frame.

Note Default configuration can differ from switch to switch, but commonly the default option is src-dst-ip. It is not possible to have different load-balancing methods for different EtherChannels on one switch. If the load-balancing method is changed, it is applicable for all EtherChannels.

If only one address or port number is hashed, a switch looks at one or more low-order bits of the hash value. The switch then uses those bits as index values to decide over which links in the bundle to send the frames.

If two or more addresses or port numbers are hashed, a switch performs an XOR operation.

A four-link bundle uses a hash of the last 2 bits. A bundle of eight links uses a hash of the last 3 bits.

Table 3-7 shows results of an XOR on a two-link bundle, using the source and destination addresses.

Table 3-7 XOR for Two-Link EtherChannels

Example IP Addresses	IPs in Binary	XOR Result	Forward Frame over Link with Index
Source: 192.168.1.2 Destination: 192.168.1.4	Source: ...xxxxx0 Destination: ...xxxxx0	...xxxxx0	0
Source: 172.16.1.20 Destination: 172.16.1.21	Source: ...xxxxx0 Destination: ...xxxxx1	...xxxxx1	1
Source: 192.168.1.1 Destination: 192.168.1.2	Source: ...xxxxx1 Destination: ...xxxxx0	...xxxxx1	1
Source: 10.1.1.101 Destination: 10.1.1.103	Source: ...xxxxx1 Destination: ...xxxxx1	...xxxxx0	0

A conversation between two devices is sent through the same EtherChannel link because the two endpoint addresses stay the same. Only when a device talks to several other devices does traffic get distributed evenly over the links in the bundle.

When one pair of hosts has a much greater volume of traffic than the other pair, one link will be much more utilized than others. To fix the imbalance, consider using some other load-balancing mechanisms, such as source and destination port number, that will redistribute traffic much differently.

If most of the traffic is IP, it makes sense to load balance according to IP addresses or port numbers. For non-IP traffic, the hash uses MAC addresses to calculate the path.

To achieve the optimal traffic distribution, always bundle an even number of links. For example, if you use four links, the algorithm will take the last 2 bits. These 2 bits mean four indexes: 00, 01, 10, and 11. Each link in the bundle will get assigned one of these indexes. If you bundle only three links, the algorithm still needs to use 2 bits to make decisions. One of the three links in the bundle will be used more than the other two.

With four links, the algorithm strives to load balance traffic in a 1:1:1:1 ratio. A three-link algorithm strives to load balance traffic in a 2:1:1 ratio.

Note You cannot control the port that a particular flow uses. You can only influence the load balance with a frame distribution method that results in the greatest variety.

Configuring EtherChannel in a Switched Network

This section shows you how to configure the Layer 2 EtherChannel and explains its load-balancing behavior. Configure a port channel between SW1 and SW2 shown in Figure 3-26.




Figure 3-26 *EtherChannel Configuration Topology*

Table 3-8 shows device information.

Table 3-8 *Device Information*

Device	IP Address	Interface	Neighbor	Interface on the Neighbor
PC 1	172.16.1.101/24	Ethernet 0/0	Switch 1	Ethernet 0/1
PC 2	172.16.1.102/24	Ethernet 0/0	Switch 1	Ethernet 0/2
PC 3	172.16.1.203/24	Ethernet 0/0	Switch 2	Ethernet 0/1
Switch 1	<i>No IP address</i>	Ethernet 1/1	Switch 2	Ethernet 0/2
Switch 1	<i>No IP address</i>	Ethernet 1/2	Switch 2	Ethernet 0/3

EtherChannel Configuration and Load Balancing

Complete the following steps to configure EtherChannel on Switch 1. Switch 2 has EtherChannel preconfigured.

- Step 1.** On Switch 1, configure the two ports that connect to Switch 2 to use channel group 1 and LACP active mode:

```
Switch1# configure terminal
Switch1(config)# interface range Ethernet 1/1-2
Switch1(config-if-range)# channel-group 1 mode active
Creating a port-channel interface Port-channel 1
```

Now the two interfaces are bundled into channel group 1. Because you chose the **active** keyword, LACP will work as the negotiation protocol. Because Switch 2 has its ports bundled and activated for LACP passive mode, EtherChannel should come right up.

```
%LINEPROTO-5-UPDOWN: Line protocol on Interface Port-channel1,
changed state to up
```

Notice that by assigning the two ports to a port channel, the switch has created a port channel 1 interface.

Issue the **show ip interface brief** command. Port channel 1 will be listed as just another interface at the very bottom of the list.

- Step 2.** Enter interface configuration mode for the newly created port channel interface and configure it for trunk mode using dot1Q:

```
Switch1(config)# interface port-channel 1
Switch1(config-if)# switchport trunk encapsulation dot1q
Switch1(config-if)# switchport mode trunk
```

The configuration applied to the port channel will also reflect on physical interfaces that are bundled into that port channel. You can investigate the running configuration and see that EtherChannel 1/1 and EtherChannel 1/2 both have had the trunking configuration applied.

Step 3. On Switch 1, enter the `show etherchannel summary` command:

```
Switch1# show etherchannel summary
Flags: D - down          P - bundled in port-channel
       I - stand-alone  s - suspended
       H - Hot-standby (LACP only)
       R - Layer3       S - Layer2
       U - in use       f - failed to allocate aggregator

       M - not in use, minimum links not met
       u - unsuitable for bundling
       w - waiting to be aggregated
       d - default port

Number of channel-groups in use: 1
Number of aggregators:          1

Group  Port-channel  Protocol    Ports
-----+-----+-----+-----+-----
1      Po1(SU)          LACP        Et1/1(P)  Et1/2(P)
```

Group 1 port channel is a Layer 2 EtherChannel that is in use (SU flag). The negotiation protocol in use is LACP, and the ports bundled (notice the P flag) are Ethernet 1/1 and Ethernet 1/2.

If a port comes up but cannot join the port channel, it is denoted with an I flag (for “independent”).

Step 4. Enter the `show etherchannel load-balance` command to verify which information EtherChannel uses to load balance traffic:

```
Switch1# show etherchannel load-balance
EtherChannel Load-Balancing Configuration:
      src-dst-ip

EtherChannel Load-Balancing Addresses Used Per-Protocol:
Non-IP: Source XOR Destination MAC address
      IPv4: Source XOR Destination IP address
      IPv6: Source XOR Destination IP address
```

Notice that the default configuration for load balancing is `src-dst-ip`. This means the source and destination IP address are used for hash input.

Step 5. For testing how much traffic goes over each link, as shown in Figure 3-27, clear interface counters on Switch 1 using the `clear counters` command:

```
Switch1# clear counters
```

```
Clear "show interface" counters on all interfaces [confirm] [Enter]
```

By clearing the counters, you are setting up to test how much traffic goes over each link.




Figure 3-27 *EtherChannel Load-Balancing Configuration Option*

Step 6. Perform an extended ping from PC 1 to PC 3:

```
PC1# ping
Protocol [ip]: [Enter]
Target IP address: 172.16.1.203
Repeat count [5]: 10000
Datagram size [100]: 1500
Timeout in seconds [2]: [Enter]
Extended commands [n]: [Enter]
Sweep range of sizes [n]: [Enter]
Type escape sequence to abort.
Sending 10000, 1500-byte ICMP Echos to 172.16.1.203, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
<... output omitted ...>
```

In the next step, you check over which interface all the traffic went.

Step 7. Verify counters on Switch 1 for both interfaces:

```
Switch1# show interface ethernet 1/1 | i packets output
      10094 packets output, 15146494 bytes, 0 underruns
Switch1# show interface ethernet 1/2 | i packets output
      13 packets output, 1664 bytes, 0 underruns
```

Notice that most of the traffic went over the Ethernet 1/1 interface.

But what about if you ping from PC 2 to PC 3? Will traffic go over the other interface in EtherChannel bundle?

Step 8. Clear interface counters on Switch 1 using the `clear counters` command:

```
Switch1# clear counters
Clear "show interface" counters on all interfaces [confirm] [Enter]
```

Step 9. Perform an extended ping from PC 2 to PC 3:

```
PC2# ping
Protocol [ip]: [Enter]
Target IP address: 172.16.1.203
Repeat count [5]: 10000
Datagram size [100]: 1500
Timeout in seconds [2]: [Enter]
Extended commands [n]: [Enter]
Sweep range of sizes [n]: [Enter]
Type escape sequence to abort.
Sending 10000, 1500-byte ICMP Echos to 172.16.1.203, timeout is 2
seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
<... output omitted ...>
```

Step 10. Verify counters on Switch 1 for both interfaces:

```
Switch1# show interface ethernet 1/1 | i packets output
      29 packets output, 2201 bytes, 0 underruns
Switch1# show interface ethernet 1/2 | i packets output
     10003 packets output, 15140537 bytes, 0 underruns
```

So, with the ping from PC 1 to PC 3, traffic went over Ethernet 1/1. With the ping from PC 2 to PC 3, traffic went over Ethernet 1/2. This is for the default load-balancing method that takes destination and source IP address for calculating the hash.

Step 11. Change the load-balancing behavior on Switch 1 from `src-dst-ip` to `dst-ip`:

```
Switch1(config)# port-channel load-balance dst-ip
How will traffic get distributed over the two links now?
```

Step 12. Verify that the load-balancing behavior has changed:

```
Switch1# show etherchannel load-balance
EtherChannel Load-Balancing Configuration:
    dst-ip

EtherChannel Load-Balancing Addresses Used Per-Protocol:
Non-IP: Source XOR Destination MAC address
  IPv4: Source XOR Destination IP address
  IPv6: Source XOR Destination IP address
```

Step 13. Clear the interface counters on Switch 1 by using the `clear counters` command:

```
Switch1# clear counters
Clear "show interface" counters on all interfaces [confirm] [Enter]
```

Step 14. Perform an extended ping from PC 1 to PC 3:

```
PC1# ping
Protocol [ip]: [Enter]
Target IP address: 172.16.1.203
Repeat count [5]: 10000
Datagram size [100]: 1500
Timeout in seconds [2]: [Enter]
Extended commands [n]: [Enter]
Sweep range of sizes [n]: [Enter]
Type escape sequence to abort.
Sending 10000, 1500-byte ICMP Echos to 172.16.1.203, timeout is 2
seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
<... output omitted ...>
```

Step 15. Verify the counters on Switch 1 for both interfaces:

```
Switch1# show interface ethernet 1/1 | i packets output
    32 packets output, 2108 bytes, 0 underruns
Switch1# show interface ethernet 1/2 | i packets output
    10002 packets output, 15140188 bytes, 0 underruns
```

The majority of the traffic went over the Ethernet 1/2 port.

Step 16. Clear the interface counters on Switch 1 by using the `clear counters` command:

```
Switch1# clear counters
Clear "show interface" counters on all interfaces [confirm] [Enter]
```

Step 17. Perform an extended ping from PC 2 to PC 3:

```
PC2# ping
Protocol [ip]: [Enter]
Target IP address: 172.16.1.203
Repeat count [5]: 10000
Datagram size [100]: 1500
Timeout in seconds [2]: [Enter]
Extended commands [n]: [Enter]
Sweep range of sizes [n]: [Enter]
Type escape sequence to abort.
Sending 10000, 1500-byte ICMP Echos to 172.16.1.203, timeout is 2
seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
<... output omitted ...>
```

Step 18. Verify counters on Switch 1 for both interfaces:

```
Switch1# show interface ethernet 1/1 | i packets output
31 packets output, 2329 bytes, 0 underruns
Switch1# show interface ethernet 1/2 | i packets output
10004 packets output, 15140597 bytes, 0 underruns
```

Now that the load balancing is based on destination IP, the behavior has changed. Because the only input information for calculation of the hash is destination IP address, it does not matter whether you ping PC 3 from PC 1 or PC 2. In both cases, the hash function will be the same, and traffic will go over the same link (in this example, Ethernet ½).

Note The default method of load balancing usually works for most scenarios, but you can change it based on the traffic needs in the network.

EtherChannel Guard

The EtherChannel Guard feature is used to detect EtherChannel misconfigurations between the switch and a connected device.

EtherChannel misconfiguration occurs when the channel parameters do not match on both sides of the EtherChannel, resulting in the following message:

```
%PM-SP-4-ERR_DISABLE: channel-misconfig error detected on Po3, putting E1/3 in
err-disable state
```

The EtherChannel Guard feature can be enabled by using the **spanning-tree etherchannel guard misconfig** global configuration command.

However, EtherChannel Guard is enabled by default. To verify whether it is configured, use the **show spanning-tree summary** command, as demonstrated in Example 3-13.

Example 3-13 *Show VTP Status and Show VLAN outputs from SW1 and SW3*

```
Switch1# show spanning-tree summary
Switch is in pvst mode
Root bridge for: VLAN0001
Extended system ID          is enabled
Portfast Default            is disabled
PortFast BPDU Guard Default is disabled
Portfast BPDU Filter Default is disabled
Loopguard Default          is disabled
EtherChannel misconfig guard is enabled
<...output omitted...>
```

Study Tips

- VLAN provides logical grouping of the hosts to restrict the broadcast domain.
- VLANs are usually categorized into local and end-to-end VLANs, and each has its own pros and cons.
- With the help of trunking, VLANS can be easily extended over a single physical link.
- ISL and 802.1Q are two trunking protocols, with dot1Q the industry standard.
- Dot1Q frames insert 4 bytes and recalculate the CRC.
- Native VLAN is not encapsulated in dot1Q trunking, and it is important to have same native VLAN on both sides of the switches.
- VTP is used to distribute VLAN databases. It has multiple versions and modes. VTP works in server, client, and transparent mode.
- Any switch with a higher revision number can overwrite the VLAN database. Insert the new switch with caution and follow the recommended steps.
- EtherChannel is a technology that was originally developed by Cisco as a LAN switch-to-switch technique of grouping several Fast or Gigabit Ethernet ports into one logical channel.
- PagP and LACP are the two main protocols for EtherChannel.
- For EtherChannel, it is highly recommended to use the even number of ports in the channel to have better load balancing.

Summary

In review, a VLAN is a logical grouping of switch ports that connects nodes of nearly any type, regardless of physical location. VLAN segmentation is based on traffic flow patterns. A VLAN is usually defined as an end-to-end VLAN or a local VLAN. An end-to-end VLAN spans the entire switched network, whereas a local VLAN is limited to the switches in the building access and building distribution submodules. The creation of a VLAN implementation plan depends on the business and technical requirements.

Furthermore, a trunk is a Layer 2 point-to-point link between networking devices that can carry the traffic of multiple VLANs. ISL and 802.1Q are the two trunking protocols that connect two switches. The 802.1Q protocol is an open standard protocol also used for VLAN trunking.

VTP is used to distribute and synchronize information about VLANs configured throughout a switched network. VTP pruning helps to stop flooding of unnecessary traffic on trunk links. VTP configuration sometimes needs to be added to small network deployments, whereas VTP transparent mode is usually privileged for larger networks. When configuring VLANs over several switches, ensure that the configuration is compatible throughout switches in the same domain.

To increase bandwidth and provide redundancy, use EtherChannel by aggregating individual, similar links between switches. EtherChannel can be dynamically configured between switches using either the Cisco proprietary PAgP or the IEEE 802.3ad LACP. EtherChannel is configured by assigning interfaces to the EtherChannel bundle and configuring the resulting port channel interface. EtherChannel load balances traffic over all the links in the bundle. The method that is chosen directly impacts the efficiency of this load-balancing mechanism.

Review Questions

Use the questions here to review what you learned in this chapter. The correct answers are found in Appendix A, “Answers to Chapter Review Questions.”

1. True or False: It is important to have the same native VLAN on both switch link partners for ISL trunking.
2. True or False: The Cisco Catalyst 6500 supports up to 1024 VLANs in the most recent software releases.
3. True or False: When removing the native VLAN from a trunk port, CDP, PAgP, and DTP, use the lowest-numbered VLAN to send traffic.
4. True or False: In VTP client mode, switches can add and delete VLANs.
5. True or False: Token Ring support is available in VTP Version 1.

Questions 6 through 8 are based on the configuration in Example 3-14.

Example 3-14 Configuration Example for Questions 6 Through 8

```
Catalyst6500-IOS# show run interface gigabitEthernet 3/9
Building configuration...
Current configuration : 137 bytes !
interface GigabitEthernet3/9
mtu 9216
no ip address
switchport
switchport access vlan 5
switchport trunk encapsulation dot1q
end
```

6. If the interface in Example 3-14 negotiates trunking, what would be the native VLAN?
 - a. VLAN 1
 - b. VLAN 5
 - c. VLAN 9216
 - d. No native VLAN if the port negotiated trunking
7. Under what condition can the interface in Example 3-14 negotiate ISL trunking?
 - a. If the port is a member of an EtherChannel.
 - b. If the link partner defaults to ISL trunking for negotiated ports.
 - c. If the link partner is configured for trunking in the on mode.
 - d. The interface cannot negotiate trunking because it is configured statically for 802.1Q trunking.
8. Which statement is true for the configuration of the interface in Example 3-14?
 - a. The interface is a member of VLAN 5 and may negotiate to a trunk port.
 - b. The interface may negotiate to an ISL trunk with a native VLAN of 5.
 - c. The interface may negotiate to an 802.1Q trunk and operate with a native VLAN of 1.
 - d. The interface will not negotiate to a trunk port because it is configured in access VLAN 5.
 - e. If a host workstation is connected to the interface, it must be configured for trunking.

Questions 9 through 11 are based on the configuration in Example 3-15.

Example 3-15 *Configuration Example for Questions 9 Through 11*

```
svs-san-6509-2# show interfaces gigabitEthernet 3/9 switchport
Name: Gi3/9
Switchport: Enabled
Administrative Mode: dynamic auto
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 2 (VLAN0002)
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL
```

9. What is the trunk native VLAN based on in Example 3-15?
 - a. VLAN 1
 - b. VLAN 2
 - c. VLAN 5
 - d. No Native VLAN if the port negotiated trunking
10. Based on Example 3-15, what statement is true if the link partner (peer switch) is configured for the dynamic trunking mode?
 - a. The interface cannot negotiate to a trunk port because it is configured for dot1Q encapsulation.
 - b. The interface cannot negotiate to a trunk port because the native VLAN and access VLANs are mismatched.
 - c. The interface can negotiate to a trunk port if the peer is configured for the dynamic desirable trunking mode.
 - d. The interface can negotiate to a trunk port if access VLAN is the same on both sides.

11. What is the interface's access mode VLAN in Example 3-15?
 - a. VLAN 1
 - b. VLAN 2
 - c. VLAN 5
 - d. VLAN 1001
12. How does implementing VLANs help improve the overall performance of the network?
 - a. By isolating problem employees
 - b. By constraining broadcast traffic
 - c. By grouping switch ports into logical communities
 - d. By forcing the Layer 3 routing process to occur between VLANs
13. What are the advantages of using local VLANs over end-to-end VLANs? (Choose two.)
 - a. Eases management
 - b. Eliminates the need for Layer 3 devices
 - c. Allows for a more deterministic network
 - d. Groups users by logical commonality
 - e. Keeps users and resources on the same VLAN
14. Which prompt indicates that you are in VLAN configuration mode of Cisco IOS?
 - a. Switch#
 - b. Switch(vlan)#
 - c. Switch(config)#
 - d. Switch(config-vlan)#
15. Which switch port mode unconditionally sets the switch port to access mode regardless of any other DTP configurations?
 - a. Access
 - b. Nonegotiate
 - c. Dynamic auto
 - d. Dynamic desirable

- 16.** What information is contained in the FCS of an ISL-encapsulated frame?
 - a.** CRC calculation
 - b.** Header encapsulation
 - c.** ASIC implementation
 - d.** Protocol independence
- 17.** 802.1Q uses an internal tagging mechanism, where a tag is inserted after the _____ field.
 - a.** Type
 - b.** SA
 - c.** Data
 - d.** CRC
- 18.** Which command correctly configures a port with ISL encapsulation in Cisco IOS?
 - a.** Switch(config-if)# `switchport mode trunk isl`
 - b.** Switch(config-if)# `switchport mode encapsulation isl`
 - c.** Switch(config-if)# `switchport trunk encapsulation isl`
 - d.** Switch(config-if)# `switchport mode trunk encapsulation isl`
- 19.** Which command correctly sets the native VLAN to VLAN 5?
 - a.** `switchport native vlan 5`
 - b.** `switchport trunk native 5`
 - c.** `switchport native trunk vlan 5`
 - d.** `switchport trunk native vlan 5`
- 20.** If the Layer 2 interface mode on one link partner is set to dynamic auto, a trunk will be established if the link partner is configured for which types of interface modes in Cisco IOS? (Choose two.)
 - a.** Trunk
 - b.** Access
 - c.** Nonegotiate
 - d.** Dynamic auto
 - e.** Dynamic desirable

- 21.** What is the default VTP mode for a Catalyst switch?
- a.** Client
 - b.** Access
 - c.** Server
 - d.** Transparent
- 22.** When is a consistency check performed with VTP Version 2?
- a.** When information is read from NVRAM
 - b.** When the digest on a received VTP message is correct
 - c.** When new information is obtained from a VTP message
 - d.** When you enter new information through the CLI or SNMP
- 23.** Which command correctly sets the VTP version to version 1 in Cisco IOS global configuration mode?
- a.** `vtp v1-mode`
 - b.** `vtp v2-mode`
 - c.** `no vtp version`
 - d.** `no vtp version 2`
- 24.** Which of the following are valid VTP Version 1 and 2 modes? (Choose all that apply.)
- a.** Primary server mode
 - b.** Server mode
 - c.** Client mode
 - d.** Transparent mode
- 25.** After you complete the VTP configuration, which command should you use to verify your configuration?
- a.** `show vtp status`
 - b.** `show vtp counters`
 - c.** `show vtp statistics`
 - d.** `show vtp status counters`

- 26.** What command might correct a problem with incorrect VTP passwords?
- a.** `password vtp 0`
 - b.** `clear vtp password`
 - c.** `clear password vtp`
 - d.** `vtp password password_string`
- 27.** True or False: The EtherChannel would come up if one side of the EtherChannel mode is set to auto and the other to on.
- 28.** Which of the following solutions are provided by EtherChannel? (Choose two.)
- a.** Provide redundancy
 - b.** Help to overcome bandwidth limitation
 - c.** Because of EtherChannel, can transmit more than one VLAN over the links between switches
 - d.** Can limit the broadcast to the local switches
- 29.** Which statement identifies network benefits provided by VLANs?
- a.** VLANs allow you to group stations without regard to the physical location of the users.
 - b.** VLANs help to isolate problem employees.
 - c.** VLANs reduce the impact of network problems.
 - d.** VLANs can transmit frames to all ports in all VLANs.
- 30.** Match each command to its explanation.
- a.** `Switch(config-if)# switchport voice vlan vlan-id`
 - b.** `Switch(config-if)# switchport mode access`
 - c.** `Switch(config-if)# switchport access vlan vlan-id`
 - d.** `Switch(config-if)# switchport trunk native vlan vlan-id`
 - e.** `Switch(config-if)# switchport trunk allowed vlan add vlan-id`
 - f.** `Switch(config-if)# switchport mode trunk`
- 1.** Configures the port to be assigned only to a single VLAN
 - 2.** Configures the port to be assigned to multiple VLANs
 - 3.** Configures a VLAN to be added to trunk port
 - 4.** Configures a native VLAN for the trunk
 - 5.** Configures a port to be a part of voice VLAN
 - 6.** Configures a port to be a part of data VLAN

- 31.** How can you reset VTP revision number on a switch? (Choose two.)
- a.** Set the switch to transparent mode and then to server mode.
 - b.** Set the switch to client mode and then to server mode.
 - c.** Change the VTP domain name to a nonexistent VTP domain and then back to the original name.
 - d.** Reload the switch.
- 32.** Which statement about transparent VTP mode is true?
- a.** Creates, modifies, and deletes VLANs on all switches in VTP domain
 - b.** Creates, modifies, and deletes local VLANs only
 - c.** Does not forward advertisements to other switches in VTP domain
 - d.** Synchronizes VLAN configurations from other switches in VTP domain
- 33.** What is the correct command for configuring load balancing on an EtherChannel link?
- a.** Switch(config)# **channel-group number load-balance *method***
 - b.** Switch(config-if)# **channel-group number load-balance *method***
 - c.** Switch(config-if)# **port-channel number load-balance *method***
 - d.** Switch(config)# **port-channel load-balance *method***
- 34.** Which of the following EtherChannel modes does not send or receive any negotiation frames?
- a.** Passive
 - b.** Active
 - c.** On
 - d.** Desirable auto

This page intentionally left blank

Index

Numerics

- 80/20 rule, 47
- 802.1Q trunking, 43-44, 49-52
 - 802.1Q Ethernet frames, 51
 - architectural advantages over ISL, 51
 - CRC, 51
 - native VLAN, 52
 - VLAN ranges, 54
- 802.1X authentication, 316-318
 - configuring, 318-319

A

- AAA (authentication, authorization, and accounting), 305-315
 - accounting, 306, 314-315
 - authentication, 305-308
 - RADIUS*, 309-310
 - TACACS+*, 310-311
 - authorization, 306, 313-314
 - configuring, 311
- access layer, 12-13
 - Layer 3 switching, 17-19
- Access mode, 53
- accounting, 306, 314-315
- accurate time, need for, 320
- ACI (Application Centric Infrastructure), 411
- ACLs (access control lists), 26
 - PACLs, 424-425
 - VACLs, 448-451
- active routers (HSRP), 252
- active state (HSRP), 253
- advertisement requests (VTP), 77
- advertisements, VTP, 75-77
- aggressive mode (UDLD), 358
- aligning HSRP with STP topology, 254-255
- alternate ports (RSTP), 134
- answers to chapter review questions, 469-472
- application intelligence, 3
- ARP (Address Resolution Protocol), 63
- ARP spoofing, mitigating, 437-443
- assigning switch ports to VLANs, 56-57
- attacks, 416-417
- authentication, 305-308
 - 802.1X, 316-319
 - GLBP, 295
 - HSRP, configuring, 271-272
 - IP SLA, configuring, 382
 - RADIUS*, 309-310
 - TACACS+*, 310-311
 - VRRP, configuring, 279-280
 - VTP, 75

authorization, 306, 313-314
 autostate exclude command, 220-221
 availability
 FHRPs, 247
 need for, 248-249
 HA, 393-394
 local switching architectures, need for, 394-395
 redundant switch supervisors, 401-405
 StackWise, 395-397
 VSS, 397-401
 HSRP, 250-253
 authentication, configuring, 271-272
 comparing with VRRP, 275
 configuring, 255-263
 interface tracking, 266-268
 load sharing, 263-265
 object tracking, 268-271
 state transition, 253-254
 timers, configuring, 272-273
 versions, 274

B

baby giant frames, 52
 backbone, 14-17
 BackboneFast, 154-156
 backup ports (RSTP), 134
 bandwidth
 comparing hubs and switches, 2
 load balancing, EtherChannel, 100-102
 best practices
 for Cisco switch security configuration, 411-412
 for MST, 194-196
 for SNMP, 339-340
 for STP stability mechanisms, 175-178
 for trunking, 65-66
 for UDLD, 170-171
 for VLANs, 65-66
 for VTP implementation, 93

blocking state (STP), 129
 BPDU Filter, 159-161
 BPDU Guard, 158-159
 BPDUs (bridge protocol data units), 121, 124
 inferior BPDUs, 155
 bridges, 2-3
 broadcast domains, 3
 in VLANs, 43
 BVI (bridge virtual interface), 206

C

CAM (content addressable memory) table, 5
 campus networks, 2
 Cisco enterprise campus architecture, 19-20
 hierarchical network design, 10-19
 access layer, 12-13
 core layer, 14-17
 distribution layer, 13-14
 ECMP, 14
 FHRP, 14
 monitoring
 RSPAN, 370-371
 SPAN, 368-371
 switches
 broadcast domains, 3
 CAM table, 5
 collision domains, 24
 CRC, 24
 features, 3
 functions of, 5-6
 MLS, 8
 security, 410-411
 store-n-forward mode, 24
 TCAM, 26
 trunking, 7
 vulnerabilities, 415-417
 trunking, 49-54
 configuring, 64-65

- VLANs, 6, 42-48
 - configuring*, 61-64
 - end-to-end VLANs*, 44-45
 - inter-VLAN routing*, 43
 - local VLANs*, 45-46
 - ports*, 43
 - segmentation*, 44
- vulnerabilities, 412-419
 - MAC flooding attacks*, 417-419
 - rogue access*, 412-415
- capabilities of access layer**, 12-13
- Catalyst switches.** *See also* Cisco switches
 - Catalyst 2960-X series, 23
 - Catalyst 3850-X series, 23
 - Catalyst 6500 switches, 23
 - Catalyst 6800-X series, 23
 - load balancing, 32-33
 - PoE, 360-364
 - route caching, 30-31
 - SDM templates, 364-368
 - show mac address-table command, 27-28
 - topology-based switching, 31-33
 - VLANs, implementing, 43
- CDP (Cisco Discovery Protocol)**, 52, 352
 - comparing with LLDP, 352
- CEF (Cisco Express Forwarding)**, 30, 31-33. *See also* topology-based switching
- centralized switching**, 33
- chapter review questions**
 - answers to chapter review questions, 469-472
 - Chapter 2, 35-39
 - Chapter 3, 110-117
 - Chapter 4, 200-201
 - Chapter 5, 242-246
 - Chapter 8, 385-391
- Cisco ASA (Adaptive Security Appliance)**, 410
- Cisco enterprise campus architecture**, 19-20
- Cisco IOS**
 - global configuration mode, creating VLANs, 55-56
 - trunking modes, 53-54
- Cisco STP Toolkit**, 151-152
 - BackboneFast, 154-156
 - BPDU Filter, 159-161
 - BPDU Guard, 158-159
 - FlexLinks, 171-175
 - Loop Guard, 164-166
 - PortFast, 156-158
 - Root Guard, 161-163
 - stability mechanisms recommendations, 175-178
 - UDLD, 166-171
 - UplinkFast, 153-154
- Cisco switches**
 - Catalyst switches, 22-23
 - Catalyst 2960-X series*, 23
 - Catalyst 3850-X series*, 23
 - fixed configuration switches, 23
 - modular switches, 23
 - Nexus switches, 22
 - security configuration best practices, 411-412
 - trunking modes, 53-54
 - WLAN solutions, 69-70
- Cisco WLAN solutions**, 69-70
- clear counters command**, 105
- client mode (VTP)**, 72
- collision domains**, 24
- commands**
 - autostate exclude command, 220-221
 - clear counters command, 105
 - erase startup-config command, 73
 - monitor session commands, 373
 - port-channel load-balance command, 101
 - show interfaces command, 59
 - show ip protocol command, 229-230
 - show ip route command, 229-230
 - show mac address-table command, 27-28
 - show running-config command, 59
 - show spanning-tree command, 140-142
 - show standby command, 259
 - show vlan command, 57-59, 62-63, 81-87
 - show vtp status command, 79-81

- spanning-tree backbonefast command, 156
 - switchport host command, 56-57
 - traceroute command, 210
 - community ports, 454**
 - comparing**
 - end-to-end and local VLANs, 46-47
 - hierarchical and flat network design, 10-11
 - HSRP and VRRP, 276
 - hubs and switches, 2
 - Layer 2 and Layer 3 switches, 24
 - RADIUS and TACACS+, 308
 - STP and RSTP port states, 135-136
 - configuration revision numbers, 75-77**
 - configuring**
 - 802.1X authentication, 318-319
 - AAA, 311
 - BackboneFast, 154-156
 - BPDU Filter, 159-161
 - BPDU Guard, 158-159
 - DAI, 440-443
 - DHCP
 - DHCP relay, 239*
 - in multilayer switched networks, 233-238*
 - options, 239-240*
 - EtherChannel, 99-100
 - Layer 3 EtherChannel interfaces, 226-229*
 - in switched networks, 102-108*
 - FlexLinks, 171-175
 - GLBP, 285-294
 - HSRP, 255-263, 271-272
 - timers, 272-273*
 - inter-VLAN routing
 - using external router, 206-211*
 - using SVI, 212-214*
 - IP SLA, 377-379
 - IPSG, 435-439
 - LLDP, 353-355
 - Loop Guard, 164-166
 - MHSRP, 263-265
 - MST, 185-190
 - path cost, 192-193*
 - port priority, 193*
 - PoE, 363-364
 - port security, 420-422
 - PortFast, 156-158
 - PVLANs, 454-456
 - RADIUS, 311-312
 - Root Guard, 161-163
 - RSPAN, 372-374
 - SNMP, 340-344
 - SPAN, 371-372
 - storm control, 427-429
 - STP, 140-151
 - SVI, 221-222
 - TACACS+, 312-313
 - trunks, 64-65
 - UDLD, 166-171, 358-360
 - UplinkFast, 153-154
 - VACLs, 450-451
 - VLANs, 61-64
 - VRRP, 276-280
 - VTP, 78-87
 - control plane, 29**
 - controller-based WLAN solution, 69-70**
 - convergence**
 - in access layer, 13
 - STP, 148-149
 - core layer, 14-17**
 - need for, 20-22
 - CRC (cyclic redundancy check), 24**
 - 802.1Q trunking, 51
 - creating VLANs, 55-56**
 - CST (Common Spanning Tree), 121**
-
- ## D
- DAI (Dynamic ARP Inspection), configuring, 440-443**
 - data centers, 2, 10**
 - Nexus switches, 22
 - Data field (802.1Q frames), 51**

- Data field (Ethernet frames), 5
 - debugging STP topology events, 148-149
 - default behaviors, UDLD, 359
 - deleting VLANs, 56
 - designated ports, 123
 - election process, 128-129
 - RSTP, 134
 - designing networks
 - 80/20 rule, 47
 - VLANs, 65-66
 - Dest field (802.1Q frames), 51
 - destination address field (Ethernet frames), 5
 - deterministic traffic flow, 48
 - devices
 - bridges, 2-3
 - broadcast domains, 3
 - hubs, 2
 - IP phones, voice VLAN, 67-69
 - planes of operation, 28-29
 - switches
 - CAM table*, 5
 - CRC*, 24
 - fixed configuration switches*, 23
 - full duplex mode*, 24
 - functions of*, 5-6
 - MLS*, 8
 - modular switches*, 23
 - port channels*, 7-8
 - store-n-forward mode*, 24
 - TCAM*, 26
 - trunking*, 7
 - VLANs*, 6
 - DFA (Dynamic Fabric Allocation), 19
 - DHCP (Dynamic Host Configuration Protocol), 48, 231-232
 - DHCP relay, configuring, 239
 - messages, 236
 - multilayer switched networks, configuring, 233-238
 - options, configuring, 239-240
 - spoofing attacks, 430-434
 - DHCP snooping, 432-434
 - direct topology changes (STP), 132
 - disabled ports (RSTP), 134
 - disabled state (4.237), 130
 - discarding state (RSTP), 136
 - discovery protocols
 - CDP, 352
 - comparing with LLDP*, 352
 - LLDP, 352-356
 - configuring*, 353-355
 - implementation properties*, 353
 - neighbor discovery*, 355-356
 - TLVs*, 353
 - displaying MAC address table information, 60-61
 - distributed hardware forwarding, 28-29
 - distributed switching, 33
 - distribution layer, 13-14
 - ECMP, 14
 - FHRP, 14
 - DTP (Dynamic Trunking Protocol), 53-54
 - duplex mismatches, troubleshooting, 196-197
 - Dynamic auto mode, 53
 - Dynamic desirable mode, 53
- ## E
-
- ECMP (equal-cost multipathing), 14
 - election process
 - designated ports, 128-129
 - HSRP preemption, enabling, 258-263
 - root bridge election, 124-126
 - root port election, 126-127
 - end-to-end VLANs, 44-45
 - comparing with local VLANs, 46-47
 - enterprise campus architecture, implementing with VLANs, 47-48
 - erase startup-config command, 73
 - EtherChannel, 7-8, 41, 94, 225-226
 - Layer 2 configuration, 99-100
 - Layer 3 EtherChannel interfaces, configuring, 226-229

- load balancing, 100-102
 - links, bundling*, 102
 - XOR operation*, 101
- mode interactions
 - LACP*, 97-98
 - PAgP*, 98-99
- need for, 94-97
- switched network configuration, 102-108
- troubleshooting, 108-109
- EtherChannel Guard**, 108-109
- Ethernet**
 - frame format, 4-5
 - PoE, 360-364
 - components*, 362
 - power classes*, 362-363
 - standards*, 362
- example troubleshooting plan,
 - troubleshooting inter-VLAN routing, 223-225
- extended pings, performing, 105-106
- Extended System ID, 185
- external routers
 - advantages of, 211
 - inter-VLAN routing, 206-211

F

- fast switching. *See* route caching
- FCS field (802.1Q frames), 51
- features of switches, 3
- FHRPs (first-hop routing protocols), 14, 247
 - GLBP, 282-300
 - authentication*, 295
 - comparing with HSRP*, 283-284
 - configuring*, 285-294
 - load balancing options*, 294-295
 - tracking*, 296-300
 - virtual forwarder states*, 285
 - virtual gateway states*, 285
 - weighting options*, 298-300

- HSRP
 - authentication, configuring*, 271-272
 - interface tracking*, 266-268
 - load sharing*, 263-265
 - object tracking*, 268-271
 - timers, configuring*, 272-273
- need for, 248-249
- VRRP, 274-281
 - configuring*, 276-280
 - tracking*, 280-281

fields

- of 802.1Q Ethernet frames, 51
- of BPDUs, 124
- of Ethernet frames, 4-5
- of show vlan command, 57
- fixed configuration switches, 23
- Flags field (BPDUs), 124
- flat enterprise campus networks, 10-11
- FlexLinks, 171-175
- Forward Delay field (BPDUs), 124
- forward delay time (STP), 148
- forwarding plane, 29
- forwarding state (RSTP), 136
- forwarding state (STP), 130
- frame corruption (STP), troubleshooting, 197-198
- frames
 - 802.1Q Ethernet, 51
 - baby giant frames, 52
 - BPDUs, fields, 124
 - Ethernet, 4-5
 - rewrites, 28
 - tagging, 7
- full-duplex mode, 24
- functions of switches, 5-6

G

- GLBP (Gateway Load Balancing Protocol)**, 14, 282-300
 - authentication, 295
 - comparing with HSRP, 283-284

- configuring, 285-294
- load balancing options, 294-295
- tracking, 296-300
- virtual forwarder states, 285
- virtual gateway states, 285
- weighting options, 298-300

global configuration mode, creating VLANs, 55-56

H

HA (high availability), 393-394

- in access layer, 119
- local switching architectures, need for, 394-395
- redundant switch supervisors, 401-405
 - supervisor redundancy modes, 402-405*
- StackWise, 395-397
- VSS, 397-401

half-duplex mode, 24

Hello Time field (BPDUs), 124

hello time (STP), 148

hierarchical network design, 10-19

- access layer, 12-13
 - Layer 3 switching, 17-19*
- Cisco enterprise campus architecture, 19-20
- core layer, 14-17
 - need for, 20-22*
- distribution layer, 13-14
 - ECMP, 14*
 - FHRP, 14*
- VLANs, mapping, 47-48

high availability, in access layer, 12-13

HSRP (Hot Standby Routing Protocol), 14, 250-253

- aligning with STP topology, 254-255
- authentication, configuring, 271-272
- comparing with VRRP, 275
- configuring, 255-263
- interface tracking, 266-268
- load sharing, 263-265

- MHSRP, configuring, 263-265
- object tracking, 268-271
- preemption, enabling, 258-263
- state transition, 253-254
- timers, configuring, 272-273
- versions, 274

hubs, 2

I

identity-based networking, 316-319

IEEE 802.1Q trunking, 49-52

- native VLAN, 52
- VLAN ranges, 54

IEEE 802.1w standard. *See* RSTP (Rapid Spanning Tree Protocol)

IEEE 802.3 standard, Ethernet frame format, 4-5

indirect topology changes (STP), 132

inferior BPDUs, 155

ingress queues, 25

initial state (HSRP), 253

insignificant topology changes (STP), 133

interface tracking, 266-268

- VRRP, 280-281

inter-VLAN routing, 43, 204-206

- MLS, 217-220
- router-on-a-stick, 206
- troubleshooting, 222-225
- using external router, 206-211
- using routed ports, 214-222
- using SVI, 212-214

IP addressing, DHCP, 231-232

- DHCP relay, configuring, 239
- messages, 236
- multilayer switched networks, configuring, 233-238

IP phones, 42

- voice VLAN, 67-69

IP SLA (Service Level Agreement), 374-377

- authentication, configuring, 382
- configuring, 377-379

- ICMP Echo test example, 375
- responders, 377, 379-381
- sources, 377
- time stamps, 381-382
- UDP jitter example, 383-384
- IPSG (IP Source Guard), 435-439
- ISL (Inter-Switch Link), 49
 - comparing with 802.1Q trunking, 51
- isolated ports, 454
- IST (internal spanning tree), 183

J-K

- jitter example of IP SLA, 383-384

- key chains, configuring MD5
 - authentication for HSRP, 272

L

- LACP (Link Aggregation Control Protocol), 97-98
- LAN switching, 2
- Layer 1
 - hubs, 2
- Layer 2 EtherChannel configuration, 99-100
- Layer 2 switching
 - ACLs, 26
 - ingress queues, 25
 - MAC address forwarding, 24-25
 - MAC table, 26
 - QoS, 26
- Layer 3 redundancy
 - GLBP, 282-300
 - authentication*, 295
 - comparing with HSRP*, 283-284
 - configuring*, 285-294
 - load balancing options*, 294-295
 - tracking*, 296-300
 - virtual forwarder states*, 285
 - virtual gateway states*, 285
 - weighting options*, 298-300

- VRRP, 274-281
 - comparing with HSRP*, 275
 - configuring*, 274-280
 - millisecond timers*, 275
 - tracking*, 280-281

- Layer 3 switching, 26-27
 - in the access layer, 17-19
 - centralized switching, 33
 - comparing with Layer 2 switching, 24
 - distributed switching, 33
 - EtherChannel interfaces, configuring, 226-229
 - QoS, 27

layered network design

- Cisco enterprise campus architecture, 19-20
- hierarchical network design, 11-22
 - access layer*, 12-13
 - core layer*, 14-17
 - distribution layer*, 13-14

- learning state (RSTP), 136

- learning state (STP), 129

- Len/Etype field (802.1Q frames), 51

- Length/Type field (Ethernet frames), 5

- links, RSTP, 138-139

- listen state (HSRP), 253

- listening state (STP), 129

- LLC (Logical Link Control), 70

- LLDP (Link Layer Discovery Protocol), 352-356

- comparing with CDP, 352

- configuring, 353-355

- implementation properties, 353

- neighbor discovery, 355-356

- TLVs, 353

load balancing, 32-33

- EtherChannel, 100-102

- configuring*, 102-108

- links, bundling*, 102

- XOR operation*, 101

- GLBP, 282-300

- load sharing, HSRP, 263-265

local switching architectures, need for, 394-395

local VLANs, 45-46

comparing with end-to-end VLANs, 46-47

Loop Guard, 164-166

leveraging with UDLD, 360

M

MAC addresses, 4

displaying MAC table information, 60-61

Ethernet frame format, 4-5

Layer 2 switching, 24-25

MAC flooding attacks, 417-419

management plane, 28

mandatory TLVs (LLDP), 353

manual system clock configuration, 320-322

mapping VLANs to hierarchical networks, 47-48

Max Age field (BPDUs), 124

max age time (STP), 148

MD5 authentication, configuring for HSRP, 272

memory, TCAM, 26

SDM templates, 364-368

system resource configuration, 367-368

Message Age field (BPDUs), 124

Message Type field (BPDUs), 124

messages

BPDUs, 121, 124

inferior BPDUs, 155

DHCP, 236

VTP, 77

MHSRP (Multigroup HSRP), configuring, 263-265

MIB (Management Information Base), 337

millisecond timers, VRRP, 275

mitigating spoofing attacks

ARP spoofing, 437-443

DHCP spoofing, 430-434

IPSG, 435-439

MLS (multilayer switching), 8, 26-27

comparing with Layer 2 switches, 24

inter-VLAN routing, 217-220

planes of operation, 28-29

QoS, 27

modes, VTP, 71-73

modifying STP behavior, 140-151

modular switches, 23

monitor session commands, 373

monitoring campus networks

RSPAN, 370-371

SPAN, 368-371

MST (Multiple Spanning Tree), 179-196

best practices, 194-196

configuring, 185-190

Extended System ID, 185

path cost, configuring, 192-193

port priority, configuring, 193

protocol migration, 194

regions, 182-183

STP instances, 183-185

verifying, 190-191

MTU (maximum transmission unit), 51

N

native VLAN (802.1Q), 52

need for EtherChannel, 94-97

neighbor discovery, LLDP, 355-356

network management, SNMP, 336-344

best practices, 339-340

MIB, 337

SNMPv3 configuration example, 340-344

traps, 338

versions, 339

Nexus switches, 22

SDM templates, 364-368

show mac address-table command, 27-28

nondesignated ports, 123

nonedge port links (RSTP), 138

Nonegotiate mode, 53

normal mode (UDLD), 358

NSF (Nonstop Forwarding), 404-405
 NTP (Network Time Protocol), 323-335
 design principles, 329-331
 example, 326-329
 modes, 324-326
 securing, 331-333
 source address, 333
 versions, 333-335

O

object tracking (HSRP), 268-271
 OOB (out-of-band) ports, 28
 optional TLVs (LLDP), 353
 OSI model, Layer 1-2
 overwriting VTP configuration, 87-93

P

packets, rewrites, 28
 PAgP (Port Aggregation Protocol), 98-99
 path cost (MST), configuring, 192-193
 path manipulation, STP, 145-147
 performing extended pings, 105-106
 plain-text authentication, configuring for HSRP, 271
 planes of operation, 28-29
 management plane, 28
 PoE (Power over Ethernet), 70, 360-364
 components, 362
 configuring, 363-364
 power classes, 362-363
 standards, 362
 verifying, 363-364
 point-to-point links (RSTP), 138
 port channels, 7-8
 Port ID field (BPDUs), 124
 port priority (MST), configuring, 193
 port-channel load-balance command, 101
 PortFast, 156-158
 ports
 error conditions (port security), 422-424
 OOB, 28

 PACLs, 424-425
 routed ports, 206, 214-215
 inter-VLAN routing, 214-222
 STP, 129-130
 designated ports, 123
 nondesignated ports, 123
 root ports, 123
 trunk ports, 7
 VLAN ports, 43
 power classes (PoE), 362-363
 preamble field (Ethernet frames), 4
 preemption (HSRP), enabling, 258-263, 273
 promiscuous ports, 454
 Protocol ID field (BPDUs), 124
 protocol migration (MST), 194
 pruning (VTP), 74-75
 PTP (Precision Time Protocol), 336
 PVLANS (private VLANs), 451-458
 across multiple switches, 457-458
 configuring, 454-456
 port types, 453-454
 protected port feature, 458
 verifying, 456-457
 PVST+ (Per-VLAN STP Plus), 130-131

Q-R

QoS (quality of service), 26-27
 RADIUS, 309-310
 configuring, 311-312
 limitations of, 315
 recommendations for STP stability mechanisms, 175-178
 redundancy
 FHRPs, 247
 need for, 248-249
 GLBP, 282-300
 authentication, 295
 comparing with HSRP, 283-284
 configuring, 285-294
 load balancing options, 294-295

- tracking*, 296-300
- virtual forwarder states*, 285
- virtual gateway states*, 285
- HSRP, 250-253
 - authentication, configuring*, 271-272
 - configuring*, 255-263
 - interface tracking*, 266-268
 - load sharing*, 263-265
 - object tracking*, 268-271
 - state transition*, 253-254
 - timers, configuring*, 272-273
 - versions*, 274
- switch supervisors, 401-405
 - supervisor redundancy modes*, 402-405
- VRRP, 274-281
 - comparing with HSRP*, 276
 - configuring*, 276-280
 - millisecond timers*, 275
 - tracking*, 280-281
- regions (MST), 182-183
- resource errors (STP), troubleshooting, 198
- responders (IP SLA), 377-381
- rewrites, 28
- rogue access, 412-415
- root bridge election (STP), 124-126
- Root Bridge ID field (BPDUs), 124
- root bridge, verifying, 144
- Root Guard, 161-163
- Root Path Cost field (BPDUs), 124
- root ports, 123
 - RSTP, 134
- route caching, 30-31
- routed ports, 206, 214-215
 - inter-VLAN routing, 214-222
- router-on-a-stick, 206
- routers
 - broadcast domains, 3
 - inter-VLAN routing with external router, 206-211

- routing protocols, verifying, 229-230
- RSPAN (Remote SPAN), 370-371
 - configuring, 372-374
- RSTP (Rapid Spanning Tree Protocol), 133-134
 - convergence, 150-151
 - links, 138-139
 - modifying behavior, 140-151
 - port roles, 134-135
 - port states, comparing with STP, 135-136
 - STP priority, 143-145
 - topology changes, 136-138

S

- SDM (Switching Database Manager)
 - templates, 364-368
 - selecting, 367
- security
 - in access layer, 13
 - authentication, 305-308
 - 802.1X*, 316-319
 - GLBP*, 295
 - HSRP, configuring*, 271-272
 - RADIUS*, 309-310
 - TACACS+*, 310-311
 - VRRP, configuring*, 279-280
 - VTP*, 75
 - in campus networks, 410-411
 - IP SLA, configuring, 382
 - port security, 419-425
 - configuring*, 420-422
 - PACLs*, 424-425
 - port error conditions*, 422-424
 - vulnerabilities of campus networks, 412-419
 - rogue access*, 412-415
- segmentation, VLANs, 44
- selecting SDM templates, 367
- Sender Bridge ID field (BPDUs), 124
- server mode (VTP), 72
- shared links (RSTP), 138

- show interfaces command, 59
- show ip protocol command, 229-230
- show ip route command, 229-230
- show mac address-table command, 27-28
- show running-config command, 59
- show spanning-tree command, 140-142
- show standby command, 259
- show vlan command, 57-59, 62-63, 81-87
- show vtp status command, 79-81
- SNAP (Subnetwork Access Protocol), 70
- SNMP (Simple Network Management Protocol), 336-344
 - best practices, 339-340
 - MIB, 337
 - SNMPv3 configuration example, 340-344
 - traps, 338
 - versions, 339
- SNTP (Simple Network Time Protocol), 335-336
- SOC (switch-on-chip), 33
- source address (NTP), 333
- source address field (Ethernet frames), 5
- sources (IP SLA), 377
- SPAN (Switch Port Analyzer), 368-371
 - configuring, 371-372
- spanning-tree backbonefast command, 156
- speak state (HSRP), 253
- spoofing attacks, mitigating
 - ARP spoofing, 437-443
 - DHCP spoofing, 430-434
 - IPSG, 435-439
- Src field (802.1Q frames), 51
- SSO (Stateful Switchover), 403-404
- StackWise, 395-397
- standalone WLAN solution, 69-70
- standby routers (HSRP), 252
- standby state (HSRP), 253
- start-of-frame field (Ethernet frames), 4
- state transition, HSRP, 253-254
- store-n-forward mode, 24
- storm control, 425-429
 - configuring, 427-429
- STP (Spanning Tree Protocol), 6-7, 119-123
 - aligning with HSRP, 254-255
 - behavior, modifying, 140-151
 - BPDU Filter, 159-161
 - BPDU Guard, 158-159
 - BPDUs, 121, 124
 - inferior BPDUs*, 155
 - Cisco STP Toolkit
 - BackboneFast*, 154-156
 - UplinkFast*, 153-154
 - designated ports, 123
 - election process*, 128-129
 - FlexLinks, 171-175
 - Loop Guard, 164-166
 - leveraging with UDLD*, 360
 - MST, 179-196
 - best practices*, 194-196
 - configuring*, 185-190
 - Extended System ID*, 185
 - path cost, configuring*, 192-193
 - port priority, configuring*, 193
 - protocol migration*, 194
 - regions*, 182-183
 - STP instances*, 183-185
 - verifying*, 190-191
 - need for, 120-121
 - nondesignated ports, 123
 - path manipulation, 145-147
 - PortFast, 156-158
 - ports, 129-130
 - priority, modifying, 143-145
 - root bridge
 - election process*, 124-126
 - verifying*, 144
 - Root Guard, 161-163
 - root ports, 123
 - election process*, 126-127
 - stability mechanisms recommendations, 175-178
 - standards, 121-122, 130-131
 - timers, 148-151

- topology changes, 131-133
- topology events debugging, 148-149
- troubleshooting
 - duplex mismatches*, 196-197
 - frame corruption*, 197-198
 - PortFast configuration errors*, 198
 - resource errors*, 198
 - unidirectional link failure*, 197
- UDLD, 166-171
- structure of campus networks, hierarchical network design, 10-19**
 - access layer, 12-13
 - core layer, 14-17
 - distribution layer, 13-14
 - FHRP, 14
 - versus flat network design, 10-11
- subinterfaces, 207**
- subset advertisements (VTP), 77**
- summary advertisements (VTP), 77**
- supervisor redundancy modes, 402-405**
 - NSF, 404-405
 - SSO, 403-404
- SVI (switch virtual interface), 206**
 - advantages of, 214
 - configuring, 221-222
 - inter-VLAN routing, 212-214
 - routed ports, 214-215
- switch ports**
 - assigning to VLANs, 56-57
 - end-to-end VLANs, 44-45
- switch spoofing, 444-446**
- switches. *See also* switching methods**
 - and bridges, 2-3
 - broadcast domains, 3
 - CAM table, 5
 - Catalyst switches
 - Catalyst 3850-X series*, 23
 - Catalyst 6500 switches*, 23
 - Catalyst 6800-X series*, 23
 - Cisco switches, 22
 - trunking modes*, 53-54
 - collision domains, 24
 - CRC, 24
 - end-to-end VLANs, 44-45
 - EtherChannel
 - configuring*, 102-108
 - LACP*, 97-98
 - load balancing*, 100-102
 - troubleshooting*, 108-109
 - features, 3
 - fixed configuration switches, 23
 - full-duplex mode, 24
 - functions of, 5-6
 - half-duplex mode, 24
 - and hubs, 2
 - Layer 2 switching
 - ACLs*, 26
 - ingress queues*, 25
 - MAC addresses*, 24-25
 - MAC table*, 26
 - QoS*, 26
 - MAC addresses, displaying MAC table information, 60-61
 - MLS, 8, 26-27
 - centralized switching*, 33
 - distributed switching*, 33
 - inter-VLAN routing*, 217-220
 - planes of operation*, 28-29
 - modular switches, 23
 - redundant switch supervisors, 401-405
 - supervisor redundancy modes*, 402-405
 - route caching, inter-VLAN routing, 43
 - security, 410-411
 - SOC, 33
 - store-n-forward mode, 24
 - STP, 6-7, 119-120
 - BPDUs*, 124
 - designated ports*, 123
 - election process*, 128-129
 - nondesignated ports*, 123
 - path manipulation*, 145-147
 - ports*, 129-130
 - root bridge election*, 124-126

- root port election, 126-127*
- root ports, 123*
- standards, 121-122*
- timers, 148-151*
- topology changes, 131-133*
- topology events debugging, 148-149*
- TCAM, 26
- trunking, 7
 - 802.1Q trunking, 49-52*
- VLANs, 6, 42-48
 - best practices, 65-66*
 - configuring, 61-64*
 - creating in global configuration mode, 55-56*
 - deleting, 56*
 - ISL, 49*
 - local VLANs, 45-46*
 - port channels, 7-8*
 - ports, 43*
 - PVLANs, 451-458*
 - segmentation, 44*
 - verifying configuration, 57-61*
 - VTP, 41
- vulnerabilities, 415-417
- wireless network support, 69-70
- switching methods
 - route caching, 30-31
 - topology-based switching, 31-33
- switchport host command, 56-57
- symmetric mode (NTP), 370
- system clock
 - manual configuration, 320-322
 - NTP, 323-335
 - design principles, 329-331*
 - example, 326-329*
 - modes, 324-326*
 - securing, 331-333*
 - source address, 333*
 - versions, 333-335*
- PTP, 336
- SNTP, 335-336

T

- TACACS+, 310-311
 - configuring, 312-313
 - limitations of, 315
- Tag field (802.1Q frames), 51
- tagging the frame, 7
- TCAM (ternary content-addressable memory), 26
 - SDM templates, 364-368
 - selecting, 367*
 - system resource configuration, 367-368
- time
 - accuracy, need for, 320
 - NTP, 323-335
 - design principles, 329-331*
 - example, 326-329*
 - modes, 324-326*
 - securing, 331-333*
 - source address, 333*
 - versions, 333-335*
 - PTP, 336
 - SNTP, 335-336
 - system clock, manual configuration, 320-322
- time stamps (IP SLA), 381-382
- timers
 - HSRP, tuning, 272-273
 - millisecond timers, VRRP, 275
 - STP, 148-151
- TLVs (LLDP), 353
- topology changes, RSTP, 136-138
- topology-based switching, 31-33
 - load balancing, 32-33
- traceroute command, 210
- tracking
 - interfaces, 266-268
 - VRRP, enabling, 280-281
- traffic storms, storm control, 425-429
 - configuring, 427-429

transparent bridges, 2-3
transparent mode (VTP), 72
traps (SNMP), 338
troubleshooting
 EtherChannel, 108-109
 inter-VLAN routing, 222-225
 STP
 duplex mismatches, 196-197
 frame corruption, 197-198
 PortFast configuration errors, 198
 resource errors, 198
 unidirectional link failure, 197
trunk links, 50
Trunk mode, 53
trunk ports, DTP, 53-54
trunking, 7, 49-54
 802.1Q trunking, 43-44, 49-52
 architectural advantages over ISL,
 51
 native VLAN, 52
 best practices, 65-66
 configuring, 64-65
 DTP, 53-54
 IEEE 802.1Q trunking, VLAN ranges, 54
 switch spoofing, 444-446
 trunk links, 50
 VLAN hopping, 446-448
 VTP
 advertisements, 75-77
 authentication, 75
 best practices, 93
 configuring, 78-87
 messages, 77
 modes, 71-73
 overwriting configuration, 87-93
 pruning, 74-75
 versions, 73-74
tuning HSRP, 255-263

U

UDLD (Unidirectional Link Detection),
166-171, 357-360
 configuring, 358-360
 default behaviors, 359
 leveraging with STP Loop Guard, 360
UDP (User Datagram Protocol) jitter, IP
 SLA example, 383-384
unidirectional link failure, troubleshooting,
197
unified network services, 3
UplinkFast, 153-154
UTC (Coordinated Universal Time), 320

V

VACLs (VLAN ACLs), 448-451
verifying
 GLBP, 285-294
 MST, 190-191
 PoE, 363-364
 PVLANS, 456-457
 root bridge, 144
 routing protocols, 229-230
 StackWise, 396-397
 VLAN configuration, 57-61
 VSS, 399-401
Version field (BPDUs), 124
versions
 of HSRP, 274
 of NTP, 333-335
 of SNMP, 339
 of VTP, 73-74
VID (VLAN ID), 49
viewing
 Layer 2 forwarding table, 27-28
 MAC address table information, 60-61
virtual forwarder states (GLBP), 285
virtual gateway states (GLBP), 285

- virtual routers (HSRP), 251
 - VLAN hopping, protecting against, 446-448
 - VLANs, 42-48
 - best practices, 65-66
 - broadcast domains, 43
 - configuring, 61-64
 - creating in global configuration mode, 55-56
 - deleting, 56
 - end-to-end VLANs, 44-45
 - inter-VLAN routing, 43, 204-206
 - MLS, 217-220
 - router-on-a-stick, 206
 - troubleshooting, 222-225
 - using external router, 206-211
 - using routed ports, 214-222
 - using SVI, 212-222
 - ISL, 49
 - local VLANs, 45-46
 - mapping to hierarchical network, 47-48
 - port channels, 7-8
 - ports, 43
 - PVLANs, 451-458
 - across multiple switches, 457-458
 - configuring, 454-456
 - port types, 453-454
 - protected port feature, 458
 - verifying, 456-457
 - segmentation, 44
 - STP, 6-7
 - switch ports, assigning, 56-57
 - trunking, 7, 49-54
 - 802.1Q trunking, 43-44
 - best practices, 65-66
 - configuring, 64-65
 - DTP, 53-54
 - switch spoofing, 444-446
 - trunk links, 50
 - VLAN hopping, 446-448
 - VACLs, 448-451
 - verifying configuration, 57-61
 - voice VLAN, 67-69
 - VTP, 41
 - advertisements, 75-77
 - authentication, 75
 - best practices, 93
 - configuring, 78-87
 - messages, 77
 - modes, 71-73
 - overwriting configuration, 87-93
 - pruning, 74-75
 - voice VLAN, 67-69
 - vPC (Virtual Port Channel), 96
 - VRRP (Virtual Router Redundancy Protocol), 14, 274-281
 - comparing with HSRP, 276
 - configuring, 276-280
 - tracking, 280-281
 - VSS (Virtual Switching System), 96, 397-401
 - VTP (VLAN Trunking Protocol), 41, 70-71
 - advertisements, 75-77
 - authentication, 75
 - best practices, 93
 - configuring, 78-87
 - messages, 77
 - modes, 71-73
 - overwriting configuration, 87-93
 - pruning, 74-75
 - versions, 73-74
 - vulnerabilities of campus networks, 412-419
 - MAC flooding attacks, 417-419
 - rogue access, 412-415
-
- ## W
-
- weighting, GLBP, 298-300
 - wireless networks, switches, 69-70
 - WLANs (wireless LANs), Cisco solutions, 69-70
-
- ## X-Y-Z
-
- XOR operation, EtherChannel load balancing, 101